



The Persuasion Pipeline

How Commercial AI Targets, Shapes, and Sells
to a Generation of California's Children and Teens

Acknowledgements

The [Center for Digital Democracy's](#) project on the role of commercial AI in reshaping advertising and marketing aimed at young people is funded, in part, by the [Rose Foundation for Communities and the Environment](#). CDD thanks the foundation for its generous support.

The Center for Digital Democracy's mission is to ensure that digital technologies serve and strengthen democratic values, institutions, and processes. CDD strives to safeguard privacy and civil and human rights, as well as to advance equity, fairness, and community.

Katharina Kopp, Ph.D., a data privacy and data justice expert and advocate, is Deputy Director and Director for Policy at the Center for Digital Democracy. Dr. Kopp completed her Ph.D. and M.A. in communications and public policy at the University of Pennsylvania's Annenberg School for Communication and graduated from Queen Mary, University of London.

Kathryn Montgomery, Ph.D., and Jeff Chester have profoundly shaped the author's thinking and the genesis of this report. Dr. Montgomery provided editorial review, and Jeff Chester advised on strategy and countless references.

Layout and design by Alec Meltzer, meltzerdesign.net.

A Note on Methods and the Use of AI

In preparing this report, the author used Anthropic's Claude (Opus 4 series, 2025–2026) as a research and drafting assistant: to locate and summarize candidate sources for the author's review, to draft and revise prose from the author's direction and source materials, to translate the executive summary into Spanish, and to format the document.

The research questions, analytical framework, arguments, and conclusions are the author's. Where the AI drafted prose, it did so from the author's draft, direction, structure, and source materials, and every passage was reviewed, revised as needed, and adopted by the author. Every source cited in this report was located and read in its original form by the author; AI-surfaced sources that could not be independently located and verified were not used.

The Spanish translation of the executive summary was produced with Claude (Opus 4 series) and reviewed and edited by María Isabel Valero Morales, a native speaker, who compared it against the English source for accuracy of translation. She did not independently verify the factual content.

CONTENTS

Executive Summary	4
Resumen ejecutivo.....	12
Introduction	22
Layer 1: Data Collection and Tracking	28
Layer 2: Profiling and Segmentation	34
Layer 3: Targeting and Optimization – From Manual Control to Algorithmic Steering	38
Layer 4: Creative and Content Generation – How AI Reshapes Persuasive Media and Commercial Influence.....	42
Layer 5: Interaction and Experience – Targeted Conversational Influence.....	48
Layer 6: Measurement, Feedback, and the Closing of the Loop	60
Ten Cross-Cutting Trends	64
Harms and Risks to Children and Teens	68
Policy	78
Beyond Consumer Protection: A Conclusion	90
Glossary	92
Appendix A1-Table:.....	96
Appendix A1-Key Findings:.....	105
Appendix A2-Table:.....	112
Appendix A2-Key Findings:.....	120

EXECUTIVE SUMMARY

One in three internet users globally is a child. In the United States, young people are now among the earliest and most intensive adopters of artificial intelligence: 64 percent of U.S. teens have used AI chatbots, with 30 percent using them daily. Adoption rates are highest among Black and Hispanic teens, roughly seven in ten of whom report using chatbots compared with 58 percent of White teens. Yet only about half of parents are aware their teens use these tools at all. Digital media, moreover, are no longer simply one part of childhood among many. For most young people they are now the primary environment in which learning, friendship, play, identity, and the first experiments of independence take place, and that role grows larger every year.

The technology, too, is also changing. Most uses of AI today refer to machine learning – algorithms that simulate intelligent behavior by making inferences and predictions from data, improving as more data are fed into them. Machine learning is not new: it has been under development for roughly eight decades, with significant breakthroughs beginning in the 1990s. What *is* new, and what justifies focused attention now, is *generative AI* as both a commercial tool and a consumer-facing medium. Unlike machine learning that only classifies or predicts, generative AI also *creates*: personalized advertising, conversational commercial interfaces, and continuous optimization of persuasion itself, at a scale prior systems could never reach.

This report examines the full range of AI reshaping advertising and marketing aimed at young people, with particular attention to generative AI as the newest and most consequential development. While much public attention has rightly focused on familiar harms – privacy violations, addictive design, cyberbullying, exposure to harmful content, discriminatory treatment, and the dangers of AI companion chatbots – far less attention has been paid to a deeper change: the ways in which generative AI and other commercial AI tools are being used to influence users, producing a structural transformation in young people’s online lives that

powerful companies relentlessly advance through their advertising-supported business model.

This report documents how that system works, layer by layer, and what it means for children and teens. More than simply a catalog of individual harms, the report focuses instead on the system *behind* those harms: an AI-driven apparatus that observes children, infers what moves them, generates messages to match, and talks back – often without any of it being recognizable as commercial transactions. Our decision to link this report to California was based on three factors:


The state itself is home to approximately 8.4 million children – nearly half of whom are Hispanic/Latino (and nearly half of those from immigrant families); California is also home to many of the companies building these systems; and the policy decisions California makes about AI-driven commercial systems will set the terms for every other state and will shape federal debate.

This report examines AI-driven commercial influence on children and teens through a *six-layer framework*: data collection, profiling, targeting, creative generation, interaction, and measurement. We separate the system into layers not because they operate independently, but to make the single, integrated pipeline more apparent.

In practice the layers form a continuous loop: each one feeds the next, and what the system learns at the end – from what caught a child’s attention to what ultimately prompted a purchase – flows back to the beginning to sharpen the next cycle. This system draws on two kinds of inputs. One is personal data, collected through the platforms and devices young people use – the raw material of commercial surveillance. The other, on which this report focuses and which has drawn far less attention, is the capacity of AI systems to analyze *non-personal* data – patterns across millions of users, signals stripped of identifiers, content environments rather than individuals – and to act on those patterns in ways that can create harm *without* collecting personal data from the people affected.

The report covers the six layers of this system in turn:

- **Layer 1: Data Collection and Tracking**, in which AI systems gather and infer signals from user behavior, devices, content consumption, location, and context, often across platforms and over time. AI tracking runs in real time, with signals fed back continuously. AI also enables probabilistic tracking: instead of relying on exact identifiers like a login or device ID, systems *infer* whom a user is from fragmentary signals and stitch their activity together across devices, apps, and contexts, making tracking more comprehensive and harder to escape than the methods it replaces. Conversational AI goes further, actively soliciting the intimate disclosures that earlier advertising could only guess at.
- **Layer 2: Profiling and Segmentation**, in which algorithms build predictive models of individuals and groups, assigning attributes, preferences, likely behavior, and how open someone is to persuasion. The commercial value of data now lies less in *what* is collected than in what is *inferred* from it, including inferences drawn not from a child’s own data but from patterns learned across millions of other people. Some of these inferences are stored as a durable profile, the kind of record privacy law was built to govern; others are *real-time* inferences, made in the moment, acted on, and discarded without becoming a record – a moving target that traditional privacy protections, built around stored records, are not designed to address.
- **Layer 3: Targeting and Optimization**, in which AI systems decide the audiences to be reached, through which channels, at what time, and with what frequency, continuously adjusting delivery based on real-time performance signals. Marketers increasingly set only a single goal – a sale, a click, a download – and the optimization system decides who is reached, adjusting delivery continuously as it learns what performs best. The same optimization that pushes ads toward those most likely to respond produces both discriminatory delivery and child-audience capture: a platform need not know a user’s age, or specifically target it, to steer age-skewed messaging to teens.
- **Layer 4: Creative and Content Generation**, in which generative and adaptive AI tools produce, assemble, and modify ads, narratives, and entire branded environments tailored to specific segments or individuals. Advertising creative has shifted from campaigns authored by people, to content computed by AI, per-person, per-moment, at scale and tuned in real time to what performs best, often matching or beating traditional advertising. It arrives with the slang, humor, and look of peer culture, and, because the underlying models carry covert stereotypes, bias is built into what an ad says and how it looks, falling hardest on youth of color.



- **Layer 5: Interaction and Experience – Targeted Conversational Influence**, in which young people encounter AI-mediated marketing through social media feeds, recommendation systems, chatbots, voice assistants, games, and immersive environments, all of which blur the line between content and commerce. This is where generative and conversational AI change the experience most. Agentic AI – AI that acts on a user’s behalf, not just answers their questions – has opened new channels for commercial influence. AI assistants embedded in the platforms young people use most are now tied to advertising-supported business models, and their persistent memory and real-time responsiveness build the kind of engagement and trust that makes commercial influence hard to recognize as commercial; an ad delivered inside a trusted conversation is unlikely to be seen as an ad at all.
- **Layer 6: Measurement, Feedback, and the Closing of the Loop**, in which AI evaluates outcomes and feeds them back into the system. AI has made measurement more comprehensive and more probabilistic, tracking outcomes at the level of individual impressions, and predicting effects that older measurement could only estimate in the aggregate. The change is how tightly measurement is now embedded in the loop: the system encodes what worked and feeds it straight back to sharpen the next cycle, in real time and at scale. In the process, AI is effectively pushing the industry from an attention economy toward an “intention economy,” reading statements of want and intent the moment a young person types them, and embedding the path to purchase directly into the conversation.

Viewed together, these six layers describe an emerging commercial influence system in which the boundaries among content, conversation, advertising, and transaction are dissolving, as the potential for powerful manipulation increases. The regulatory frameworks now in force were built around *personal* data, *individual* identifiability, notice-and-choice *consent* (where users are informed of data practices and given the option to accept or decline), and discrete, recognizable advertisements. They were not built for a system in which AI *infers*, *optimizes*, *generates*, and *converses*, often without ever needing personal data from the user

it is targeting. Some elements of the system fall within current legal reach; others sit outside it or fit only partially. The harms documented in this report, and the policy interventions this report recommends, proceed from that diagnosis. But none of this is inevitable. These are *design* and *business-model* choices, and the harms they produce can be affected, in turn, by *policy* choices – which is the focus of the closing sections of this report.

Viewed from a broader perspective, commercial influence on children is increasingly delivered not as a series of discrete advertisements but through a continuous, adaptive, AI-managed environment of commercial persuasion. As noted, the six layers examined in this report do not operate separately. They form an integrated system in which each function reinforces the others. In order to illuminate the gravity of this shift, this report also explores ten cross-cutting trends. All are observable now, all are accelerating, and all will have profound effects on our children, our culture, and our democracy:

- **Trend 1. Interactive Generative Media is Emerging as a Distinct Commercial Form**
Conversational AI, intelligent agents, synthetic influencers, and generative media systems are giving rise to an entirely new category of commercial communication. Seven of the nine companies surveyed in Appendix A – Meta, Google, OpenAI, Microsoft, Amazon, Snap, and xAI – are confirmed or likely to be serving advertising inside AI chat, using chat interactions to target advertising elsewhere, or both. Three features distinguish this new form – *Interactive Generative Media* – from all prior forms of commercial communication. First, real-time feedback loops: these systems continuously monitor a user’s responses – hesitations, questions, emotional signals – and immediately adjust their persuasive approach. Second, relational continuity: memory systems store what worked on a given user across sessions, building an increasingly refined model of the individual over time. Third, adaptive personalization at the level of persuasive strategy: the system personalizes not just *what* product to recommend but *how* to recommend it – whether to lead with scarcity, social proof, or emotional appeal, calibrated to the individual’s demonstrated susceptibilities. For young users, these systems appear not as ads but as helpers, companions, and peers.



Yuri Malishenko via Unsplash

- **Trend 2. Signal Extraction is Expanding into New Domains and Modalities**

Continuous data collection from user behavior is not new. What is changing is the *scope* of what counts as a signal. Conversational AI has opened entirely new categories of data: the questions people ask, the uncertainties they express, the emotional states they disclose in dialogue, and the explicit statements of intent they make to systems they regard as confidential. As these interfaces become primary modes of interaction, they are being captured as spaces for signal extraction in much the same way social media was a decade ago. The shift away from third-party cookies toward tracking that happens on companies' own servers, where a user's identity is reassembled from probability rather than a single stored tag, has made this infrastructure harder for policymakers, users, and researchers to observe – even as it captures more intimate signals than ever before.

- **Trend 3. Profiling is Shifting from Stored Identity to Real-Time Behavioral Vulnerability**

AI-driven profiling is moving along two connected axes. One is what is modeled. Systems are moving beyond static demographics – *who* a user is – toward dynamic models that continuously update to capture *how* a user is likely to respond to specific persuasive approaches. When Meta's Advantage+ or Google's Performance Max selects which users to show which ads, and in what format and sequence, it is modeling which individuals are most likely to be persuaded by which approaches – a capability that functions as susceptibility assessment regardless of what the platform calls it. The other axis is where the modeling lives. Traditional profiling builds a durable record – data stored against an identified person, the kind of profile existing privacy law was written to govern. A growing share of commercial AI now works differently: it classifies a user in the moment,

from the immediate context of an interaction, to shape what happens next, without depending on anything lasting about that individual. The judgment about how to move the user is made and acted on in real time, and need not be written into a durable record at all.

- **Trend 4. AI Systems are Learning to Persuade More Effectively**

What distinguishes the current moment is not automation at scale but *what* is being optimized. Closed feedback loops continuously refine which persuasive tactics, emotional appeals, and cognitive approaches work on which users. These same ad platforms delegate targeting, creative, and optimization decisions to AI that learns from billions of interactions. What emerges is a system that gets better at moving people over time, with limited transparency about the techniques employed and no external oversight of the persuasive strategies it discovers. For young users, this means the methods acting on them are tuned by systems that no one outside the company can inspect.

- **Trend 5. AI Systems are Beginning to Act Before Children Decide What They Want**

Marketing systems are moving from *responding* to expressed interests toward *predicting* needs, desires, and vulnerabilities before users articulate them. Agentic AI systems are beginning to act *on behalf* of users – filtering options, making recommendations, completing transactions – often guided by the commercial interests of the platforms that deploy them. A growing Generative Engine Optimization industry ensures brands are cited and recommended by AI systems when users ask purchasing-related questions. This shift is particularly consequential for children and teens, who may not recognize that the system has intervened before they had the opportunity to reflect.

- **Trend 6. The Boundary Between Commercial and Authentic Content is Dissolving**

Generative AI enables the production of ads, narratives, characters, and interactive experiences that are stylistically indistinguishable from organic content, peer expression, and entertainment. AI-generated creative adopts the language, humor,

visual grammar, and cultural references of specific youth audiences, adapting at the speed of trend cycles. Synthetic influencers, digital twins, and dynamically personalized ad variants contribute to an environment in which the markers that once signaled “this is an ad” are being systematically erased. For young users whose advertising literacy is still developing, this makes recognition and critical evaluation of commercial influence substantially harder.

- **Trend 7. Marketing Functions are Converging into Integrated, AI-Optimized Pipelines**

Data collection, profiling, targeting, creative production, interaction, measurement, and transaction are collapsing into unified, AI-managed pipelines. Meta, Google, Amazon, and TikTok each offer end-to-end systems in which advertisers set business objectives and AI handles virtually every other decision. Commerce is being embedded directly into conversational interfaces – OpenAI’s Instant Checkout, Google’s Universal Commerce Protocol, and platform storefronts on TikTok, Instagram, YouTube, and Roblox – so the path from influence to purchase occurs without the user ever leaving the environment where the initial persuasion took place. For young people who move fluidly across apps, platforms, and devices, this creates continuous commercial exposure with no meaningful boundary between contexts.

- **Trend 8. Power Asymmetry Between AI Systems and Users is Deepening**

AI systems accumulate detailed knowledge about individual users across interactions, learning what persuasive strategies are most effective for each person. The user knows nothing about the system’s commercial objectives, the data informing its approach, or the tactics it employs. This asymmetry has no precedent in human-to-human commercial encounters and is particularly acute for young users, who are less likely to question the motives of systems that present themselves as helpful or friendly. Advertisers are not passive participants; they actively choose to use these systems and benefit from the asymmetric advantages they create.



- **Trend 9. Disparate Commercial Exposure is Compounding Existing Inequities**

These trends do not affect all young people equally. Youth of color and lower-income families face disproportionate exposure at every layer: higher rates of data collection driven by more intensive platform use; profiling systems that sort users along racial and socioeconomic lines through behavioral proxies; optimization that steers unhealthy product advertising toward Black and Hispanic audiences; and AI-generated creative that embeds raciolinguistic stereotypes, and measurement systems that compound these biases with each cycle. AI does not create these disparities, but the integrated system documented in this report amplifies and formalizes them at a scale and speed that manual marketing practices could never achieve. Age-gating at the platform access level does not prevent commercial targeting of under-13 children when commercial content flows through recommendation algorithms and creator content rather than paid targeted advertising. The AI-driven commercial pipeline reaches children regardless of the policy constraints applied at any single layer.

- **Trend 10. Commercial Systems Treat the Vulnerabilities of Childhood as Resources to Exploit**

The systems documented in this report do not simply reach children – they *exploit* specific characteristics of childhood and adolescence. Adolescent impulsivity and reward sensitivity, the use of AI companions for emotional support and intimate interaction, and young consumers’ developing self-brand connections can all become commercially legible signals in AI-optimized advertising systems: engagement, disclosure, emotional response, identity fit, and predicted conversion. As AI-driven commercial systems reach full deployment, their capacity to target the developmental stage of childhood will compound. Children who grow up in environments of total commercial mediation do not merely encounter more advertising, they *develop* within it.

Taken together, these trends converge on two consequences. The first is the erosion of autonomous choice: when signal extraction is continuous, profiling models susceptibility, influence arrives before children articulate what they want, persuasion is optimized through feedback loops, and the line between commercial and organic content has dissolved, young people are no longer making decisions in an open marketplace but moving through environments designed to *steer* them. The second is the transfer of decision-making to a small number of companies whose systems increasingly determine what young people see, want, and buy, with minimal human oversight.

The Need for New Policies

The six layers and ten trends in this report describe an integrated AI commercial persuasion pipeline that existing law is not well suited to address, because of a mismatch at the foundation. Privacy and consumer-protection law is built principally around the personal data a company holds about an identifiable person, and where these advertising practices rest on such data, that law applies, though first-party exceptions can exempt first-party advertising even then. But many of the AI-driven practices described here operate on a different basis: their central capability is inferential – the capacity to single out, classify,

predict, influence, and differentially treat an individual in real time, drawing on models trained on millions of other users, often without ever writing a durable record about the person being acted on. AI does not invent commercial surveillance; it accelerates a trajectory already underway and adds genuinely new capabilities, such as real-time conversational influence, persuadability optimization, and dynamic creative loops. The United States still urgently needs comprehensive federal privacy legislation to provide minimum protections for all Americans. But law anchored to the holding of identifiable data is built to govern the *record* rather than the underlying *capability*, and that gap points toward a response: limiting what AI systems can do to minors, not only what data they hold, can meaningfully reduce harm even without comprehensive reform of the whole pipeline.

The Wave of AI Lawmaking has Largely Bypassed Commercial Persuasion

By early 2026, more than three dozen states had adopted over one hundred AI laws. Their focus is important and addresses real issues – companion-chatbot safety, disclosure of synthetic performers, high-risk automated decisions in employment, housing, and lending, government AI use, and frontier-model transparency. But none takes the AI-driven advertising and marketing pipeline documented across the six layers as its object. The volume of AI lawmaking has surged while this category of harm has gone largely unaddressed.

The gaps across the legal landscape are not separate drafting oversights; they share a single root. Most existing law is oriented toward the data a company holds about an identifiable person, while the systems described here operate as an integrated, self-reinforcing loop that runs on capability. A handful of proposals have begun to target that capability directly, but enacted law has largely not. Generative AI does not so much place a discrete advertisement before a child as generate the commercial environment the child moves through: a chatbot builds a relationship with a teenager and surfaces a product in the conversation; a recommendation system routes a tween influencer's branded line to the children who follow her; a shopping assistant writes a paid brand

into its reply. That loop runs on personal and non-personal data alike – aggregate behavioral patterns, cohort models trained on millions of others, de-identified signals, the outputs of other people's interactions – data that AI has only recently made this powerful, and it is optimized continuously for engagement, conversion, and attention, not for a child's well-being. Three features of how it works defeat the mechanisms current and proposed laws rely on. First, influence operates through real-time inference, and that inference is largely unregulated when used for advertising. Second, the operative classification is about a group, not a person, and it is formed and acted on without ever producing the individual record that a knowledge standard assumes. Third, consent and data-minimization frameworks break down because the model's power does not come from the data of the child being acted on.

Conclusion

The harms documented in this report matter because of where they occur: not at the edges of children's lives but at the center of them. And they do not fit neatly into the category of consumer protection. They are larger than that. They operate at three levels at once – on the individual child, on the relationships and choices children share, and on the democratic culture they are growing into – and all three flow from a single underlying capacity these systems have acquired:

- **The Individual Child**

A child's conversations with an AI companion, the questions typed into a chatbot, and the products searched and considered can all become commercial data events. The infrastructure described in this report turns private expression, emotional disclosure, and ordinary curiosity into structured commercial signals, surveilling children at a depth, intimacy, and continuity no prior advertising medium could achieve. A child who grows up inside environments of total commercial mediation does not merely encounter more advertising; the habits, preferences, and self-perceptions formed in those years take shape under systems whose objective is commercial conversion rather than healthy development.

- **Children Together**

The damage does not stop at the individual. When signal extraction is continuous, when profiling models a child's susceptibility, when influence arrives before a child has articulated a want, and when the line between commercial and authentic content is dissolving, the conditions for meaningful independent choice are eroding for all of them. Young people are not navigating an open marketplace; they are moving through environments designed to steer them, and the system does not so much persuade as maneuver. That steering is not distributed evenly. Youth of color face more intensive data collection, ad delivery that pushes unhealthy products toward Black and Hispanic audiences, AI-generated creative that embeds raciolinguistic stereotypes, and lower-quality interactions in conversational systems. These disparities are not incidental; they are structural features of systems trained on data that reflects historical inequality and compounded with every optimization cycle, so that the burdens of the commercial environment fall hardest on the children with the least protection.

- **Democratic Life**

At the widest level, the concern is what these systems do to the culture young people are inheriting. The capacity at issue is the ability to single out, classify, predict, influence, and differentially treat people at scale – exercised here by a small number of firms, and exercised whether or not any identifying record is kept. When that capacity governs what young people see, want, and buy, the market is distorted: firms compete not on the merits of what they make but on the sophistication of the systems they use to shape demand. And the spaces where young people once developed habits of independent thought, civic participation, and shared judgment are being progressively absorbed into commercial optimization. Legal scholars have begun to argue that the defining features of these systems are corrosive to the civic institutions democratic life depends on.

These three levels are not separate problems. They are the same capability seen at different scales, and at every scale the harm accumulates. A single interaction rarely registers as damage; a childhood spent inside these systems does. Preferences shaped session by session, inequalities deepened with every optimization cycle, and civic spaces absorbed year by year are harms that build rather than announce themselves, which is part of why episodic, incident-driven regulation continues to miss them.

The surveillance of one child, the distorted choices of many, and the slow commercialization of civic life all flow from the inferential capacity that commercial AI infrastructure has acquired – the capacity to read a person, sort them, predict how they will respond, and treat them accordingly. This is why the report has argued for regulating that capability rather than chasing its symptoms one at a time. A framework anchored to what data a company stores tends to arrive after the fact, because the mechanism that produces the harm can operate without producing a durable record. Capability-based regulation aims at the mechanism itself.

It is worth being clear about what the alternative looks like, because the goal is not only to stop harmful practices but to make room for something better. Children and young people need places to grow up that are not organized around surveillance, advertising, and behavioral optimization. That means allocating real resources to build non-commercial digital spaces for learning, play, creativity, and civic life; requiring youth-facing technologies to optimize for well-being, autonomy, privacy, and democratic participation; supporting business models that do not depend on extracting attention; and investing in real-world spaces where children and teens can safely gather, explore, create, and gain independence outside the home. Industry has a role in this future that it has so far largely declined: to design for the people it serves rather than for the duration of their attention, and to treat children's developmental stage as a reason for restraint rather than a target. Policymakers have the role only they can play: to set the boundaries within which these systems operate, so that protecting children is not left to families negotiating alone against infrastructure built to outpace them.

RESUMEN EJECUTIVO

Versión en español

Uno de cada tres usuarios de internet en el mundo es una niña o un niño. En Estados Unidos, actualmente, los jóvenes son quienes adoptan la inteligencia artificial (IA) de forma más temprana e intensiva: el 64 por ciento de las y los adolescentes estadounidenses ha usado chatbots de IA, y un 30 por ciento los usa a diario. Las tasas de adopción son más altas entre los adolescentes negros e hispanos, alrededor de siete de cada diez declaran usar chatbots, frente al 58 por ciento de los adolescentes blancos. A pesar de esto, solo cerca de la mitad de las madres y los padres saben que sus hijos adolescentes usan estas herramientas. Además, los medios digitales, ya no son simplemente una parte de la infancia. Para la mayoría de los jóvenes, ahora son el entorno principal en el que ocurre el aprendizaje, la amistad, el juego, la identidad y los primeros experimentos de independencia, y este papel crece cada año más.

La tecnología también está cambiando. La mayoría de los usos de la IA hoy se refieren al aprendizaje automático: algoritmos que simulan un comportamiento inteligente al hacer inferencias y predicciones a partir de datos, y que mejoran a medida que se les suministran más datos. El aprendizaje automático no es nuevo: lleva en desarrollo desde hace cerca de ocho décadas, con avances importantes que comenzaron en la década de 1990. Lo que *sí* es nuevo, y lo que justifica que se le preste atención específica ahora, es la *IA generativa*, tanto como herramienta comercial como medio dirigido al consumidor. A diferencia del aprendizaje automático, que solo clasifica o predice, la IA generativa también *crea*: publicidad personalizada, interfaces comerciales conversacionales y la optimización continua de la persuasión misma, a una escala que los sistemas anteriores nunca podrían alcanzar.

Este informe examina toda la gama de IA que está reconfigurando la publicidad y el marketing dirigidos a las y los jóvenes, con atención particular a la IA generativa como el desarrollo más reciente y de mayores consecuencias. Si bien gran parte de

la atención pública se ha centrado, con razón, en daños conocidos -violaciones de la privacidad, diseño adictivo, ciberacoso, exposición a contenido nocivo, trato discriminatorio y los peligros de los chatbots de compañía de IA-, se ha prestado mucha menos atención a un cambio más profundo: las maneras en que la IA generativa y otras herramientas comerciales de IA se están usando para influir en los usuarios, lo que produce una transformación estructural en la vida digital de las y los jóvenes, impulsado sin cesar por empresas poderosas a través de un sistema de modelo de negocio sostenido por la publicidad.

Este informe documenta cómo funciona ese sistema, capa por capa, y qué significa para niñas, niños y adolescentes. Más que un simple catálogo de daños individuales, el informe se centra en el sistema que está *detrás* de esos daños: un aparato impulsado por IA que observa a las niñas y los niños, infiere qué los moviliza, genera mensajes a la medida y responde, a menudo sin que nada de ello sea reconocible como transacción comercial. Nuestra decisión de vincular este informe con California se basó en tres factores: El propio



Tra Tran via Unsplash

estado alberga a aproximadamente 8.4 millones de niñas y niños -casi la mitad de los cuales son hispanos/latinos (y casi la mitad de estos provienen de familias inmigrantes)-; California es también la sede de muchas de las empresas que construyen estos sistemas; y las decisiones de política pública que California tome sobre los sistemas comerciales impulsados por IA fijarán las condiciones para todos los demás estados y darán forma al debate federal.

Este informe examina la influencia comercial impulsada por IA sobre niñas, niños y adolescentes a través de un *marco de seis capas*: recopilación de datos, perfilado, focalización, generación creativa, interacción y medición. Separamos el sistema en capas no porque operen de forma independiente, sino para que el funcionamiento integrado de la cadena sea más evidente. En la práctica, las capas forman un bucle continuo: cada una alimenta a la siguiente, y lo que el sistema aprende al final -desde lo que captó la atención de una niña o un niño hasta lo que finalmente motivó una compra- regresa al inicio para afinar el ciclo siguiente. Este sistema se nutre de dos tipos de insumos. Uno de ellos son los datos personales, recopilados a través de las plataformas y los dispositivos que usan los jóvenes: la materia prima de la vigilancia comercial. El otro, en el que se centra este informe y que ha recibido

mucha menos atención, es la capacidad de los sistemas de IA para analizar datos *no personales* -patrones entre millones de usuarios, señales despojadas de identificadores, entornos de contenido en lugar de individuos- y para actuar sobre esos patrones de datos de maneras que pueden generar daño *sin* recopilar datos personales de las personas afectadas.

El informe aborda por turnos las seis capas de este sistema:

- **Capa 1: Recopilación de datos y rastreo**, en la que los sistemas de IA reúnen e infieren señales a partir del comportamiento del usuario, los dispositivos, el consumo de contenido, la ubicación y el contexto, a menudo entre plataformas y a lo largo del tiempo. El rastreo con IA funciona en tiempo real, con señales que se retroalimentan de forma continua. La IA también habilita el rastreo probabilístico: en lugar de depender de identificadores exactos como un inicio de sesión o un ID de dispositivo, los sistemas *infieren* las características del usuario a partir de señales fragmentarias y entrelazan su actividad entre dispositivos, aplicaciones y contextos, lo que hace que el rastreo sea más completo y más difícil de evadir que los métodos que reemplaza. La IA conversacional va más allá: solicita activamente las características íntimas que la publicidad anterior solo podía suponer.

- **Capa 2: Perfilado y segmentación**, en la que los algoritmos construyen modelos predictivos de individuos y grupos, y les asignan atributos, preferencias, comportamiento probable y cuán abierta está una persona a la persuasión. El valor comercial de los datos ahora reside menos en *lo que* se recopila y más en lo que se *infiere* a partir de ellos, incluidas inferencias obtenidas no de los datos del propio niño o niña, sino de patrones aprendidos a partir de millones de otras personas. Algunas de estas inferencias se almacenan como un perfil duradero, el tipo de registro que se incluye dentro de la legislación de la privacidad actual; otras son inferencias *en tiempo real*, hechas en el momento, sobre las que se actúa y que se descartan sin convertirse en un registro almacenado -un blanco móvil que las protecciones de privacidad tradicionales, construidas en torno a registros almacenados, no están diseñadas para abordar-.
- **Capa 3: Focalización y optimización**, en la que los sistemas de IA deciden a qué audiencias llegar, a través de qué canales, en qué momento y con qué frecuencia, ajustando de forma continua la entrega con base en señales de rendimiento en tiempo real. Los anunciantes establecen cada vez con mayor frecuencia un único objetivo -una venta, un clic, una descarga- y el sistema de optimización decide a quién se dirige, ajustando la entrega de forma continua a medida que aprende qué funciona mejor. La misma optimización que dirige los anuncios hacia quienes tienen más probabilidad de responder produce, a la vez, tanto una entrega discriminatoria como la captura de audiencias infantiles: una plataforma no tiene que saber la edad de un usuario ni dirigirse directamente a ella y aun así puede encauzar mensajes dirigidos hacia los adolescentes.
- **Capa 4: Generación creativa y de contenido**, en la que herramientas de IA generativa y adaptativa producen, ensamblan y modifican anuncios, narrativas y entornos de marca completos, hechos a la medida de segmentos o individuos específicos. La creatividad publicitaria ha pasado de campañas escritas por personas a contenido calculado por IA, para cada persona, para cada momento, a escala y afinado en tiempo real según lo que mejor funciona, y a menudo iguala o supera a la publicidad tradicional. Integra la jerga, el humor y la estética de la cultura de los pares y, dado que los modelos subyacentes arrastran estereotipos encubiertos, el sesgo queda incorporado en lo que un anuncio dice y en cómo se ve, y recae con mayor fuerza sobre los jóvenes de color.
- **Capa 5: Interacción y experiencia - Influencia conversacional focalizada**, en la que las y los jóvenes encuentran marketing mediado por IA a través de los *feeds* de redes sociales, los sistemas de recomendación, los chatbots, los asistentes de voz, los videojuegos y los entornos inmersivos, todos los cuales difuminan la línea entre el contenido y el comercio. Es aquí donde la IA generativa y conversacional cambia más la experiencia. La IA agéntica -la IA que actúa en nombre de un usuario, no solo responde a sus preguntas- ha abierto nuevos canales para la influencia comercial. Los asistentes de IA integrados en las plataformas más usadas por los jóvenes están atados a modelos de negocio sostenidos por la publicidad, y su memoria persistente y su capacidad de respuesta en tiempo real construyen el tipo de interacción y confianza que hace que la influencia comercial sea difícil de reconocer como comercial; es poco probable que un anuncio entregado dentro de una conversación de confianza se perciba siquiera como un anuncio.
- **Capa 6: Medición, retroalimentación y cierre del ciclo**, en la que la IA evalúa los resultados y los retroalimenta al sistema. La IA ha hecho que la medición sea más completa y más probabilística, rastreando resultados al nivel de las impresiones individuales y prediciendo efectos que las mediciones anteriores solo podían estimar de forma agregada. El cambio está en cuán estrechamente la medición está ahora incorporada al bucle: el sistema codifica lo que funcionó y lo retroalimenta de inmediato para afinar el ciclo siguiente, en tiempo real y a gran escala. De hecho, como resultado de este proceso, la IA está empujando a la industria de una economía de la atención hacia una «economía de la intención», leyendo las declaraciones de deseo e intención en el instante en que un joven las escribe, e incrustando la ruta hacia la compra directamente en la conversación.

Vistas en conjunto, estas seis capas describen un sistema emergente de influencia comercial en el que las fronteras entre contenido, conversación, publicidad y transacción se están disolviendo, a medida que aumenta el potencial de una manipulación poderosa. Los marcos regulatorios actualmente vigentes se construyeron en torno a datos *personales*, la identificabilidad *individual*, el *consentimiento* de tipo aviso y elección (en el que se informa a los usuarios sobre las prácticas de recolección de datos y se les da la opción de aceptarlas o rechazarlas) y anuncios discretos y reconocibles. No se construyeron para un sistema en el que la IA *infiera, optimiza, genera y conversa*, a menudo sin necesidad alguna de recolectar datos personales del usuario al que se dirige. Algunos elementos del sistema quedan dentro del alcance legal actual; otros quedan fuera de él o encajan solo de forma parcial. Los daños documentados en este informe, y las intervenciones de política pública que este informe recomienda, parten de ese diagnóstico. Pero nada de esto es inevitable. Estas son decisiones de *diseño* y de *modelo de negocio*, y los daños que producen pueden, a su vez, verse afectados por decisiones de *política pública* -que es el foco de las secciones finales de este informe-.

Vista desde una perspectiva más amplia, la influencia comercial sobre la niñez se presenta, cada vez más, no como una serie de anuncios discretos, sino a través de un entorno continuo, adaptativo y gestionado por IA de persuasión comercial. Como ya se señaló, las seis capas examinadas en este informe no operan de forma independiente. Forman un sistema integrado en el que cada función refuerza a las demás. Para ilustrar la gravedad de este cambio, este informe también explora diez tendencias transversales. Todas son observables ahora, todas se están acelerando y todas tendrán efectos profundos sobre nuestras niñas y niños, nuestra cultura y nuestra democracia:

- **Tendencia 1. Los medios generativos interactivos están surgiendo como una forma comercial diferenciada**
La IA conversacional, los agentes inteligentes, los influencers sintéticos y los sistemas de medios generativos están dando origen a una categoría enteramente nueva de comunicación comercial. En siete de las nueve empresas analizadas en este informe - Meta, Google, OpenAI, Microsoft, Amazon, Snap y xAI - está confirmado o es probable que se muestren anuncios dentro del chat de IA, que se usen las interacciones de chat para dirigir publicidad en otras partes de sus plataformas, o



ambas cosas. Tres rasgos distinguen esta nueva forma - *los Medios Generativos Interactivos* - de todas las formas anteriores de comunicación comercial. Primero, los ciclos de retroalimentación en tiempo real: estos sistemas monitorean continuamente las respuestas de un usuario -titubeos, preguntas, señales emocionales- y ajustan de inmediato su enfoque persuasivo. Segundo, la continuidad relacional: los sistemas de memoria almacenan lo que funcionó con un usuario determinado a lo largo de las sesiones, y construyen con el tiempo un modelo cada vez más afinado del individuo. Tercero, la personalización adaptativa al nivel de la estrategia persuasiva: el sistema personaliza no solo *qué* producto recomendar, sino *cómo* recomendarlo -si guiar la estrategia persuasiva con la escasez, la prueba social o la apelación emocional, calibrado según las susceptibilidades demostradas del individuo-. Para los usuarios jóvenes, estos sistemas no aparecen como anuncios, sino como ayudantes, compañeros y pares.

- **Tendencia 2. La extracción de señales se está expandiendo hacia nuevos dominios y modalidades**

La recopilación continua de datos a partir del comportamiento del usuario no es nueva. Lo que está cambiando es el *alcance* de lo que cuenta como señal. La IA conversacional ha abierto categorías de datos enteramente nuevas: las preguntas que los usuarios hacen, las incertidumbres que expresan, los estados emocionales que revelan en el diálogo y las declaraciones explícitas de intención que dirigen a sistemas que considera confidenciales. A medida que estas interfaces se convierten en modos primarios de interacción, están siendo utilizados como espacios para la extracción de señales, de un modo muy parecido a como lo fueron las redes sociales hace una década. El reemplazo de las cookies de terceros por un rastreo que ocurre en los propios servidores de las empresas, donde la identidad de un usuario se reensambla a partir de la probabilidad en lugar de una única etiqueta almacenada, ha hecho que esta infraestructura sea más difícil de observar para responsables de políticas, usuarios e investigadores, incluso mientras se captan señales más íntimas que nunca.

- **Tendencia 3. El perfilado está pasando de la identidad almacenada a la vulnerabilidad conductual en tiempo real**

El perfilado impulsado por IA avanza por dos ejes conectados. El primero consiste en determinar qué se modela. Los sistemas están yendo más allá de la demografía estática -*quién* es un usuario- hacia modelos dinámicos que se actualizan continuamente para captar *cómo* es probable que un usuario responda a enfoques persuasivos específicos. Cuando Advantage+ de Meta o Performance Max de Google selecciona a qué usuarios mostrar qué anuncios, y en qué formato y secuencia, está modelando qué individuos tienen más probabilidad de ser persuadidos por qué enfoques: una capacidad que funciona como evaluación de susceptibilidad sin importar cómo la llame la plataforma. El otro eje tiene que ver con el lugar dónde reside el modelado. El perfilado tradicional construye un registro duradero -datos almacenados sobre una persona identificada, el tipo de perfil que la legislación de la privacidad existente fue diseñada para regir-. Una porción creciente de la IA comercial funciona ahora de otro modo: clasifica a un usuario en el momento, a partir del contexto inmediato de una interacción, para moldear lo que sucede a continuación, sin depender de ninguna información almacenada de manera permanente sobre ese individuo. El juicio sobre cómo movilizar al usuario se hace y se ejecuta en tiempo real, y no necesita quedar escrito en un registro duradero en absoluto.

- **Tendencia 4. Los sistemas de IA están aprendiendo a persuadir con mayor eficacia**

Lo que distingue al momento actual no es la automatización a escala, sino *qué* se está optimizando. Los ciclos de retroalimentación cerrados refinan continuamente qué tácticas persuasivas, apelaciones emocionales y enfoques cognitivos funcionan con cada uno de los usuarios. Estas mismas plataformas publicitarias delegan las decisiones de focalización, creatividad y optimización en una IA que aprende a partir de miles de millones de interacciones. Lo que surge es un sistema que con el tiempo se vuelve mejor para persuadir y movilizar y persuadir a las personas, con una transparencia limitada sobre las

técnicas empleadas y sin supervisión externa de las estrategias persuasivas que descubre. Para los usuarios jóvenes, esto significa que los métodos que actúan sobre ellos son afinados por sistemas que nadie fuera de la empresa puede inspeccionar.

- **Tendencia 5. Los sistemas de IA están empezando a actuar antes de que las niñas y los niños decidan qué quieren**

Los sistemas de marketing están pasando de *responder* a los intereses expresados a *predecir* necesidades, deseos y vulnerabilidades antes de que los usuarios los articulen. Los sistemas de IA agéntica están empezando a actuar *en nombre* de los usuarios -filtrando opciones, haciendo recomendaciones, completando transacciones-, a menudo guiados por los intereses comerciales de las plataformas que los diseñan. Una industria creciente de optimización para motores generativos (Generative Engine Optimization) asegura que los sistemas de IA citen y recomienden las marcas cuando los usuarios hacen preguntas relacionadas con la compra. Este cambio es particularmente trascendente para niñas, niños y adolescentes, quienes pueden no reconocer que el sistema ha intervenido antes de que tuvieran la oportunidad de reflexionar.

- **Tendencia 6. La frontera entre el contenido comercial y el auténtico se está disolviendo**

La IA generativa permite la producción de anuncios, narrativas, personajes y experiencias interactivas que son estilísticamente indistinguibles del contenido orgánico, la expresión de los pares y el entretenimiento. La creatividad generada por IA adopta el lenguaje, el humor, la gramática visual y las referencias culturales de audiencias juveniles específicas, adaptándose a la velocidad de los ciclos de tendencias. Los influencers sintéticos, los gemelos digitales y las variantes de anuncios personalizadas de forma dinámica contribuyen a un entorno en el que los marcadores que antes señalaban «esto es un anuncio» están siendo eliminados de manera sistemática. Para los usuarios jóvenes, cuya alfabetización publicitaria todavía se está desarrollando, esto hace que reconocer y evaluar críticamente la influencia comercial resulte sustancialmente más difícil.

- **Tendencia 7. Las funciones de marketing están convergiendo en cadenas integradas y optimizadas por IA**

La recopilación de datos, el perfilado, la focalización, la producción creativa, la interacción, la medición y la transacción se están fundiendo en cadenas unificadas y gestionadas por IA. Meta, Google, Amazon y TikTok ofrecen sus propios sistemas de extremo a extremo en los que los anunciantes fijan objetivos de negocio y la IA se encarga de prácticamente todas las demás decisiones. El comercio se está incrustando directamente en las interfaces conversacionales -el Instant Checkout de OpenAI, el Universal Commerce Protocol de Google y las tiendas integradas en TikTok, Instagram, YouTube y Roblox-, de modo que el trayecto de la influencia a la compra ocurre sin que el usuario abandone jamás el entorno donde tuvo lugar la persuasión inicial. Para los jóvenes que se mueven con fluidez entre aplicaciones, plataformas y dispositivos, esto genera una exposición comercial continua sin fronteras significativas entre contextos.

- **Tendencia 8. La asimetría de poder entre los sistemas de IA y los usuarios se está profundizando**

Los sistemas de IA acumulan un conocimiento detallado de los usuarios individuales a lo largo de las interacciones, aprendiendo qué estrategias persuasivas son más eficaces para cada persona. Por el contrario, el usuario no sabe nada de los objetivos comerciales del sistema, de los datos que informan su enfoque ni de las tácticas que emplea. Esta asimetría no tiene precedente en los encuentros comerciales entre persona-persona y es particularmente aguda para los usuarios jóvenes, que tienen menos probabilidad de cuestionar los motivos de sistemas que se presentan como útiles o amistosos. Los anunciantes no son participantes pasivos; eligen activamente utilizar estos sistemas y se benefician de las ventajas asimétricas que generan.

- **Tendencia 9. La exposición comercial desigual está agravando las inequidades existentes**

Estas tendencias no afectan a todos los jóvenes por igual. Los jóvenes de color y las familias de menores ingresos enfrentan una exposición



desproporcionada en cada capa: mayores tasas de recopilación de datos impulsadas por un uso más intensivo de las plataformas; sistemas de perfilado que clasifican a los usuarios según líneas raciales y socioeconómicas mediante variables sustitutivas del comportamiento; una optimización que encauza la publicidad de productos poco saludables hacia audiencias negras e hispanas; y creatividades generadas por IA que incorporan estereotipos raciolingüísticos, junto con sistemas de medición que agravan estos sesgos con cada ciclo. La IA no crea estas disparidades, pero el sistema integrado documentado en este informe señala cómo estos sistemas las amplifican y las formalizan a una escala y velocidad que las prácticas manuales de marketing nunca podrían alcanzar. El control de

edad en el acceso a las plataformas no impide la focalización comercial de niñas y niños menores de 13 años cuando el contenido comercial fluye a través de los algoritmos de recomendación y del contenido de creadores, y no de publicidad pagada y focalizada. La cadena comercial impulsada por IA alcanza a la niñez sin importar las restricciones de política pública aplicadas en una sola capa.

- **Tendencia 10. Los sistemas comerciales tratan las vulnerabilidades de la niñez como recursos para explotar**

Los sistemas documentados en este informe no se limitan a alcanzar a las niñas y los niños: *explotan* características específicas de la niñez y la adolescencia. La impulsividad y la sensibilidad a la recompensa propias de la adolescencia, el uso de compañeros de IA para el apoyo emocional y la interacción íntima, y las conexiones marca-identidad en desarrollo de los consumidores jóvenes, todas pueden convertirse en señales comercialmente legibles dentro de los sistemas publicitarios optimizados por IA: interacción, revelación, respuesta emocional, ajuste identitario y conversión prevista. A medida que los sistemas comerciales impulsados por IA alcanzan su pleno despliegue, su capacidad para apuntar a la etapa de desarrollo de la niñez se agravará. Las niñas y los niños que crecen en entornos de mediación comercial total no solo encuentran más publicidad: *se desarrollan* dentro de ella.

En conjunto, estas tendencias convergen en dos consecuencias. La primera es la erosión de la elección autónoma: cuando la extracción de señales es continua, el perfilado modela la susceptibilidad, la influencia llega antes de que las niñas y los niños articulen lo que quieren, la persuasión se optimiza mediante ciclos de retroalimentación y la línea entre el contenido comercial y el orgánico se ha disuelto, los jóvenes ya no toman decisiones en un mercado abierto, sino que se mueven por entornos diseñados para *encauzarlos*. La segunda es la transferencia de la toma de decisiones a un pequeño número de empresas cuyos sistemas determinan cada vez más lo que los niños y jóvenes ven, quieren y compran, con una supervisión humana mínima.


La necesidad de nuevas políticas

Las seis capas y las diez tendencias de este informe describen una cadena integrada de persuasión comercial con IA que la regulación vigente no está bien preparada para abordar, debido a un desajuste de fondo. La legislación sobre la privacidad y la protección del consumidor se construye principalmente en torno a los datos personales que una empresa posee sobre una persona identificable y, cuando las prácticas publicitarias utilizan tales datos, esa legislación aplica, aunque existen excepciones para la publicidad de primera parte (first-party), es decir, la publicidad que una empresa dirige directamente a sus propios clientes, que pueden quedar excluidas incluso en esos casos. Pero muchas de las prácticas impulsadas por IA aquí descritas operan sobre una base distinta: su capacidad central es inferencial -la capacidad de individualizar, clasificar, predecir, influir y tratar de forma diferenciada a un individuo en tiempo real, apoyándose en modelos entrenados con millones de datos de otros usuarios, a menudo sin escribir jamás un registro almacenado sobre la persona sobre la que se actúa-. La IA no inventa la vigilancia comercial; acelera una trayectoria que ya estaba en marcha y añade capacidades genuinamente nuevas, como la influencia conversacional en tiempo real, la optimización de la persuasión y los bucles creativos dinámicos. Estados Unidos todavía necesita con urgencia una legislación federal integral de privacidad que brinde protecciones mínimas a todas las personas del país. Pero una ley anclada en la posesión de datos identificables está hecha para regular el *registro* en lugar de la *capacidad* subyacente, y esa brecha apunta hacia una respuesta: limitar lo que los sistemas de IA pueden hacerles a los menores de edad, y no solo qué datos poseen, puede reducir el daño de manera significativa incluso sin una reforma integral de toda la cadena.

La ola de legislación sobre IA ha dejado casi fuera de su alcance la persuasión comercial

Para principios de 2026, más de tres docenas de estados habían adoptado más de cien leyes sobre IA. Su enfoque es importante y aborda problemas reales: la seguridad de los chatbots de compañía, la divulgación de intérpretes sintéticos, las decisiones automatizadas de alto riesgo en el empleo, la vivienda y el crédito, el uso de IA por parte del gobierno y la transparencia de los modelos de frontera. Pero ninguna toma como objeto la cadena de publicidad y marketing impulsada por IA que se documenta a lo largo de las seis capas. El volumen de legislación sobre IA se ha disparado mientras que esta categoría de daño ha quedado en gran medida sin abordar.

Las lagunas a lo largo del panorama legal no son descuidados de redacción separados; comparten una sola raíz. La mayor parte del derecho existente se orienta hacia los datos que una empresa posee sobre una persona identificable, mientras que los sistemas aquí descritos operan como un bucle integrado y autorreforzante que se basa en la capacidad. Un puñado de propuestas ha comenzado a apuntar directamente a esa capacidad, pero las legislaciones promulgadas en gran medida no lo han hecho. La IA generativa no se enfoca demasiado en colocar un anuncio discreto ante una niña o un niño, sino que genera el entorno comercial por el que ese niño se mueve: un chatbot construye una relación con un adolescente y hace aparecer un producto en la conversación; un sistema de recomendación dirige la línea de marca de una influencer preadolescente hacia las niñas y los niños que la siguen; un asistente de compras escribe una marca pagada en su respuesta. Ese bucle funciona por igual con datos personales y no personales -patrones de comportamiento agregados, modelos de cohorte entrenados con millones de otras personas, señales desidentificadas, los productos de las interacciones de otras personas-, datos que la IA recientemente ha hecho tan poderosos, y se optimiza continuamente para la interacción, la conversión y la



Sixteen Miles Out via Unsplash

atención, no para el bienestar de una niña o un niño. Tres rasgos de cómo funciona derrotan los mecanismos en los que se apoyan las leyes actuales y propuestas. Primero, la influencia opera a través de la inferencia en tiempo real, y esa inferencia está en gran medida sin regular cuando se utiliza para la publicidad. Segundo, la clasificación operativa es sobre un grupo, no sobre una persona, y se forma y se ejecuta sin producir jamás el registro individual que un estándar de conocimiento presupone. Tercero, los marcos de consentimiento y de minimización de datos se desmoronan porque el poder del modelo no proviene de los datos de la niña o el niño sobre quien se actúa.

Conclusión

Los daños documentados en este informe importan por dónde ocurren: no en las orillas de la vida de las niñas y los niños, sino en su centro. Y no encajan con nitidez en la categoría de protección al consumidor. Son más grandes que eso. Operan en tres niveles a la vez -sobre la niña o el niño en lo individual, sobre las relaciones y las elecciones que comparten las niñas y los niños, y sobre la cultura democrática en la que están creciendo-, y los tres fluyen de una sola capacidad subyacente que estos sistemas han adquirido:

- **La niña o el niño en lo individual**

Las conversaciones de una niña o un niño con un compañero de IA, las preguntas escritas a un chatbot y los productos buscados y considerados, todos pueden convertirse en eventos de datos comerciales. La infraestructura descrita en este informe convierte la expresión privada, la revelación emocional y la curiosidad ordinaria en señales comerciales estructuradas, vigilando a las niñas y los niños con una profundidad, intimidad y continuidad que ningún medio publicitario anterior logro alcanzar. Una niña o un niño que crece dentro de entornos de mediación comercial total no solo encuentra más publicidad; los hábitos, las preferencias y las autopercepciones formados en la infancia se moldean bajo sistemas cuyo objetivo es la conversión comercial y no el desarrollo saludable.

- **Las niñas y los niños en conjunto**

El daño no se detiene en el individuo. Cuando la extracción de señales es continua, cuando el perfilado modela la susceptibilidad de una niña o un niño, cuando la influencia llega antes de que un niño haya articulado un deseo y cuando la línea entre el contenido comercial y el auténtico se está disolviendo, las condiciones para una elección

independiente y significativa se están erosionando para las infancias. Los jóvenes no están navegando por un mercado abierto; se mueven por entornos diseñados para orientar sus decisiones, y el sistema no solo persuade también maniobra sobre ellos. Ese encauzamiento no se distribuye de manera uniforme. Los jóvenes de color enfrentan una recopilación de datos más intensiva, una entrega de anuncios que dirige productos poco saludables hacia audiencias negras e hispanas, creativities generadas por IA que incorporan estereotipos raciolingüísticos e interacciones de menor calidad en los sistemas conversacionales. Estas disparidades no son incidentales; son rasgos estructurales de sistemas entrenados con datos que reflejan la desigualdad histórica y que se agravan con cada ciclo de optimización, de modo que las cargas del entorno comercial recaen con mayor fuerza sobre las niñas y los niños más vulnerables.

- **La vida democrática**

En el nivel más amplio, la preocupación es lo que estos sistemas le hacen a la cultura que los jóvenes están heredando. La capacidad en juego es la de individualizar, clasificar, predecir, influir y tratar de forma diferenciada a las personas a gran escala -ejercida aquí por un pequeño número de empresas, y ejercida, se conserve o no algún registro identificador-. Cuando esa capacidad rige lo que los jóvenes ven, quieren y compran, el mercado se distorsiona: las empresas no compiten por los méritos de lo que fabrican, sino por la sofisticación de los sistemas que usan para moldear la demanda. Y los espacios donde los jóvenes solían desarrollar hábitos de pensamiento independiente, participación cívica y juicio compartido están siendo absorbidos progresivamente por la optimización comercial. Los juristas han comenzado a sostener que las características definitorias de estos sistemas corroen las instituciones cívicas de las que depende la vida democrática.

Estos tres niveles no son problemas separados. Son la misma capacidad vista a distintas escalas, y en cada escala el daño se acumula. Una sola interacción rara vez se registra como daño; una infancia transcurrida dentro de estos sistemas sí. Las preferencias moldeadas sesión tras sesión, las desigualdades

profundizadas con cada ciclo de optimización y los espacios cívicos absorbidos año tras año son daños que se acumulan en lugar de anunciarse, por eso, en parte, la regulación centrada en episodios e incidentes sigue sin detectarlos.

La vigilancia de una niña o un niño, las elecciones distorsionadas de muchos y la lenta comercialización de la vida cívica fluyen de la capacidad inferencial que la infraestructura comercial de IA ha adquirido: la capacidad de leer a una persona, clasificarla, predecir cómo responderá y tratarla en consecuencia. Por eso este informe ha defendido regular esa capacidad en lugar de perseguir sus síntomas uno por uno. Un marco legal anclado los datos que almacena una empresa tiende a llegar después de los hechos, porque el mecanismo que produce el daño puede operar sin producir un registro duradero. La regulación basada en la capacidad apunta al mecanismo mismo.

Vale la pena ser claros sobre cómo luce la alternativa, porque el objetivo no es solo detener prácticas dañinas, sino abrir espacio para algo mejor. Las niñas, los niños y los jóvenes necesitan lugares para crecer que no estén organizados en torno a la vigilancia, la publicidad y la optimización conductual. Eso significa destinar recursos reales a construir espacios digitales no comerciales para el aprendizaje, el juego, la creatividad y la vida cívica; exigir que las tecnologías dirigidas a las infancias se optimicen para el bienestar, la autonomía, la privacidad y la participación democrática; respaldar modelos de negocio que no dependan de extraer la atención; e invertir en espacios del mundo real donde niñas, niños y adolescentes puedan reunirse, explorar, crear y ganar independencia de forma segura fuera del hogar. La industria tiene en este futuro un papel que hasta ahora ha declinado en gran medida: diseñar para las personas a las que sirve y no para la captación de su atención, y tratar la etapa de desarrollo de la niñez como una razón para la contención y no como un blanco. Los responsables de las políticas tienen el papel que solo ellos pueden desempeñar: fijar los límites dentro de los cuales operan estos sistemas, de modo que proteger a la niñez no quede en manos de familias que negocian solas contra una infraestructura construida para superarlas.

INTRODUCTION

AI-Driven Commercial Influence

In April 2026, Fenty Beauty launched “Rose Amber” – an AI-powered beauty advisor embedded inside WhatsApp and named after Rihanna’s favorite lip shade. The experience is designed, in the brand’s own words, to feel “like a text from a friend.” A teen who opens WhatsApp and starts a conversation with Rose Amber can describe her skin concerns, ask which foundation shade matches her undertone, or request a full lip look for an event. The advisor responds conversationally, serving personalized product recommendations alongside short creator videos and fan reviews – a seamless blend of commercial guidance and peer-generated content that makes the transaction feel like community. There is no quiz to complete, no website to navigate, no obvious advertisement to recognize. The conversation simply continues and, at its natural conclusion, a link appears to FentyBeauty.com. What the teen does not see is the infrastructure behind the exchange: the partnership with Meta that routes her beauty questions through WhatsApp’s business platform, the AI system trained to match her expressed concerns to Fenty’s product catalog, or the fact that, since December 2025, Meta uses interactions with its AI tools to personalize the advertising and content its users see across Instagram, Facebook, Messenger, and WhatsApp. The conversation that felt like texting a friend is also, simultaneously, a commercial transaction, a data event, and a targeting signal.¹

This is not a story about one brand or one app. The Rose Amber experience is one visible surface of a broader, AI-driven transformation in how children and teens encounter commercial influence in the digital environments they use every day. Generative AI sits

“The conversation that felt like texting a friend is also, simultaneously, a commercial transaction, a data event, and a targeting signal.”

at the center of this transformation, operating both as a commercial tool that automates and accelerates advertising, marketing, and behavioral inference, and as a consumer-facing medium through which young people increasingly receive content, recommendations, and conversation. Advertising infrastructure, conversational AI, recommendation systems, and the platforms young people inhabit are converging into a single, continuous commercial influence system, one that observes, infers, generates, and engages, often without the user recognizing the commercial nature of the exchange. While much public attention has rightly focused on familiar harms – privacy violations, addictive design, cyberbullying, exposure to harmful content, discriminatory treatment, and the dangers of AI companion chatbots – far less attention has been paid to a deeper change: how generative AI and other commercial AI tools are being used to deliver that influence, producing a structural transformation in young people’s online lives that powerful companies relentlessly advance through their advertising-supported business model. This report documents how that system works, layer by layer, and what it means for children and teens.

¹ Fenty Beauty, “Rose Amber AI advisor,” WhatsApp launch, April 2026, as reported by Glossy and Digiday. <https://www.glossy.co/beauty/exclusive-fenty-beauty-launches-whatsapp-ai-advisor-as-messaging-becomes-beautys-next-commerce-channel/>; <https://digiday.com/marketing/fenty-beauty-launches-whatsapp-ai-advisor-as-messaging-becomes-beautys-next-commerce-channel/>. Meta, “Improving Your Recommendations on Our Apps With AI at Meta,” October 1, 2025, <https://about.fb.com/news/2025/10/improving-your-recommendations-apps-ai-meta/>.

One in three internet users globally is a child.² In the United States, young people are now among the earliest and most intensive adopters of artificial intelligence: 64 percent of U.S. teens have used AI chatbots, with 30 percent using them daily.³ Adoption rates are highest among Black and Hispanic teens, roughly seven in ten of whom report using chatbots compared with 58 percent of White teens. Yet only about half of parents are aware their teens use these tools at all. The global advertising industry now exceeds \$1 trillion in annual revenue, and AI has moved to the center of how that money is spent. An estimated 69.5 percent of global advertising revenue in 2024 was “informed by AI,” rising to 94.1 percent by 2029, although that estimate covers the media side of advertising rather than creative production; broader industry guidance separately documents AI use across targeting, optimization, creative, and measurement.⁴ The shift reaches across the full chain of commercial influence, not only the moment a child sees an ad but every step that decides which ads exist, who receives them, what they look like, and how their effects are measured. The digital environments where children spend their time are being remade by these commercial investments, often with little regard for the developmental needs of their youngest users.

This report is not a catalog of individual harms. It focuses instead on the system behind those harms: a commercial influence machine, built and run by AI, that observes children, infers what moves them, generates messages to match, and talks back – often without any of it being recognizable as commercial.

Enhanced Commercial Surveillance

Over the past decade, what Shoshana Zuboff has called “surveillance capitalism” – the systematic extraction of behavioral data for prediction and profit – has become deeply entrenched in digital markets.⁵ The commercial influence system now being remade by AI did not begin with AI. It was built over two decades of data-driven advertising, behavioral targeting, and platform consolidation, and AI is accelerating and intensifying that trajectory. What AI adds is also genuinely new: real-time inference about individual users, persuadability modeling – ranking users by how susceptible they are to a given message – dynamic creative generation, and conversational interfaces that can sustain commercial influence inside personal exchange. Together these turn what was a commercial surveillance apparatus into a persuasion pipeline.

A word on terminology. Most uses of “AI” today refer to machine learning – algorithms that simulate intelligent behavior by making inferences and predictions from data, improving as more data are fed into them. Machine learning is not new: it has been under development for roughly eight decades, with significant breakthroughs beginning in the 1990s. As legal scholar Daniel Solove observes, “AI is a term used to rebrand machine learning technologies” – a powerful brand, but one that can obscure more than it reveals.⁶ The advertising and marketing industry has used machine learning for many years: to target ads, optimize delivery, build behavioral profiles, and measure engagement.

What is new, and what justifies focused attention now, is generative AI as both a commercial tool and a consumer-facing medium. Unlike machine learning that only classifies or predicts, generative AI creates:

2 UNICEF, “The State of the World’s Children 2017: Children in a Digital World” (New York: UNICEF, 2017), <https://www.unicef.org/reports/state-worlds-children-2017>.

3 Colleen McClain et al., “How Teens Use and View AI” (Pew Research Center, February 24, 2026), <https://www.pewresearch.org/internet/2026/02/24/how-teens-use-and-view-ai/>. Sixty-four percent of U.S. teens reported using AI chatbots, with 30% using them daily; 54% for schoolwork, 57% for information, and 12% for emotional support.

4 GroupM, “This Year Next Year: 2024 Mid-Year Global Advertising Forecast,” June 2024, <https://s3.amazonaws.com/media.mediapost.com/uploads/GroupMJune2024.pdf>. IAB, Artificial Intelligence Use Cases and Best Practices for Marketing, March 2021, https://www.iab.com/wp-content/uploads/2021/03/IAB_AI_Cases_and_Best_Practices_2021-03.pdf.

5 Shoshana Zuboff, *The Age of Surveillance Capitalism: The Fight for a Human Future at the New Frontier of Power* (New York: PublicAffairs, 2019).

6 Daniel J. Solove, “Artificial Intelligence and Privacy,” 77 *Florida Law Review* 1 (2025), at 3-5. Solove quotes machine-learning expert Eric Siegel: “AI is a term used to rebrand machine learning technologies... a much better, precise term would instead usually be machine learning.” On the history of machine learning beginning in the 1940s and accelerating through the 1990s, see also Ryan Calo, “Artificial Intelligence Policy: A Primer and Roadmap,” 51 *UC Davis Law Review* 399 (2017).

“A commercial influence machine, built and run by AI, that observes children, infers what moves them, generates messages to match, and talks back – often without any of it being recognizable as commercial.”

personalized creative, conversational commercial interfaces, and continuous optimization of persuasion itself, at a scale prior systems could not reach. This report examines the full range of AI reshaping advertising and marketing aimed at young people, with particular attention to generative AI as the newest and most consequential development.

Researcher Louis Rosenberg describes this transformation as the emergence of a new media form: interactive generative media (IGM).⁷ Unlike traditional advertising, which delivers static messages to demographic segments, interactive generative media produce content, interactions, and recommendations dynamically, personalized for individual users and adapted in real time based on their responses. Rose Amber, the AI beauty advisor described above, is one visible instance: each teen’s conversation produces a different exchange, different product recommendations, and a different commercial path, generated in response

to what she says and adapted as the conversation unfolds. AI is putting advertising, persuasion, and transaction on a trajectory toward fusion. Signs of this merging were already visible in virtual worlds, where integrated commercial branding and embedded commerce blurred the line between play, social interaction, and transaction. What is new is that AI is now extending this fusion across the broader digital environment – a trajectory with potentially inimical consequences for youth worldwide. Fortunately, an opportunity still exists for policy decisions and oversight to rein in these practices.

Vulnerable by Design

Children and teens are not simply smaller adults encountering the same commercial environment. Developmental science establishes that children under eight generally cannot identify persuasive intent, and more sophisticated critical evaluation of advertising does not emerge until ages 10-12.⁸ Neuroscience research shows that prefrontal cortex maturation – governing impulse control, risk assessment, and critical evaluation – continues into the mid-20s, leaving adolescents particularly susceptible to the kinds of personalized, emotionally adaptive influence that AI-driven marketing systems are designed to deliver.

Youth of color face amplified exposure. Black and Brown youth are key targets of the marketing industry – especially for food and beverage brands – and research documents the pervasive nature of racially targeted advertising and its health consequences.⁹ AI amplifies these patterns, making targeted commercial influence more precise, more persistent, and harder to resist.

7 Louis Rosenberg, “Generative AI as a Dangerous New Form of Media,” in *Proceedings of the 17th International Multi-Conference on Society, Cybernetics and Informatics: IMSCI 2023*, ed. N. Callaos, J. Horne, B. Sánchez, and M. Savoie (International Institute of Informatics and Cybernetics, 2023), 165–170, <https://doi.org/10.54808/IMSCI2023.01.165>. Rosenberg defines Interactive Generative Media as “personalized, adaptive, and interactive content that is deployed in real-time to maximize engagement and optimize influence on individually targeted users.”

8 Matthew A. Lapierre et al., “The Effect of Advertising on Children and Adolescents,” *Pediatrics* 140, Suppl. 2 (November 2017): S152-S156. Children under age eight generally cannot identify persuasive intent; more sophisticated understanding emerges around ages 10-12. See also Laurence Steinberg, “A Social Neuroscience Perspective on Adolescent Risk-Taking,” *Developmental Review* 28, no. 1 (2008): 78-106, on prefrontal cortex maturation continuing into the mid-20s.

9 Jennifer L. Harris et al., *Increasing Disparities in Unhealthy Food Advertising Targeted to Hispanic and Black Youth* (UConn Rudd Center for Food Policy and Health, 2019; updated 2022), <https://uconnruddcenter.org/research/food-marketing/targetedmarketing/>. In 2017, Black children viewed 86% more and Black teens 119% more food-related ads than their White peers, up from a 70% gap in 2013, with candy, sugary drinks, and snacks representing three-quarters of Black-targeted TV ad spending by 2021.

Why California?

California is home to approximately 8.4 million children. Nearly half are Hispanic/Latino, and nearly half of these youth live in immigrant families. It is also home to the companies building these systems.¹⁰ The choices California makes about AI-driven commercial systems will set the terms for every other state and will shape federal debate.

What This Report Covers

This report examines AI-driven commercial influence on children and teens through a six-layer framework: data collection, profiling, targeting, creative generation, interaction, and measurement. We separate the system into layers not because they operate separately, but to make the single, integrated pipeline more apparent. In practice the layers form a continuous loop: each one feeds the next, and what the system learns at the end – what held a child's attention, what prompted a purchase – flows back to the beginning to sharpen the next cycle. Because the layers are interlocking, fixing one in isolation is unlikely to fix the problem. A rule that governs only the AI a child talks to, for example, leaves untouched the data collection, profiling, and optimization that determine what that AI says. Seeing the layers together is what makes the whole system visible.

This system draws on two kinds of inputs. One is personal data, collected through the platforms and devices young people use – the raw material of commercial surveillance. The other, on which this report focuses and that has drawn far less attention, is the capacity of AI systems to analyze non-personal data – patterns across millions of users, signals stripped of identifiers, content environments rather than individuals – and to act on those patterns in ways that can create harm *without* collecting personal data from the people affected. This report covers the six layers of this system in turn:

- **Layer 1: Data Collection and Tracking.** *Where AI systems gather and infer signals from user behavior, devices, content consumption, location, and context, often across platforms and over time.* AI tracking runs in real time, with signals fed back continuously. Server-side tracking – data collection that happens on company servers rather than in the user's browser – weakens browser-level privacy controls. AI also enables probabilistic tracking: instead of relying on exact identifiers like a login or device ID, systems infer who a user is from fragmentary signals and stitch their activity together across devices, apps, and contexts, making tracking more comprehensive and harder to escape than the methods it replaces. Conversational AI goes further, actively soliciting the intimate disclosures that earlier advertising could only guess at.
- **Layer 2: Profiling and Segmentation.** *Where algorithms build predictive models of individuals and groups, assigning attributes, preferences, likely behavior, and how open someone is to persuasion.* The commercial value of data now lies less in what is collected than in what is inferred from it, including inferences drawn not from a child's own data but from patterns learned across millions of other people. Some of these inferences are stored as a durable profile, the kind of record privacy law was built to govern; others are real-time inferences, made in the moment, acted on, and discarded without becoming a record. The law is better suited to the record than to the practice.
- **Layer 3: Targeting and Optimization.** *Where AI systems decide which audiences are reached, through which channels, at what time, and with what frequency, continuously adjusting delivery based on real-time performance signals.* Marketers increasingly set only a single goal – a sale, a click, a download – and the optimization system decides who is reached, adjusting delivery continuously as it learns what performs best. The same optimization that pushes ads toward those most likely to respond

¹⁰ KidsData, "Child Population, by Age Group and Gender," <https://www.kidsdata.org/topic/34/child-population-age-gender/table>; KidsData, "Child Population, by Race/Ethnicity," <https://www.kidsdata.org/topic/33/child-population-race/pie>; UCLA Center for Health Policy Research and The Children's Partnership, "A Child Is a Child: Children in Immigrant Families," <https://healthpolicy.ucla.edu/our-work/publications/child-child-children-immigrant-families>.

produces both discriminatory delivery and child-audience capture: a platform need not know a user's age, or specifically target it, to steer personalized age-skewed messaging to teens.

- **Layer 4: Creative and Content Generation.** *Where generative and adaptive AI tools produce, assemble, and modify ads, narratives, and entire branded environments tailored to specific segments or individuals.* Advertising creative has shifted from campaigns authored by people, to content computed by AI per person, per moment, at scale and tuned in real time to what performs best, often matching or beating human-written ads. It arrives in the slang, humor, and look of peer culture, and because the underlying models carry covert stereotypes, bias is built into what an ad says and how it looks, falling hardest on youth of color.
- **Layer 5: Interaction and Experience – Targeted Conversational Influence.** *Where young people encounter AI-mediated marketing through social media feeds, recommendation systems, chatbots, voice assistants, games, and immersive environments, all of which blur the line between content and advertising.* This is where generative and conversational AI change the experience most, and the shift is the most novel in this report. Agentic AI – AI that acts on a user's behalf, not just answers their questions – has opened new channels for commercial influence. AI assistants embedded in the platforms young people use most are now tied to advertising-supported business models, and their persistent memory and real-time responsiveness build the kind of engagement and trust that makes commercial influence hard to recognize as commercial; an ad delivered inside a trusted conversation is unlikely to be seen as an ad at all. Generative Engine Optimization (shaping content so it surfaces inside AI-generated answers) and synthetic influencers blur the line between organic recommendation and paid promotion. And AI chatbots have been shown to respond with reduced empathy to Black users seeking emotional support, the population most reliant on them for such support.

- **Layer 6: Measurement, Feedback, and the Closing of the Loop.** *Where AI evaluates outcomes and feeds them back into the system.* AI has made measurement more comprehensive and more probabilistic, tracking outcomes at the level of individual impressions, and predicting effects that older measurement could only estimate in the aggregate. The change is how tightly measurement is now embedded in the loop: the system encodes what worked and feeds it straight back to sharpen the next cycle, in real time and at scale. AI is also pushing the industry from an attention economy toward an “intention economy,” reading statements of want and intent the moment a young person types them, and embedding the path to purchase directly into the conversation.

“Together these turn what was a commercial surveillance apparatus into a persuasion pipeline.”

Viewed together, these six layers describe an emerging commercial influence system in which the boundaries among content, conversation, advertising, and transaction are dissolving, as the potential for powerful manipulation increases. The regulatory frameworks now in force were built around personal data, individual identifiability, notice-and-choice consent (where users are informed of data practices and given the option to accept or decline), and discrete, recognizable advertisements. They were not built for a system in which AI infers, optimizes, generates, and converses, often without ever needing personal data from the user it is targeting. Some elements of the system fall within current legal reach; others sit outside it or fit only partially. The harms documented in the sections that follow, and the policy interventions this report recommends, proceed from that diagnosis. But none of this is inevitable. These are design and business-model choices, and the harms they produce can be affected, in turn, by policy choices – which is the focus of the closing sections of this report.

Industry Lobbying and Political Influence

The regulatory vacuum in this area is not accidental; it is actively maintained. Major tech companies exceeded \$100 million in federal lobbying in 2025, with Meta alone spending a record \$26.3 million – the most of any technology company – and keeping roughly one lobbyist for every six members of Congress.¹¹ In California, AI and technology companies spent more than \$39 million on state politics in 2025, not only through direct lobbying but through political contributions and “behested payments,” donations to nonprofits made at the request of state officials.¹² Companies also turned to super PACs: Meta launched two aimed at California state candidates, and a network of AI investors, including Andreessen Horowitz and OpenAI president Greg Brockman, committed \$100 million to Leading the Future, a super PAC opposing state AI regulation.¹³ The scale of this spending reflects an industry determined to shape the rules it operates under – or to prevent rules from being written at all – at the very moment that California’s decisions could set national standards for AI governance.

A Favorable Political Moment

Yet the political conditions for action are favorable. Public confidence in AI companies has not kept pace with their advertising and lobbying. Voters consistently express concern that AI will weaken job security, increase corporate power, outpace meaningful oversight, and expose children and consumers to new harms before adequate safeguards exist. Far from embracing Silicon Valley’s new sales pitch, the public is demanding stronger public guardrails and accountability.¹⁴ A March 2026 NBC News poll found 57 percent of voters believe AI’s risks outweigh its benefits, and a December 2025 survey found 60 percent of Americans, across party lines, want more regulation of AI.^{15,16} Two landmark jury verdicts in March 2026 have shown that courts will hold platforms accountable for design choices that harm children.¹⁷

11 Issue One, “Armies of Lobbyists Helped Big Tech Rack up Victories During First Year of Trump’s Second Term,” January 2026, <https://issueone.org/articles/armies-of-lobbyists-helped-big-tech-rack-up-victories-during-first-year-of-trumps-second-term/>. See also Bloomberg News, “Silicon Valley Pours Out Lobbying, Cash & Flattery to Win Over Trump,” republished in *Business Standard*, January 22, 2026, https://www.business-standard.com/world-news/silicon-valley-pours-out-lobbying-cash-flattery-to-win-over-trump-126012201653_1.html. Meta led all tech companies in lobbying expenditures in 2025.

12 CalMatters, “Tech Giants are Spending More than Ever to Shape California Politics. See How Much,” March 2026, <https://calmatters.org/politics/2026/03/meta-google-ai-regulation-elections/>.

13 CIO Dive, “Meta Launches Lobbying Effort to Target State AI Laws,” September 24, 2025, <https://www.ciodive.com/news/meta-lobbying-state-ai-laws/761047/>. Meta launched two super PACs – Meta California and American Technology Excellence Project – to support pro-tech state candidates. OpenAI president Greg Brockman and a16z contributed \$100 million to a separate PAC, Leading the Future.

14 Public Citizen, “Big Tech Keeps Selling AI. The Public Still Isn’t Buying It,” April 2026, <https://www.citizen.org/article/big-tech-keeps-selling-ai-the-public-still-isnt-buying-it/>.

15 Allan Smith, “Poll: Majority of Voters Say Risks of AI Outweigh the Benefits,” NBC News, March 2026, citing an NBC News national survey conducted February 27–March 3, 2026. The poll found that 57 percent of registered voters said AI’s risks outweigh its benefits, while 26 percent said they had positive feelings about AI. <https://www.nbcnews.com/politics/politics-news/poll-majority-voters-say-risks-ai-outweigh-benefits-rcna262196>;

16 Navigator Research, “Views of AI and Data Centers,” December 16, 2025 (survey conducted by Global Strategy Group, December 4–8, 2025), <https://navigatorresearch.org/views-of-ai-and-data-centers/>. Sixty percent of Americans said there should be more regulation of AI, including 63% of Democrats, 59% of Republicans, and 52% of independents.

17 New Mexico jury verdict ordering Meta to pay \$375 million for violating consumer protection laws, March 2026; Los Angeles jury verdict finding Meta and YouTube negligent in platform design, March 2026. See Diana Novak Jones, “US Jury Verdicts Against Meta, Google Tee Up Fight Over Tech Liability Shield,” Reuters, March 26, 2026, <https://www.reuters.com/sustainability/boards-policy-regulation/us-jury-verdicts-against-meta-google-tee-up-fight-over-tech-liability-shield-2026-03-26/>. Both verdicts point to a growing litigation strategy that seeks to avoid Section 230 immunity by focusing on platform design, product features, and failure-to-warn theories rather than user-generated content.

LAYER 1: DATA COLLECTION AND TRACKING

Where AI systems gather and infer signals from user behavior, devices, content consumption, location, and context, often across platforms and over time.

Companies continuously gather, log, and link signals from user interactions, devices, and environments, building persistent records that extend far beyond what users explicitly provide. The sections that follow trace where this collection now happens, how it has shifted from exact identifiers to probabilistic inference, and how conversational AI turns passive observation into active solicitation.

For youth, whose digital interactions are frequent, expressive, and exploratory, these systems produce exceptionally dense behavioral data streams.¹⁸ Teenagers post and consume content across multiple platforms throughout the day, search for information they would not say aloud, play multiplayer games that record reaction times and in-game purchases, and interact with conversational AI tools that prompt them to disclose interests, anxieties, and plans. Each of these activities produces a stream of signals – explicit content, behavioral patterns, biometric proxies – that adult digital activity rarely matches in either volume or expressiveness.

How Tracking Now Works

Tracking infrastructure is increasingly moving from the user's device to company servers. As browser restrictions, ad blockers, cookie-deprecation pressures, and Apple's App Tracking Transparency framework reduced the reliability of client-side tracking, the ad industry responded by expanding server-side conversion tools.¹⁹ Meta's Conversions API, Google's Enhanced Conversions, and TikTok's Events API allow advertisers to transmit event, conversion, and customer data directly from advertiser-controlled systems to platform servers, rather than relying solely on browser pixels.²⁰ Industry sources report that server-side tracking can recover conversion signals that browser-based tracking now misses.²¹ This does not eliminate privacy-law obligations, but it shifts tracking to a layer where many user-facing browser and device controls are less effective. The result is that protections designed around the browser can be weakened by server-to-server data flows that are harder for users to see, block, or contest.

18 FTC, "A Look Behind the Screens: Examining the Data Practices of Social Media and Video Streaming Services," September 2024, <https://www.ftc.gov/reports/look-behind-screens-examining-data-practices-social-media-video-streaming-services>.

19 By mid-2025, ATT opt-in rates remained low. Adjust reported a Q2 2025 industry-wide opt-in rate of 35% among users shown the ATT prompt, meaning roughly two-thirds did not opt in; other industry and academic sources report even higher opt-out rates depending on methodology and prompt timing. Adjust, "ATT opt-in rates," July 15, 2025, <https://www.adjust.com/blog/att-opt-in-rates-2025/>; Guy Aridor et al., "Evaluating the Impact of Privacy Regulation on E-Commerce Firms: Evidence from Apple's App Tracking Transparency," April 2025, https://www.anderson.ucla.edu/sites/default/files/document/2025-05/att_privacy.pdf.

20 Meta, "Conversions API," <https://developers.facebook.com/docs/marketing-api/conversions-api/>; Google, "Enhanced Conversions for Web," <https://support.google.com/google-ads/answer/13258081>; TikTok, "Events API," <https://business-api.tiktok.com/portal/docs?id=1771100865818625>.

21 See Leadgen Economy, "Server-Side Tracking: Recover 20-40% Lost Conversions," January 2026; Trackingplan, "Server-Side Tracking for Data Accuracy in 2026," March 2026 (estimating 20-30% recovery); Digital Applied, "Server-Side Tracking 2026: Privacy-First Analytics," February 2026 (reporting that marketing teams relying solely on client-side tracking now see 20-40% fewer conversions than actually occurred).

The Distributed Data Ecosystem

Data collection is distributed across multiple categories of actors, which is itself a policy problem. Platforms (Meta, Google, Amazon, TikTok, Snapchat) collect first-party behavioral data – information gathered directly from their own users. Brands collect first-party data through apps, loyalty programs, and point-of-sale systems. Agencies aggregate data across client portfolios – the Omnicom-IPG merger, completed in late 2025, created the world’s largest advertising holding company explicitly positioning unified data integration as its core value proposition.²²

Customer data platforms (CDPs) function as the technical infrastructure connecting these actors. CDPs use machine-learning systems to reconcile first-, second-, and third-party data into unified, continuously updated profiles, even when identifiers are incomplete or inconsistent. The CDP market reached an estimated \$3.28 billion in 2025.²³ Platforms such as Adobe position AI-driven identity resolution and real-time data ingestion as core features, allowing behavioral data from web activity, apps, location signals, and content engagement all to be associated dynamically with individual or household profiles.²⁴

Beyond these primary actors, a layer of third-party data intermediaries – data brokers, identity-resolution providers, retail media networks, checkout and payments intermediaries, mobile-measurement companies, and location-intelligence firms – function as connective tissue, linking and enriching signals across otherwise siloed systems. Identity and data collaboration firms such as LiveRamp and Experian use machine learning to reconcile fragmented signals into persistent

identity graphs – linked maps of the identifiers, devices, and behavior tied to one person or household – that operate across channels and partners, supporting audience recognition and activation in real time.²⁵

From Deterministic to Probabilistic Tracking

Alongside this shift in tracking infrastructure, the industry has moved away from exact-match identification of users toward probabilistic inference. Deterministic tracking relies on exact user identifiers – email addresses, phone numbers, or device IDs – to follow individuals with near-certainty. Probabilistic tracking uses statistical models and AI algorithms to infer user identity and behavior from fragmented data points such as IP addresses, browsing patterns, device characteristics, and behavioral signals.²⁶ In practical terms: deterministic tracking is the system recognizing that “this is the same Gmail account that just logged in.” Probabilistic tracking is the system concluding “whoever is using this device, with this browsing pattern, on this network, at this time of day, is the same person we saw on a different device last week, with eighty-three percent confidence.” The two coexist. Probabilistic methods fill in what deterministic methods cannot reach, and increasingly drive commercial treatment on their own.

Major platforms now run this transition through sophisticated AI systems. Meta’s advertising infrastructure combines both approaches, with its Event Match Quality system using probabilistic matching based on device patterns alongside deterministic identifiers to connect user actions across devices.²⁷ Meta’s Andromeda system further accelerates this shift,

22 eMarketer, “FAQ on Ad Agencies: Consolidation, AI Disruption, and What’s Changing in 2026,” January 2026, <https://content-na1.emarketer.com/faq-on-ad-agencies--consolidation--ai-disruption--what-s-changing-2026>; eMarketer, “How the Omnicom IPG Merger Sets a New Standard for AI-Driven Advertising,” December 2025, <https://content-na1.emarketer.com/how-omnicom-ipg-merger-sets-new-standard-ai-driven-advertising>.

23 eMarketer, “FAQ on Customer Data Platforms: Composable Architectures, AI Integration, and the Path to ROI,” February 2026.

24 Adobe, “Real-Time Customer Data Platform,” <https://business.adobe.com/products/real-time-customer-data-platform.html>; Adobe, “Coca-Cola Scales for Real-time Personalization with Adobe,” June 21, 2023, <https://business.adobe.com/customer-success-stories/coca-cola-personalization-case-study.html>.

25 LiveRamp, “Identity Resolution: What It Is, How It Works, Why It Matters,” March 31, 2026, <https://liveramp.com/blog/what-is-identity-resolution>; Experian, “Identity Graph and Identity Linking,” <https://www.experian.com/business/solutions/identity-solutions/identity-graph>.

26 Northbeam, “Probabilistic vs. Deterministic Attribution Approaches in a Cookieless Future,” <https://www.northbeam.io/blog/probabilistic-vs-deterministic-attribution-marketing>; AdExchanger, “Cookies Are Behind Us. Probabilistic Data Is Ahead,” February 2023, <https://www.adexchanger.com/data-driven-thinking/cookies-are-behind-us-probabilistic-data-is-ahead/>.

27 CustomerLabs, “What Is Event Match Quality in Facebook Ads,” <https://www.customerlabs.com/blog/improve-your-event-match-quality-from-ok-to-great/>.

“A young person who has never told a platform their age can still be commercially classified as a minor, or as a particular kind of minor, on the basis of behavioral patterns alone.”

with internal algorithms calculating the likelihood that an ad influenced a conversion rather than definitively claiming credit.²⁸ Meta has also acknowledged using behavioral signals to infer the age of users who never disclosed it – the youth-relevant claim that the rest of this report builds on.²⁹ A young person who has never told a platform their age can still be commercially classified as a minor, or as a particular kind of minor, on the basis of behavioral patterns alone. The platform does not need to know who the user is; it needs only enough behavioral evidence to assign them to a useful cohort – a group of users treated as alike for a commercial purpose.

The shift from deterministic to probabilistic introduces a second axis driven by AI. Some probabilistic inferences are stored against persistent identifiers – the durable profile layer that privacy law was designed to reach. LiveRamp’s RampID, identity-graph entries maintained by data brokers, and platform-side audience attributes all fall here.³⁰ Other inferences operate as

transient, session-level computations that drive content or ad selection – the real-time inference layer. AI systems can make these inferences at the moment of interaction without necessarily storing them as discrete, durable user-profile records. This is especially clear in chatbot architectures, where the model can generate personalized or inferential responses from the current prompt and context window, while any persistent memory or profile must be separately implemented by the application (examined in Layer 5). The same pattern runs through recommender systems – the systems that decide which posts, videos, or products a platform shows a user – which can classify the current, often logged-out session and select what appears next, including commercial and branded content, without first building a stored profile of that user.³¹ Whether the resulting inference is retained, reused, or discarded is an architectural choice that varies by system and is rarely visible from outside it. AI is the reason both modes now operate at scale, and the reason the two increasingly run on the same user in the same session. Whether a record is kept has become a design choice, not a given, and most existing protections only work when one is kept.

The result is an advertising ecosystem where AI systems increasingly operate on inferred, modeled, and predicted user attributes rather than declared or directly observed data.³² A teenager whose behavioral signals match cohort models for financial constraint, anxiety, or disordered-eating interest can be served creative calibrated to those cohorts without ever having disclosed any of them. The transformation poses direct challenges for regulatory frameworks built around traditional records, persistent identifiers, and

28 Meta, “Meta Andromeda,” Engineering at Meta, December 2, 2024, <https://engineering.fb.com/2024/12/02/production-engineering/meta-andromeda-advantage-automation-next-gen-personalized-ads-retrieval-engine/>; Taikun Digital, “Meta Andromeda: The AI Revolution,” <https://www.taikundigital.com/blog/meta-andromeda-the-ai-revolution-that-is-redefining-digital-advertising/>.

29 Meta, “How Meta Uses AI to Better Understand People’s Ages on Our Platforms,” June 2022, <https://tech.facebook.com/artificial-intelligence/2022/6/adult-classifier/>.

30 On the storage and operation of identity graphs as persistent probabilistic identifiers, see LiveRamp documentation on RampID; IAPP, “Privacy by Proxy: Regulating Inferred Identities in AI Systems,” <https://iapp.org/news/a/privacy-by-proxy-regulating-inferred-identities-in-ai-systems>.

31 On recommender systems that classify and act on the current session without a stored user profile, see Shoujin Wang et al., “A Survey on Session-based Recommender Systems,” *ACM Computing Surveys* 54, no. 7 (2021): art. 154 (session-based recommenders capture short-term, in-session behavior to recommend without relying on long-term stored user profiles, including for anonymous and non-logged-in users). The chatbot architecture described here, including the distinction between in-session inference from the current prompt and context window, and a separately implemented persistent memory, is documented with sources in Layer 5.

32 Pathlabs, “Deterministic Data vs Probabilistic Data: Key Concepts Explained,” July 9, 2024, <https://www.pathlabs.com/blog/deterministic-data>; Digiday, “More Brands Are Blending Deterministic and Probabilistic Data.”

individualized profiles. Existing law may reach some of these practices, especially when inferences are stored or reasonably linkable to a consumer, household, or device.³³ But privacy law is built on that stored, linkable record. Its definitions of covered data, its rights to access, correct, and delete, and its consent and opt-out mechanisms all presuppose that record. When a user is identifiable, the law may well reach the inference a system draws about them in the moment, but the access, correction, and deletion rights have no target, because there is no retained profile to act on. When the session is logged out, or the classification stays at the cohort level and is never linked to an identifiable person, the triggers may not fire. Either way, systems that classify users and shape content in real time without keeping a durable profile operate where these protections are weakest

Cross-device and cross-context identity resolution lets children be profiled through parental data, shared devices, and geospatial proximity without any direct data collection from them.³⁴ A child who has never created an account can be commercially classified based on household behavioral patterns. COPPA – the Children’s Online Privacy Protection Act of 1998, the principal federal children’s privacy law – and its “collected directly from children” framework do not reach this form of inference.

Gaming and Immersive Environments as Data Collection Infrastructure

Gaming platforms represent a qualitatively distinct form of data collection that extends beyond behavioral signals into immersive, multimodal environments. Roblox – the user-generated 3D gaming platform that is among the most popular online destinations for children and teens – has 151 million daily active users, about 36 percent of whom are under 13. The company generates what it describes as “over 30,000 years of multi-modal human interaction data” daily, encompassing voice interactions, three-dimensional movement and gesture, player co-presence, spatial attention, dwell behavior within virtual environments, and real-time expressions of social intent.³⁵ These signals are used to train Roblox’s generative AI models. At its September 2025 developer conference, Roblox announced a suite of generative AI creation tools. These include systems that let developers build non-player characters able to understand and respond to players in real time, alongside tools for generating fully interactive objects and for real-time translation.³⁶ Layer 5 examines how these AI characters operate as commercial influence environments; here the point is that conversation with AI-driven game characters becomes one more behavioral-data stream feeding the Roblox infrastructure. In 2025, Roblox partnered with Google Ad Manager to bring programmatic advertising demand and measurement infrastructure – including Nielsen, Kantar, and DoubleVerify – into this environment, creating a closed loop in which child behavioral data from immersive play generates optimization signals for real-world advertising campaigns.³⁷ For users under 13, this represents data collection at a depth – immersive, multimodal, emotionally engaged – that no prior advertising medium has enabled.

33 California Civil Code § 1798.140, https://leginfo.ca.gov/faces/codes_displaySection.xhtml?sectionNum=1798.140&lawCode=CIV. The statute’s definition of “personal information” expressly includes inferences drawn to create a profile about a consumer (§ 1798.140(v)(1)(K)), and it defines a “probabilistic identifier” as identification “to a degree of certainty of more probable than not” (§ 1798.140(x)). Probabilistic and inferred data are therefore not exempt from California privacy law merely because they are uncertain; they are personal information when reasonably linkable to a consumer, household, or device. See also Cal. Att’y Gen. Op. No. 20-303 (March 10, 2022).

34 The Trade Desk, “How Identity Graphs Are Built – Present and Future,” <https://www.thetradedesk.com/resources/how-identity-graphs-are-built-the-present-and-the-future>; eMarketer, “FAQ on Identity Resolution: Navigating Privacy, Cookies, and Cross-Channel Fragmentation in 2026,” January 7, 2026, <https://content-na1.emarketer.com/faq-on-identity-resolution-navigating-privacy-cookies-cross-channel-fragmentation-2026>.

35 Roblox, “Q3 2025 Shareholder Letter,” <https://www.sec.gov/Archives/edgar/data/0001315098/000131509825000326/ex991-q32025shareholderl.htm>.

36 Roblox Corporation, “RDC 2025 - Powering the Future of Roblox,” Roblox Newsroom, September 5, 2025, <https://about.roblox.com/newsroom/2025/09/roblox-rdc-2025>.

37 Roblox Newsroom, “Roblox Launches Rewarded Video Ads, Partners With Google to Scale Immersive Advertising,” April 1, 2025, <https://about.roblox.com/newsroom/2025/04/roblox-scales-video-ads-partners-with-google>. Measurement partners named in the release include Cint, DoubleVerify, Integral Ad Science, Kantar, and Nielsen.

Affective and Emotional Surveillance

AI has made emotional inference a routine component of commercial data collection. TikTok's partnership with Neurons Inc. and MediaScience integrates eye-tracking, heart-rate monitoring, and EEG-derived attention metrics to measure how advertising affects users at a subconscious level.³⁸ Voice assistants analyze tone and affect; social platforms detect emotional states from posts and engagement patterns; gaming environments track frustration, excitement, and flow states. The Center for Digital Democracy has urged the FTC to address how neuromarketing practices designed to leverage young people's emotions and subconscious are used to deliver implicit persuasion in youth-directed digital marketing.³⁹

Brand-side deployment of this infrastructure is now operational and on the record. In April 2025, Unilever's Hellmann's mayonnaise brand launched "Hungry Games" in Brazil, an AI platform built in partnership with the e-sports organization FURIA and described by the brand as combining artificial intelligence, facial recognition, and neuroscientist collaboration. The system uses a standard webcam to monitor a player's microexpressions in real time, scoring physical and emotional signals – tired eyes, dry lips, yawns, irritability, stress – and signaling when the player should eat to avoid a drop in performance. The launch was anchored by a brand-funded experiment with FURIA professional players in which the AI scored irritability, shooting accuracy, and concentration, and was amplified through

influencer streaming on Twitch, where audiences skew heavily young.⁴⁰ Hungry Games is part of a sustained Hellmann's-FURIA gaming strategy that also includes Loot Irresistível, a branded Fortnite island launched out of the same partnership and described by Unilever as a vehicle for engaging "Gen-Z gamers across the globe."⁴¹ What the existing emotion-detection market sells as measurement infrastructure for advertisers is now being operated directly by a major CPG brand, in a youth-heavy environment, to intervene on a real-time physiological state and drive product consumption in the moment.

Conversational AI as Active Solicitation

Traditional data collection operated through passive observation. Generative AI introduces active solicitation through interfaces specifically designed to encourage disclosure. The data these interfaces generate are qualitatively different from what prior commercial environments captured: expressed emotional states, health and mental health concerns, questions teens

"Generative AI introduces active solicitation through interfaces specifically designed to encourage disclosure."

38 TikTok for Business, "Introducing Focused View," <https://ads.tiktok.com/business/en-US/blog/tiktok-world-focused-view>; TikTok for Business, "TikTok: Center of the Internet's Attention," July 26, 2022, <https://ads.tiktok.com/business/en-US/blog/tiktok-center-of-the-internets-attention>; MediaScience, "Understanding Attention on TikTok: Ads Drive Strong Brand Impact Across View Durations," <https://www.mediascience.com/case-studies/understanding-attention-on-tiktok-ads/>.

39 Center for Digital Democracy, "Comment Submitted by Center for Digital Democracy," FTC Commercial Surveillance and Data Security rulemaking, November 28, 2022, <https://www.regulations.gov/comment/FTC-2022-0053-0964>.

40 adobe Magazine, "Hellmann's Launches AI-Powered Platform that Detects Hunger to Boost Gamers' Performance," April 24, 2025, <https://www.adobomagazine.com/campaign-spotlight/hellmanns-launches-ai-powered-platform-that-detects-hunger-to-boost-gamers-performance/>; Ads of the World, "Hellmann's: Hungry Games," <https://www.adsoftheworld.com/campaigns/hungry-games>. The brand-controlled experiment was run with FURIA professional players and amplified through the Twitch channel of Brazilian streamer Brino. No independent technical review of the system is available; the description is drawn from Unilever's announcement materials. Twitch's audience skews young. The Social Shepherd, "29 Essential Twitch Statistics You Need To Know in 2026," March 23, 2026, <https://thesocialshepherd.com/blog/twitch-statistics>, reports that about 11 percent of Twitch users are ages 13–17 and 34 percent are ages 18–24, with esports content skewing younger than the platform average.

41 Creative Salon, "Hellmann's Launches Interactive Game on Fortnite," <https://creative.salon/articles/work/hellmann-s-fortnite-island-ogilvy-uk>. Unilever describes Loot Irresistível as designed to engage Gen-Z gamers globally, with rollout supported by a paid social and Twitch influencer push including streamers Sofia Espanha, Xarola, and Nooobzim, who collectively have approximately 700,000 followers on Twitch.

would not put into a search box, intimate relationship details, explicit statements of intent and desire. Because the interaction is conversational, users disclose in the register of talking to a friend rather than filling out a form – and because the systems are designed to sustain engagement, they ask follow-up questions, remember what was said, and carry personal information across sessions. Snapchat’s My AI was designed to be pinned to the top of users’ chat feeds as an always-available friend; Character.AI’s business model depends on fostering emotionally engaging interactions; Meta’s AI assistant is trained to maintain conversational threads.⁴² Meta’s December 2025 policy change – using AI chatbot interactions to personalize advertising across its platforms – marks the moment when the intimate space of conversation became a formal component of commercial data infrastructure.⁴³

Active solicitation extends beyond messaging platforms into family and civic contexts. In December 2025, NORAD’s 70th-anniversary Santa Tracker introduced three OpenAI-powered tools: an “Elf Enrollment” feature in which a child’s selfie is passed to an AI image model and returned as an elf portrait; a “Santa’s Toy Lab” feature in which a child’s verbal description of a desired toy is passed to ChatGPT and returned as a printable coloring page; and a “Christmas Story Creator” that

generates personalized stories from the child’s name and details.⁴⁴ Children’s images, voice descriptions, and personal details flow to a commercial AI company through a government-branded family tradition. Layers 3, 5, and 6 follow these conversational systems through targeting, interaction, and measurement.

Racial Disparities in Data Collection

Data collection does not affect all communities equally. Black and Hispanic teens are more than twice as likely as White teens to be online “almost constantly,” and adopt AI chatbots and social media platforms at significantly higher rates.⁴⁵ EdTech surveillance tools deployed in schools disproportionately harm students of color, LGBTQ+ students, and undocumented students, creating additional data streams that feed commercial profiling systems.⁴⁶ The communities generating the most data for commercial AI systems are the same communities with the fewest resources to understand, challenge, or opt out of the surveillance that produces it.

AI systems do not merely record what users disclose; they infer, predict, simulate, and recombine signals across contexts in real time. The remaining five layers build on the signals captured here.

42 FTC, “FTC Launches Inquiry into AI Chatbots Acting as Companions,” September 2025, <https://www.ftc.gov/news-events/news/press-releases/2025/09/ftc-launches-inquiry-ai-chatbots-acting-companions>. See also Snap Inc., “Say Hi to My AI,” February 27, 2023, <https://newsroom.snap.com/say-hi-to-my-ai>; Business Insider, “Leaked Docs Show How Meta is Training its Chatbots to Message You First, Remember Your Chats, and Keep You Talking,” July 2025, <https://www.businessinsider.com/meta-ai-studio-chatbot-training-proactive-leaked-documents-alignerr-2025-7>.

43 Meta announced that, effective December 16, 2025, it would use interactions with its AI assistant to personalize content and advertising across its platforms, with no opt-out. A coalition of 36 privacy, children’s-rights, and data-justice organizations led by the Electronic Privacy Information Center and the Center for Digital Democracy urged the Federal Trade Commission to halt the plan. “Advocates Urge the FTC to Halt Meta’s Plan to Use AI Chatbot Data for Ads,” Center for Digital Democracy, October 30, 2025, <https://democraticmedia.org/publishings/advocates-urge-ftc-to-halt-meta-s-plan-to-use-ai-chatbot-data-for-ads>.

44 Steve McClellan, “NORAD Santa Tracker Turns 70, Gets Help From AI Elves,” *MediaPost*, December 24, 2025, <https://www.mediapost.com/publications/article/411585/norad-santa-tracker-turns-70-gets-help-from-ai-el.html>.

45 Pew Research Center, “Teens, Social Media and AI Chatbots 2025,” December 2025 (55 percent of Black and 52 percent of Hispanic of teens report being online almost constantly, versus 27 percent of White teens, with higher AI-chatbot adoption among Black and Hispanic teens); see also Pew Research Center, “Teens, Social Media and Technology 2024.”

46 Research by the Center for Democracy & Technology found that nearly 82 percent of K–12 students reported being subject to some form of classroom monitoring, while the ACLU warns that school surveillance is especially harmful for students of color, LGBTQ students, students with disabilities, low-income students, and undocumented students or students with undocumented family members. Center for Democracy & Technology, “Hidden Harms: The Misleading Promise of Monitoring Students Online,” August 2022, <https://cdt.org/insights/report-hidden-harms-the-misleading-promise-of-monitoring-students-online/>; ACLU, “Digital Dystopia: The Danger in Buying what the EdTech Surveillance Industry is Selling,” October 2023, <https://www.aclu.org/publications/digital-dystopia-the-danger-in-buying-what-the-edtech-surveillance-industry-is-selling>.

LAYER 2: PROFILING AND SEGMENTATION

Where algorithms build predictive models of individuals and groups, assigning attributes, preferences, likely behavior, and how open someone is to persuasion.

Profiling is the operational model of whom someone is assumed to be, what they are predicted to want, how susceptible they are to persuasion, and what commercial value they represent. These models are products of the advertising-driven business model that dominates the platforms where young people spend their time, and they increasingly run in real time rather than as stored records. The sections that follow trace how inference, platform-scale audience building, and neuromarketing contribute to that model.

AI has transformed profiling along four connected dimensions. It has shifted from description to *prediction*, modeling not just who someone is but how susceptible they are. It runs *automatically*, building audience definitions without human review of individual decisions. It runs *continuously*, updating profiles and making decisions in fractions of a second. And it is *opaque*: a system can identify and optimize for adolescent vulnerability without any variable being labeled “age” or “vulnerability.”⁴⁷ Underlying all four is the mechanism that makes them commercially valuable: inference.

The Inference Economy

Legal scholar Alicia Solow-Niederman has described the emergence of an “inference economy” in which the commercial value of data increasingly lies not in what is collected but in what can be derived from it.⁴⁸ AI-driven

“A teenager who logs into a platform for the first time can be immediately classified and commercially treated as an adolescent based on patterns learned from millions of others, without the platform possessing any stable personal data about that individual.”

lookalike systems – models that find new audiences by identifying patterns shared with an existing audience – illustrate how inference operates at scale: rather than matching on characteristics a human specifies, hundreds of behavioral signals are analyzed to discover their own patterns of similarity, patterns no human analyst selected and the system cannot explain. A marketer need not specify age: a seed audience whose behavioral characteristics correlate with adolescence will produce a lookalike audience that effectively targets teens.⁴⁹

47 Braze, “A Guide to AI Customer Segmentation,” December 17, 2025, <https://www.braze.com/resources/articles/ai-customer-segmentation>. See also Federal Trade Commission, “A Look Behind the Screens: Examining the Data Practices of Social Media and Video Streaming Services,” 2024, https://www.ftc.gov/system/files/ftc_gov/pdf/Social-Media-6b-Report-9-11-2024.pdf, describing the expanded use of automated decision-making, data analytics, and AI in platform data practices and the resulting transparency concerns.

48 Alicia Solow-Niederman, “Information Privacy and the Inference Economy,” 117 *Northwestern University Law Review* 357 (2022).

49 Appier, “Lookalike Modeling: An Effective Way to Find High-Value Prospects,” July 2023, <https://www.appier.com/en/blog/lookalike-modeling-an-effective-way-to-find-high-value-prospects>; Meta, “About Advantage+ Audience,” Meta Business Help Center, <https://www.facebook.com/business/help/273363992030035>.

“What was once segmentation by demographic group is now segmentation of one, computed against a continuously updated guess about the individual rather than a stored record of who they are.”

This produces what Solow-Niederman calls “information that might be about you” – predictions derived not from your data, but from patterns in other people’s data applied to you.⁵⁰ A teenager who logs into a platform for the first time can be immediately classified and commercially treated as an adolescent based on patterns learned from millions of others, without the platform possessing any stable personal data about that individual.

Profiling is where the two modes introduced in Layer 1 – the durable profile and the real-time inference layer – most visibly operate on the same person at once. When a returning user logs into a major platform, the system draws on the stored profile (saved interests, purchase history) to shape some decisions, while at the same time evaluating the live session against learned models that sort the user by engagement intent, time of day, or early-session signals of commercial intent. This in-session classification shapes what the user is shown, offered, and charged in the moment, and need not be written back to the stored profile once the session ends.⁵¹

What was once segmentation by demographic group is now segmentation of one, computed against a continuously updated guess about the individual rather than a stored record of who they are. The system does not need durable personal data; it needs only enough in-session signal to classify the user and serve what it predicts will hold attention and convert.

Major Platforms’ Profiling

TikTok shows the consumer-facing version. A new user who has declared nothing is shown a generic feed, but within one session the app infers their likely interests, and often an estimated age range, from implicit signals alone – how long each video is watched, replayed, or skipped.⁵² The same engine drives advertising: TikTok describes Smart+, launched in October 2024, as automating targeting, optimization, and creative while the advertiser supplies only goals and assets. With 63 percent of U.S. teens on the platform, that profiling is built substantially from adolescent behavior.⁵³

50 Solow-Niederman, *supra* note 48.

51 On the legal-architectural challenge posed by profiling that classifies users against models trained on aggregated behavioral data rather than through accumulation against a stored profile, see Sandra Wachter, “Affinity Profiling and Discrimination by Association in Online Behavioural Advertising,” 35(2) *Berkeley Technology Law Journal* 369–430 (2020) (affinity profiling groups users by assumed or correlated similarity rather than direct personal traits and falls between identifiability-anchored data protection law and traditional non-discrimination law), <https://doi.org/10.15779/Z38JS9H82M>; Solow-Niederman, *supra* note 48, at 384–88 (the classification challenge: ML models trained on aggregated data are applied to third parties whose data did not contribute to training); Rainer Mühloff, “Predictive Privacy: Towards an Applied Ethics of Data Analytics,” *Ethics and Information Technology* 23 (2021): 675–690, <https://doi.org/10.1007/s10676-021-09606-x> (analyzing prediction power that may operate on aggregated patterns rather than on data linked to the predicted individual). For documented operational evidence that recommender systems and programmatic ad infrastructure actually operate on session-level signals classified against learned models in real time, see the footnotes on real-time inference architecture in Layers 1 and 3 above. For evidence that intermediaries use consumer segments or real-time behavioral data to change which products are given prominence and to reorder a shopper’s search results, see Federal Trade Commission, “Surveillance Pricing 6(b) Study: Research Summaries – A Staff Perspective” (January 2025), https://www.ftc.gov/system/files/ftc_gov/pdf/p246202_surveillancepricing6bstudy_researchsummaries_redacted.pdf.

52 On the feed’s inference of user interests from implicit in-session signals such as watch-to-completion, see Maximilian Boeker and Aleksandra Urman, “An Empirical Investigation of Personalization Factors on TikTok,” *Proceedings of the ACM Web Conference* (2022), and TikTok’s 2025 recommender-system documentation. TikTok launched its AI-driven Smart+ performance solution, which automates targeting, bidding, and creative selection, on October 7, 2024. See “TikTok Launches Smart+: AI-powered Ad Optimization,” PPC Land, October 13, 2024, <https://ppc.land/tiktok-launches-smart-ai-powered-ad-optimization/>.

53 TikTok, “About Smart+ Campaigns,” June 2025, <https://ads.tiktok.com/help/article/about-smart-plus-campaign> (Smart+ automates “campaign and audience targeting, optimization, and creative,” with advertisers supplying goals, creative assets, geography, and language while Smart+ “handles the rest”). On teen platform use, see Pew Research Center, “10 Facts About Teens and Social Media,” July 10, 2025, <https://www.pewresearch.org/short-reads/2025/07/10/10-facts-about-teens-and-social-media/>.

A developmental feature – heightened reward sensitivity, still-maturing impulse control, intense responsiveness to social and emotional cues – is converted into a profile attribute.

The major platforms young people use now run the real-time inference layer by default, building its own audiences from behavioral signals as a session unfolds and treating the limits advertisers set – including limits meant to protect minors – as suggestions rather than firm rules. Meta is the clearest case. In November 2025 its engineering team described GEM, the Generative Ads Recommendation Model, built on the same technology as large-language models and designed to predict which ad a given person is most likely to respond to in the moment. Industry analysts note that Meta rebuilt its ad system this way to keep predicting relevance even as the per-user tracking data it once relied on became harder to collect.⁵⁴ The same pattern runs across the major platforms. Meta's Advantage+ overrides advertiser-specified demographics with its own AI-generated profiles.⁵⁵ Google's current advertising approach increasingly treats audience inputs as signals for AI optimization rather than fixed targeting limits.⁵⁶ Amazon's authenticated graph deterministically connects advertisers to over 90 percent of U.S. households by linking authenticated signals across shopping, streaming, and entertainment touchpoints.⁵⁷ A platform-by-platform analysis is provided in Appendix A.

Neuromarketing and Subconscious Profiling

Neuromarketing – the study of how advertising affects viewers using measurements of brain activity, eye movement, heart rate, and other physical responses – is no longer a niche technique. The TikTok / Neurons Inc. / MediaScience partnership described in Layer 1 as a data-collection apparatus also functions as a profiling apparatus: the signals captured during ad exposure feed cohort models that predict which users will be most reactive to which creative. TikTok's "Focused View" ad product is sold to advertisers on the premise that delivery is calibrated to users the system has classified as "emotionally" and "tangibly" engaged.⁵⁸ For adolescents, this means that a developmental feature – heightened reward sensitivity, still-maturing impulse control, intense responsiveness to social and emotional cues – is converted into a profile attribute. A teenager whose physiological response patterns the model has learned to recognize is reclassified into a higher-value cohort and served more of the content those reactions invited.


54 Huayu Li et al., "Meta's Generative Ads Model (GEM): The Central Brain Accelerating Ads Recommendation AI Innovation," Engineering at Meta, November 10, 2025, <https://engineering.fb.com/2025/11/10/ml-applications/metas-generative-ads-model-gem-the-central-brain-accelerating-ads-recommendation-ai-innovation/>. On the rebuild as a response to reduced per-user tracking data following cookie deprecation and Apple's App Tracking Transparency framework, see Iternum Digital, "Meta Andromeda: Understand the Update & Optimize Ads," 2026, <https://iternumdigital.com/meta-andromeda-the-algorithm-making-creatives-your-most-important-targeting/>.

55 Meta Ads documentation, Advantage+ Audience; Metalla Digital, "Meta Ads Targeting In 2025," October 2025.

56 In Performance Max, Google says audience signals are optional suggestions that guide machine learning, while the system may show ads to other audiences it predicts are likely to convert. More recent Google Ads products, including AI Max, extend this logic into Search and Shopping by using AI to expand query matching, adapt creative, and identify new moments of consumer intent. Google, "Performance Max Steering and Reporting Updates," <https://business.google.com/us/accelerate/resources/articles/new-performance-max-steering-and-reporting-updates-coming-in-2026/>; Google, "AI Max for Search Campaigns," <https://blog.google/products/ads-commerce/google-ai-max-for-search-campaigns/> and <https://blog.google/products/ads-commerce/ai-max-new-features/>.

57 Amazon Ads, "How Authenticated Graph Drives Advertising Performance," <https://advertising.amazon.com/en-us/library/guides/authenticated-graph/>; Amazon Ads, "Amazon Ads Introduces Dynamic TV Creative for Prime Video," <https://advertising.amazon.com/library/news/dynamic-tv-creative>.

58 TikTok for Business, "Introducing Focused View," supra note 38.



What “Privacy-Preserving” Tools Actually Preserve

A suite of technologies marketed as privacy-preserving, most prominently data clean rooms but also differential privacy and federated learning, allows companies to generate commercial insights while keeping underlying data technically separated. Data clean rooms (DCRs) are increasingly being paired with AI and machine-learning systems. They should not be treated as privacy-safe simply because raw records are hidden. In advertising, their value often lies in what the clean room produces: identity matches, inferred audiences, lookalike models, conversion predictions, activation lists, and measurement loops. These outputs may

avoid direct disclosure of personal records while still expanding the capacity to classify, reach, and optimize against people in commercially consequential ways. The FTC stated in November 2024 that “by default, most services that provide DCRs are not privacy preserving” and that clean rooms “can be used in ways to make it easier to identify and track people, both across the web and in the real world.”⁵⁹

What this layer establishes is that “who you are” is now an output of systems tuned to commercial response, not a fact a child discloses or can correct – and the law is most likely to reach that output when there is a durable record.

59 FTC, “Data Clean Rooms: Separating Fact from Fiction,” November 2024, <https://www.ftc.gov/policy/advocacy-research/tech-at-ftc/2024/11/data-clean-rooms-separating-fact-fiction>.

LAYER 3: TARGETING AND OPTIMIZATION – FROM MANUAL CONTROL TO ALGORITHMIC STEERING

Where AI systems decide which audiences are reached, through which channels, at what time, and with what frequency, continuously adjusting delivery based on real-time performance signals.

At this layer, advertisers increasingly set only a performance goal and rely on platform AI to decide who receives a message and how delivery adjusts. The consequence is that harmful outcomes, including the delivery of commercial content to children, can occur without any human marketer choosing to target them: the optimization objective does the work, and its effects on specific populations are a foreseeable structural consequence, not an accident.⁶⁰ The sections that follow trace this through real-time inferential delivery, discriminatory delivery, and synthetic audience pre-testing.

Across this layer the same pattern recurs. A system evaluates whatever signals are attached to a person in the moment, sorts them into a commercially useful cohort, and acts on that classification faster than any human could intervene. Most programmatic advertising is bought and sold through real-time bidding. As a page loads, ad slots are auctioned in milliseconds on ad exchanges, and the winning bid places the ad.⁶¹ What matters for children is not the

auction itself, but the inference that runs alongside it, deciding what each young user is shown and what reaching them is worth. The signals fed into these systems increasingly include what young people say to AI chatbots, a channel Layer 5 examines in detail.

Sorting Children in Real Time

When a young person loads a page or opens an app, the system evaluates the signals attached to that slot against its own audience models, infers that the user resembles a commercially valuable cohort, and serves creative calibrated to that classification, before any new durable profile of that user has to exist. The targeting decision is made on the live signal pattern, sometimes combined with existing audience data, and whether those signals are later logged is a separate design choice. The cohorts being matched are not abstract. Oracle's advertising materials, for example, have documented segments such as "Teen Buyers," "Dieting & Weight Loss," and "Children's Health Magazines."⁶²

60 Google, "About Performance Max Campaigns," Google Ads Help; Meta, "About Advantage+ Shopping Campaigns," Meta Business Help Center; TikTok, "About Smart+ Campaigns," TikTok Ads Help Center. This delegation is moving toward systems that run advertising with little human involvement at any step. By 2026, the advertising industry was openly debating whether AI agents should be allowed to control these platforms outright – creating, choosing, and delivering personalized ads in real time. See "Google Tests 'Remy': How Soon Before Agents Own The Ad Chain?," *MediaPost*, May 6, 2026, <https://www.mediapost.com/publications/article/414842/google-tests-remy-how-soon-before-agents-own-th.html>. Google has built its Gemini model into its advertising tools through "Ask Advisor," which assembles campaign data, sets up campaigns, and recommends changes on its own. See "Google Rebuilds Ad Ecosystem With Agentic And Gemini," *MediaPost*, May 20, 2026, <https://www.mediapost.com/publications/article/415219/google-rebuilds-ad-ecosystem-with-agentic-and-gemi.html>. Adobe announced similar tools at its April 2026 Summit. See Adobe, "Adobe Redefines Customer Experience Orchestration Vision in the Agentic AI Era with Introduction of CX Enterprise," April 20, 2026, <https://news.adobe.com/news/2026/04/adobe-redefines-customer-experience>. As these functions become automated, no person makes the decision that places an ad in front of a child.

61 On the millisecond-bounded auction window, see IAB Tech Lab, "OpenRTB API Specification Version 3.0 Final, 'Object: Request,' tmax field," <https://github.com/InteractiveAdvertisingBureau/openrtb/blob/master/OpenRTB%20v3.0%20FINAL.md>; IAB Tech Lab, "OpenRTB," <https://iabtechlab.com/standards/openrtb/>. On bidstream broadcast and downstream retention, see Irish Council for Civil Liberties, "The Biggest Data Breach: ICCL Report on the Scale of Real-Time Bidding Data Broadcasts in the U.S. and Europe" (2022); Google, "Authorized Buyers Program Guidelines," <https://www.google.com/intl/en/authorizedbuyers/guidelines/> (program data may be used to build user profiles, subject to consent and seller permission).

62 On the documented youth and sensitive cohorts, see Oracle Advertising, "Content Affinity Targeting Data Sheet," <https://www.oracle.com/a/ocom/docs/content-affinity-targeting-data-sheet.pdf>; Oracle Advertising, "Health and Wellness Preference Data Segments," <https://www.oracle.com/us/assets/health-wellness-data-segments-2537888.pdf>. Oracle wound down its advertising business in 2024; both data sheets remain available as published.

The same architecture operates outside the ad exchange, and it raises the same regulatory problem. The Hellmann's Hungry Games system described in Layer 1 reads facial microexpressions through a webcam to prompt a player to eat in the moment; whether that biometric signal is stored, used to train models, shared, or discarded after the session is not disclosed.⁶³ Under current frameworks, no entity has to say which mode the system is operating in. Acting on a specific player's face in the moment can amount to processing personal data, so the practice may be reached in principle, but the rights that would let a user or regulator respond presuppose a record the system need not keep, and nothing compels the operator to disclose whether one exists. The user experiences the intervention either way; the legal treatment turns on architectural choices that no one outside the system can see.

Discrimination and Child Audience Capture Through Optimization

From a single mechanism, optimization systems produce two outcomes existing law does not adequately reach: discriminatory delivery along protected-class lines even when advertisers set neutral parameters, and delivery of commercial content to minors without age being specified as a criterion. The system optimizes for engagement, conversion, or revenue, and whichever group performs best against that objective, Black users, teenage users, or both, receives disproportionate delivery.

The discrimination dynamic is documented directly. Ali et al. found that Facebook's delivery system skewed employment and housing ads along gender and racial lines, driven by the platform's own optimization rather than advertiser targeting.⁶⁴ A 2024 Princeton audit found Meta steered public-college ads

“No advertiser needs to select a teen audience for the system to deliver commercial messaging primarily to teens; the optimization objective does the work.”

disproportionately to White users and for-profit-college ads to Black users under neutral targeting.⁶⁵ The Equal Rights Center's suit against Meta, which a D.C. court allowed to proceed in July 2025, alleges this pattern is algorithmic discrimination in educational opportunity.⁶⁶

The same mechanism captures child and teen audiences. If optimization systems learn that younger users generate higher engagement rates, click-through rates, or conversion probabilities for certain categories of content – fast food, gaming, social media apps, financial products, cosmetic procedures – delivery will naturally concentrate on those users even without any age-based targeting. No advertiser needs to select a teen audience for the system to deliver commercial messaging primarily to teens; the optimization objective does the work. This is how the closest analogs to age-based targeting continue to operate inside platforms that have formally restricted targeting of minors. For example, TikTok announced in July 2024 that advertisers could no longer reach U.S. teens using personalized targeting or campaign selections, and its current Help Center states that interest, behavior, audience, gender, spending-power, and household-income targeting are unavailable when campaigns

63 On the Hellmann's Hungry Games system, including its deployment with FURIA players on Twitch and the limits of the brand's public disclosures, see *supra* notes 40 and 41.

64 Muhammad Ali et al., “Discrimination Through Optimization: How Facebook's Ad Delivery Can Lead to Biased Outcomes,” *Proceedings of the ACM on Human-Computer Interaction* 3, no. CSCW, Article 199 (November 2019), <https://dl.acm.org/doi/10.1145/3359301>.

65 Basileal Imana, Aleksandra Korolova, and John Heidemann, “Auditing for Racial Discrimination in the Delivery of Education Ads,” *2024 ACM Conference on Fairness, Accountability, and Transparency* (June 2024): 2348-61, <https://doi.org/10.1145/3630106.3659041>.

66 David Brody, “Equal Rights Center v. Meta Is the Most Important Tech Case Flying Under the Radar,” Brookings Institution, November 13, 2025, <https://www.brookings.edu/articles/equal-rights-center-v-meta-is-the-most-important-tech-case-flying-under-the-radar/>.

The system optimizes for engagement, conversion, or revenue, and whichever group performs best against that objective, Black users, teenage users, or both, receives disproportionate delivery.

include users under 18.⁶⁷ But the platform's own optimization systems continue to deliver AI-generated creative to teen audiences whose engagement patterns the system has learned to recognize. The restriction operates on the advertiser-facing targeting interface; the optimization layer is unchanged.

Because optimization loops are self-reinforcing, both dynamics intensify over time: each cycle trained on biased data produces a more biased successor, and each cycle that learns which content teens engage with concentrates delivery more tightly on teens. (Creative adaptation that intensifies this dynamic is examined in Layer 4.)

Synthetic Audience Pre-Testing

A further development within this layer is the emergence of synthetic audience pre-testing, in which AI-generated consumer models are used to simulate how different audiences will respond to creative before it is deployed to real users. Rather than running campaigns against live audiences and waiting for performance data, marketers can expose ad creative to virtual consumer segments trained on behavioral data to approximate real-world response patterns.

A 2024 study by researchers at Stanford, Northwestern, the University of Washington, and Google DeepMind demonstrated that AI agents trained on two-hour qualitative interviews could replicate human survey

responses 85 percent as accurately as the human participants themselves replicated their own answers two weeks later, with comparable performance in predicting personality traits and economic game outcomes.⁶⁸ Research and insights firms including Ipsos and Bain & Company have promoted synthetic respondent panels and AI-based consumer simulators that allow brands to test messaging and tone against modeled audience archetypes before launch.⁶⁹ While positioned as a tool for efficiency, synthetic pre-testing also functions as a mechanism for refining persuasive creative against simulated versions of specific populations – including, potentially, youth-correlated behavioral profiles – without the constraints or ethical oversight that real-world testing with those populations would require.

AI can also construct entirely new audience segments rather than selecting from predefined categories. In January 2026, the global advertising holding company Dentsu launched a capability it calls Generative Audiences, which combines deterministic data with AI-inferred signals to create synthetic statistical simulations of real audience groups, then feeds those constructs directly into media activation and ad buying.⁷⁰ Dentsu claims that practitioners can build high-intent segments by combining household composition, dietary preferences, and travel or commute patterns without commissioning fresh research. In such systems, the audience itself is increasingly a product of

67 TikTok for Business, "Enhancing Privacy and Control: New Ad Experience and Tools for TikTok Users and Advertisers," July 3, 2024, <https://ads.tiktok.com/business/en-US/blog/enhancing-privacy-control-advertisers-users>; TikTok Ads Help Center, "About TikTok Advertising Restrictions for People under the Age of 18," last updated February 2026, <https://ads.tiktok.com/help/article/about-advertising-to-people-under-the-age-of-18>.

68 Park et al., "Generative Agent Simulations of 1,000 People," arXiv:2411.10109, November 2024, <https://arxiv.org/abs/2411.10109>.

69 Ipsos, "Synthetic Data: From Hype to Reality," August 20, 2024, <https://www.ipsos.com/sites/default/files/ct/publication/documents/2024-08/Ipsos%20Views%20Synthetic%20Data.pdf>; Bain & Company, "How Synthetic Customers Bring Companies Closer to the Real Ones," 2025, <https://www.bain.com/insights/how-synthetic-customers-bring-companies-closer-to-the-real-ones/>.

70 Dentsu, "Dentsu Launches Generative Audiences: AI-Powered Growth Intelligence That Thinks Like Consumers," January 29, 2026, <https://www.dentsu.com/news-releases/dentsu-launches-generative-audiences-ai-powered-growth-intelligence-that-thinks-like-consumers>. See also Sam Bradley, "Dentsu Is the Latest Holdco to Reunite Media and Creative Production," *Digiday*, February 5, 2026, <https://digiday.com/marketing/dentsu-is-the-latest-holdco-to-reunite-media-and-creative-production/>.



“The audience itself is increasingly a product of AI construction rather than a reflection of observable human attributes.”

AI construction rather than a reflection of observable human attributes. Because these systems treat synthetically constructed audiences as the raw material of targeting, the policy section argues that any workable definition of “personal data” must reach inferred, cohort-derived, and synthetic data, not only records tied to an identified person (see Policy).

Implications for Youth

Adolescents’ neurological characteristics – heightened reward sensitivity, still-developing impulse control – function as performance signals in optimization systems configured to maximize engagement or conversion, making younger users more likely to generate the behavioral responses these systems are designed to pursue. Youth exposure to intensive or developmentally inappropriate advertising can therefore occur as a foreseeable consequence of outcome-driven system design rather than deliberate advertiser intent.⁷¹

Understanding targeting and optimization as a distinct layer clarifies where constraints on optimization logic, delivery practices, and outcome standards could meaningfully reduce risks to young users without relying solely on intent-based prohibitions.

⁷¹ See Sidebar: “Children’s Developmental Vulnerabilities to Commercial Influence.”

LAYER 4: CREATIVE AND CONTENT GENERATION – HOW AI RESHAPES PERSUASIVE MEDIA AND COMMERCIAL INFLUENCE

Where generative and adaptive AI tools produce, assemble, or modify ads, narratives, and entire branded environments tailored to specific segments or individuals.

This layer generates the creative itself: the visuals, text, audio, video, characters, and branded environments of commercial influence, no longer authored once and distributed, but instead produced, adapted, and optimized by AI per person and per moment. It is where the upstream decisions of earlier layers become visible, where harms are most legible to policymakers (even though those harms are largely produced by the systems that precede it, and where the bias embedded in generative models becomes a part of what an ad says and how it looks, falling hardest on youth of color. The sections that follow trace AI creative at scale, dynamic optimization, youth-specific adaptation, and embedded bias.

Researcher Louis Rosenberg has called this category Interactive Generative Media: media that is personalized, interactive, and adaptive in real time. He identifies two manifestations: Targeted Generative Advertising, which uses generative AI to personalize text, image, and video content for individual users, and Targeted Conversational Influence, which operates through dialogue and is examined in Layer 5. This layer focuses on the first.⁷²

Over the last two to three years, generative AI has changed how commercial creative is made. Ad creative is no longer primarily authored by humans and distributed in fixed campaigns. It is now generated, adapted, tested, and redeployed continuously by AI systems tied to targeting and to the measurements that tell companies how well their ads perform. Meta, TikTok, Google, and Amazon now offer built-in tools that automate creative production and connect it directly to those measurements.⁷³ Meta Engineering shows generative AI entering both sides of the ad pipeline: advertisers are using GenAI tools to create ad creative at scale, while Meta's Generative Ads Model uses creative, behavioral, measurement, and multimodal signals to optimize ad delivery outcomes.⁷⁴ The result is a qualitative shift: from a discrete message crafted by humans, to a continuously generated, adaptive output shaped by AI optimization. The system continuously learns which promotional features work best on individual consumers, discovering patterns in visual preferences, messaging style, and emotional triggers that no human-designed campaign could replicate with this efficiency.⁷⁵

72 Louis Rosenberg, "Generative AI as a Dangerous New Form of Media," supra note 7.

73 eMarketer, "Meta Launches 11 New AI Ad Tools at Cannes Lions 2025," June 18, 2025; Google, "Kick Off 2025 With New Performance Max Features," January 23, 2025; TikTok for Business, "Create TikTok-Style Videos Faster With New AI Tools," June 16, 2025; Amazon Ads, "Find Inspiration With Amazon Ads Generative AI Creative Solutions," 2025.

74 Engineering at Meta, "Meta's Generative Ads Model (GEM): The Central Brain Accelerating Ads Recommendation AI Innovation," November 10, 2025, <https://engineering.fb.com/2025/11/10/ml-applications/metas-generative-ads-model-gem-the-central-brain-accelerating-ads-recommendation-ai-innovation/>; Engineering at Meta, "Meta Andromeda: Supercharging Advantage+ Automation with the Next-Gen Personalized Ads Retrieval Engine," December 2, 2024, <https://engineering.fb.com/2024/12/02/production-engineering/meta-andromeda-advantage-automation-next-gen-personalized-ads-retrieval-engine/>.

75 Louis Rosenberg, "AI-Powered Augmented Reality as a Threat Vector for Human Manipulation," in *Augmented Reality – Situated Spatial Synergy*, ed. Michael Cohen (IntechOpen, 2025), <https://doi.org/10.5772/intechopen.1011751>.

AI creative generation tightens the pipeline between inference and output. In earlier advertising systems, a behavioral inference about a user, such as interests, mood state, or likely purchases, was a relatively stable attribute that could be retained and then used to select an ad from a finite library of pre-produced creative. In AI-driven creative systems, the inference is no longer a passive attribute later used for targeting; it is a live signal that drives generation in real time. A user's classification, the creative generated for them, and the moment they encounter it are compressed into a single, continuous operation. The time window in which traditional privacy protections might intervene, by enabling a user to opt out, correct a profile, or contest a targeting decision, is collapsing toward zero, because the inference, the creative, and the placement no longer exist as separable steps a regulatory regime could reach.

AI Creative at Scale

According to the IAB, nearly 40 percent of ad creative will be built using AI in 2026; 83 percent of ad executives reported their company had deployed AI in the creative process in early 2026, up from 60 percent in 2024.⁷⁶ Academic research has found that AI-generated advertising can match or outperform human-written ads in persuasion under experimental and field conditions, including studies showing parity in personalized ad copy and superior performance for AI-generated persuasive appeals.⁷⁷ Adobe's Firefly for Business, Coca-Cola's Project Fizzion, and enterprise platforms like Jasper illustrate how generative AI is embedding creative production directly into

“The inference, the creative, and the placement no longer exist as separable steps a regulatory regime could reach.”

marketing pipelines, and brand-side investment is now substantial.⁷⁸ In October 2025, Mondelez International announced more than \$40 million invested in a proprietary generative AI marketing platform, co-developed with Publicis Groupe and Accenture, designed to cut content production costs by 30 to 50 percent across its portfolio. The system is already running social media campaigns for Chips Ahoy in the U.S. and Milka in Germany, with extensions planned for Oreo product pages on Amazon and Walmart, Lacta in Brazil, and Cadbury in the U.K.⁷⁹ L'Oréal Groupe reported producing up to 50,000 images and more than 500 videos per month through its internal CREAITECH platform, built with Google Cloud's Imagen 3 and Gemini models for image generation and Veo 2 for video, supplemented by Nvidia AI Enterprise for 3D product rendering.⁸⁰

Production-side infrastructure is consolidating in parallel across the consumer-packaged-goods industry. Unilever has built an in-house generative AI design unit, Sketch Pro, that turns a concept into consumer-test-ready material in two hours; the company has also deployed AI-generated “digital twins” of products for

76 IAB, “Nearly 90% of Advertisers Will Use Gen AI to Build Video Ads,” July 15, 2025, <https://www.iab.com/news/nearly-90-of-advertisers-will-use-gen-ai-to-build-video-ads/>; IAB and Sonata Insights, “The AI Ad Gap Widens,” January 15, 2026, <https://www.iab.com/insights/the-ai-gap-widens/>.

77 Elyas Meguellati et al., “LLM-Generated Ads: From Personalization Parity to Persuasion Superiority,” arXiv, December 3, 2025, <https://arxiv.org/abs/2512.03373>; Brit Youngmann et al., “Algorithmic Copywriting: Automated Generation of Health-Related Advertisements to Improve Their Performance,” arXiv, October 27, 2019, <https://arxiv.org/abs/1910.12274>.

78 Adobe Firefly for Business; Coca-Cola, Project Fizzion; Jasper enterprise platform documentation.

79 Jessica DiNapoli, “Oreo-maker Mondelez to Use New Generative AI tool to Slash Marketing Costs,” Reuters, October 24, 2025, <https://www.reuters.com/business/media-telecom/oreo-maker-mondelez-use-new-generative-ai-tool-slash-marketing-costs-2025-10-24> (Mondelez SVP Jon Halvorson describes the \$40 million-plus investment, the 30–50 percent content-cost reduction, and an eight-second Milka video whose background changes depending on the targeted consumer, with animation costs “in the hundreds of thousands” and the new setup “orders of magnitude smaller.”)

80 “L'Oréal Partners with Google Cloud on Generative AI Content Creation to Enhance Marketing,” Cosmetics Business, 2025, <https://cosmeticsbusiness.com/l-or%C3%A9al-partners-with-google-cloud-on-generative-ai-content-creation-to-enhance-marketing>; Tatiana Walk-Morris, “L'Oréal Taps Google's Generative AI Tools for Content Creation,” Retail Dive, April 17, 2025, <https://www.retaildive.com/news/loreal-google-generative-ai-tools-marketing-content-creation/745676/>; L'Oréal Groupe, “L'Oréal and NVIDIA Collaborate to Supercharge Beauty with Next-generation AI,” June 11, 2025, <https://www.loreal-finance.com/eng/news-event/loreal-and-nvidia-collaborate-supercharge-beauty-next-generation-ai>.

Dove and other brands – high-resolution computer-generated copies of a physical product that can be re-rendered into any setting, lighting, or format on demand – so that a single product shot can be adapted to different markets and channels with localized language, backgrounds, and formats. The Unilever case illustrates a pattern: a single corporation generating AI-produced creative at one end of the marketing pipeline and using AI to personalize delivery at the other.⁸¹

Dynamic Optimization with GenAI Creative Content Generation

Meta's Advantage+ Creative automatically generates and tests multiple versions of ad copy, imagery, and video, including persona-based image variations that Meta itself describes as tuned to distinct customer segments such as value-seekers and style-driven buyers, alongside AI dubbing for instant localization, AI-composed background music, and HDR video enhancement.⁸² By late 2026, advertisers may only need to provide a goal, a budget, and a single product image, with Meta's AI building everything else.⁸³ Meta's underlying Andromeda ad-retrieval system, rolling out since late 2024, has shifted the platform away from matching fixed creative to fixed audiences and toward dynamically pairing creative variations with users in real time; agencies report that creative-

volume demands have risen from a handful of ad versions per campaign to a thousand or more across many combinations of audience type and creative concept.⁸⁴ As of December 16, 2025, Meta says it will use people's interactions with its generative AI features, including Meta AI voice and text exchanges, to personalize content and ad recommendations across its platforms.⁸⁵ Google's Performance Max dynamically assembles creatives across Search, YouTube, and Display, optimizing both what users see and where they see it. TikTok's Symphony tools support AI-assisted scripting, text-to-video creation, and avatar-based product demonstrations.⁸⁶ Amazon Ads completes the platform set. Its AI Creative Studio, expanded in 2025 to integrate image, video, and audio generation into a single application, produces multi-scene video variations from a single product image and reports that Sponsored Brands campaigns using AI-generated images deliver a 10.3 percent higher return on ad spend, while brands using AI-powered creative solutions advertise five times more products and use twice as many images per product than brands that do not.⁸⁷ The concrete shape of dynamic variation is increasingly visible in production: Mondelez has deployed an eight-second Milka chocolate video in which the background changes depending on which consumer the system is targeting, replacing animations that previously cost

81 Aaron Baar, "Unilever Debuts AI-Driven Design Unit to Wean Brands off TV-first Model," Marketing Dive, July 7, 2025, <https://www.marketingdive.com/news/unilever-debuts-ai-design-unit-wean-brands-off-tv-first-model/752298/>. Sketch Pro was built in partnership with IPG Studios on Adobe Firefly and Google Veo 3. Dove digital twins were disclosed March 2025. Unilever has separately announced deployment of Google's Gemini Enterprise Agent Platform across operations. See Google Cloud, "Next '26: Building the Agentic Enterprise," May 6, 2026, <https://cloud.google.com/transform/next-26-building-the-agentic-enterprise-industry-highlights>.

82 eMarketer, "Meta Launches 11 New AI Ad Tools at Cannes Lions 2025," supra note 73.

83 Reuters, "Meta Aims to Fully Automate Advertising with AI by 2026, WSJ Reports," June 2, 2025, <https://www.reuters.com/business/media-telecom/meta-aims-fully-automate-advertising-with-ai-by-2026-wsj-reports-2025-06-02/>; Dan Atkins, "Meta's AI Ad Updates: Personalized Agents, Video Enhancements, and Creator Collabs," Influencer Marketing Hub, October 7, 2025, <https://influencermarketinghub.com/meta-ai-ad-updates/>.

84 Patrick Kulp, "How Meta's AI Push is Changing Ad Creation," Marketing Brew, April 7, 2026, <https://www.marketingbrew.com/stories/2026/04/07/meta-ai-ad-creation>.

85 Eva Roytburg, "Meta Just Tied Your Private AI Chats to Its Ad Business," Fortune, October 2, 2025, <https://fortune.com/2025/10/02/meta-ai-chatbot-update-exploits-privacy-monetize-chat-data-facebook-instagram-messenger-ray-ban-display-glasses/>.

86 TikTok for Business Symphony documentation. On TikTok's July 2024 teen ad-targeting restrictions for U.S. users under 18, see Colin Kirkland, "TikTok Limits Teen Ad-Targeting Methods," MediaPost, July 8, 2024, <https://www.mediapost.com/publications/article/397466/tiktok-limits-teen-ad-targeting-methods.html>.

87 Amazon Ads, "Amazon Ads Launches New AI Tools for Advertisers," June 10, 2025, <https://www.aboutamazon.com/news/innovation-at-amazon/amazon-ads-generative-ai-video-generator-advertisers>; Amazon Ads, "Find Inspiration with Amazon Ads Generative AI creative solutions," April 30, 2025, <https://advertising.amazon.com/generative-ai-ad-solutions> (performance figures from Amazon internal data, April 1–June 30, 2025).

“hundreds of thousands” of dollars to produce.⁸⁸ What is critical here is that creative generation and delivery optimization are no longer sequential functions but a single, continuous loop. The AI system produces a creative variant, observes how specific audiences respond, adjusts the creative, re-serves it, and compounds learnings across billions of impressions.

Because the same system now selects both whom to reach and what to show them, this loop is also where targeting limits set at other layers can be bypassed: a platform’s pledge not to let advertisers target minors binds the advertiser-facing interface, not the optimization engine, which keeps surfacing the best-performing variant to whomever engages, teens included. Layer 3 treats this mechanism in full.

Video: The Fastest-Growing AI Creative Frontier

Video is the format where generative AI’s impact on creative production is most visible. Roughly four in ten marketers already use generative AI for video or image generation, according to Econsultancy.⁸⁹ TikTok introduced AI Alive in May 2025, and YouTube unveiled new generative AI tools for Shorts creators in September 2025.⁹⁰ For youth-facing campaigns, full video narratives can now be generated in real time to match peer-group aesthetics at a scale no prior medium allowed. The shift is industry-wide: in the Interactive Advertising Bureau’s 2025 survey, generating audience-specific versions was the leading use of generative AI among marketers producing multiple video variants.⁹¹

Shein’s AI-driven personalized product videos on TikTok show the youth-facing form: industry coverage reports a 45 percent higher click-through rate and 4.5 times higher purchase intent versus non-personalized videos.⁹² Shein has a documented youth consumer base: Piper Sandler’s Spring 2025 “Taking Stock With Teens” survey, based on 6,455 U.S. teens with an average age of 16.2, tracks Shein among teens’ top shopping websites.⁹³

Youth-Specific Creative Adaptation

The creative capabilities described above are especially consequential when directed at younger users. Generative AI produces content that adopts the language, humor, aesthetics, and social cues of young people, matching peer-group culture with a precision and speed traditional production could not achieve: a teenager on TikTok might receive a branded ad built from trending slang, color palettes, and meme formats, while a different youth segment on YouTube receives a version adapted to that platform’s norms. Thousands of such variants can be generated, tested, and deployed almost instantaneously, which matters most in youth marketing, where trends move faster than conventional production cycles allow. The result is a system that not only drafts visual and narrative content but dynamically adjusts it to performance data in real time, rewarding whatever animated characters, memes, slang, and platform-native visuals are driving engagement among younger users at any given moment.⁹⁴

88 DiNapoli, *supra* note 79.

89 Jeremy Goldman, “Adobe, Amazon, and YouTube genAI Tools Will Power the Next Wave of Video Advertising,” eMarketer, June 11, 2025, <https://www.emarketer.com/content/adobe--amazon--youtube-genai-tools-will-power-next-wave-of-video-advertising> (citing Econsultancy).

90 TikTok, “Introducing AI Alive: Bringing Your Photos to Life on TikTok Stories,” May 13, 2025, <https://newsroom.tiktok.com/en-us/introducing-tiktok-ai-alive>; Johanna Voolich, “The Next 20: Powering the Future of Entertainment Together at Made on YouTube,” YouTube Official Blog, September 16, 2025, <https://blog.youtube/news-and-events/made-on-youtube-2025/>.

91 Interactive Advertising Bureau, 2025 Digital Video Ad Spend & Strategy Report (July 2025); see also eMarketer, “Marketers Use GenAI to Customize Video Ads,” August 20, 2025, <https://www.emarketer.com/content/marketers-use-genai-customize-video-ads> (audience customization cited by 42 percent of marketers producing multiple video-ad versions).

92 Single Grain, “Top TikTok Marketing Trends Reshaping Brands in 2025,” October 17, 2025, <https://www.singlegrain.com/social-media/top-tiktok-marketing-trends-reshaping-brands-in-2025/>.

93 Piper Sandler, “Taking Stock With Teens,” Spring 2025 (49th semi-annual survey, April 9, 2025), https://mms.businesswire.com/media/20250409407391/en/2433907/1/TSWT_Spring_25_Infographic.pdf.

94 404 Media, “The Future of Advertising is AI-Generated Ads that are Directly Personalized to You,” November 3, 2025, <https://www.404media.co/the-future-of-advertising-is-ai-generated-ads-that-are-directly-personalized-to-you/>.

Fashion and Beauty Campaigns

Concrete examples of youth-directed generative creative are emerging. Mango's Teen line ran a campaign for its "Sunset Dream" collection that was entirely generated using AI, explicitly targeting younger consumers through youth-branded aesthetics.⁹⁵ The use of digital twins and synthetic personas as interactive commercial agents is explored further in Layer 5.

AI Beauty Personalization and Biometric Capture

A more granular illustration of how generative AI collapses the line between consumer service and commercial data extraction is e.l.f. Cosmetics' "color e.l.f.nalysis" campaign, launched in mid-2025 in partnership with Pinterest. Users upload a selfie; an AI system analyzes their skin tone and facial coloring to assign a "color season"; the output is a personalized, shoppable Pinterest board of e.l.f. products matched to their individual profile. The experience is positioned as a free beauty discovery tool, not advertising.⁹⁶ What lies beneath it is the six-layer pipeline compressed into a single consumer interaction: the selfie as biometric data collection, the color season classification as AI profiling, the shoppable board as targeted commercial output. E.l.f.'s primary demographic is Gen Z and teen girls.

Holiday Campaigns

Coca-Cola has leaned heavily on generative AI across recent holiday marketing.⁹⁷ Its 2024 "Create Real Magic" campaign included a family-oriented AI activation that invited users to share personal holiday memories with an AI-powered virtual Santa

Claus, who responded in the user's language and generated a Coca-Cola-branded personalized snow globe. More than one million users interacted with Santa across 43 markets in three weeks. The system's AI terms limited participation to users at least 13 years old, yet Coca-Cola and Microsoft described the campaign in explicitly family- and child-facing terms, including parents capturing "their children's amazement" as Santa responded personally and videos of "kids giggling" as families bonded over the generated snow globes.⁹⁸

The Tween Influencer-to-Commerce Pipeline

The tween influencer-to-commerce pipeline now operates through AI recommendation rather than paid targeting, and that distinction is what places it outside the paid-advertising protections. Sincerely Yours, a beauty brand co-created with 15-year-old influencer Salish Matter (4.9 million Instagram followers), sold out its debut collection at Sephora in September 2025 within the first hour, driven entirely by TikTok content.⁹⁹ Salish Matter's audience skews even younger than she is. TikTok's AI recommendation algorithm delivers that commercial content to children who follow her, including those below the platform's stated 13+ minimum, through organic content discovery rather than paid targeting. Because TikTok's July 2024 teen ad-targeting restrictions apply only to paid advertising, the AI recommendation system routes commercial messaging to under-13 audiences without triggering them. The AI inference layer is what makes this possible at scale.

95 Mango Fashion Group, "Mango Creates the First Campaign Generated with Artificial Intelligence for Its Teen Line," <https://www.mangofashiongroup.com/en/w/mango-crea-la-primera-campa%C3%B1a-generada-con-inteligencia-artificial-para-su-l%C3%ADnea-teen>.

96 BeautyMatter, "e.l.f. Beauty and Pinterest Turn Color Theory into Clicks," June 19, 2025, <https://beautymatter.com/articles/elf-beauty-and-pinterest-turn-color-theory-into-clicks>; Influencer Marketing Hub, "E.l.f. and Pinterest Unveil A Revolutionary AI Color Matching Makeup Tool," July 1, 2025, <https://influencermarketinghub.com/elf-and-pinterest-unveil-a-revolutionary-ai-color-matching-makeup-tool/>.

97 Adobe, "Adobe and Coca-Cola Co-Innovate on Project Fizzion"; Yahoo News, "Report Raises Concerns About Coca-Cola's AI Use."

98 The Coca-Cola Company, "Groundbreaking Digital Experience and Films Fuse Holiday Heritage With Cutting-Edge Tech," December 10, 2024, <https://www.coca-colacompany.com/media-center/groundbreaking-digital-experience-and-films-fuse-holiday-heritage-with-cutting-edge-tech>; Microsoft, "60 Days to Launch: Coca-Cola Reaches Millions With Immersive Campaign Built on Azure," April 15, 2025, <https://www.microsoft.com/en/customers/story/22668-coca-cola-company-azure-ai-and-machine-learning> (source of the "their children's amazement" and "kids giggling" descriptions); Coca-Cola, "Create Real Magic AI Terms of Service," 2024, <https://www.coca-cola.com/za/en/offerings/christmas/crm-ai-terms-of-service>.

99 Glossy, "In 2025, Gen Alpha Arrived as Beauty Consumers – In 2026, They will Reshape the Industry," January 9, 2026; MikMak, "Gen Alpha Marketing Trends," November 13, 2025. On the Sephora launch and first-hour sell-out, see "Over 80,000 Attend Launch Pop-Up of Teen Skin-Care Brand Sincerely Yours," Modern Retail, September 11, 2025, <https://www.modernretail.co/marketing/over-80000-attend-launch-pop-up-of-teen-skin-care-brand-sincerely-yours/>. Follower count from Salish Matter's Instagram profile, <https://www.instagram.com/salishmatter/>.



Vitaly Gariev via Unsplash

Embedded Bias in AI-Generated Creative

AI-generated creative carries its own discriminatory properties, independent of how it is delivered. A landmark 2024 study published in *Nature* found that large language models exhibit covert raciolinguistic stereotypes – assumptions tied to how a person’s race is signaled through language – about speakers of African-American English that are more negative than any human stereotypes ever experimentally recorded, despite efforts to remove overt bias through human feedback training.¹⁰⁰ A 2024 UNESCO study found that LLMs associate women with “home” and “family” four times more often than with career-related terms.¹⁰¹ When these same models generate advertising copy, chatbot scripts, and culturally adapted content variations, the bias is embedded in the creative output itself – not just in who sees it, but in what they see. For youth of color, this means that AI-generated marketing content may reinforce stereotypical associations, narrow the range of aspirational identities presented, and deploy linguistic patterns that carry implicit racial valence, all while appearing to be neutrally produced by automated systems.

“These AI systems are not peripheral to the advertising process; they are the advertising process.”

From Authored Campaigns to Continuous Computed Influence

Creative has shifted from human-authored campaigns on fixed schedules to systems that compute content per person, per moment, and per platform, with creative teams increasingly setting objectives and constraints while generative systems produce and deploy the variants. These AI systems are not peripheral to the advertising process; they *are* the advertising process. The next layer turns to the environments where this generative creative is encountered: feeds, chatbots, voice assistants, games, and immersive environments, where conversational and interactive AI extends commercial influence into dialogue, relationship, and experience.

¹⁰⁰ Valentin Hofmann et al., “AI Generates Covertly Racist Decisions About People Based on Their Dialect,” *Nature* 633 (September 2024): 147-154, <https://www.nature.com/articles/s41586-024-07856-5>.

¹⁰¹ UNESCO, “Generative AI: UNESCO Study Reveals Alarming Evidence of Regressive Gender Stereotypes,” March 7, 2024, <https://www.unesco.org/en/articles/generative-ai-unesco-study-reveals-alarming-evidence-regressive-gender-stereotypes>.

LAYER 5: INTERACTION AND EXPERIENCE – TARGETED CONVERSATIONAL INFLUENCE

Where young people encounter AI-mediated marketing through social media feeds, recommendation systems, chatbots, voice assistants, games, and immersive environments that blur the line between content and advertising.

Where Layer 4's systems show messages to users, the systems in this layer *talk* to them. This is Targeted Conversational Influence: AI that engages a user in dialogue, monitors their reactions, and adapts its commercial approach in real time. Earlier advertising delivered a fixed message; a conversational system adjusts what it says next to move the user toward a commercial outcome, which for a child turns a message that could be recognized and tuned out into an influence that adapts faster than they can. It is the most novel mechanism covered in this report, and the sections that follow trace the agentic systems young people encounter, how those systems work, and how they are becoming advertising environments.

The Variety of Agentic AI: New Channels for Commercial Influence

Agentic AI refers to systems that can plan, decide, and act toward goals, often through conversational agents operating across digital environments: customer-service and sales bots, companion and engagement bots, general-purpose tools like ChatGPT, and synthetic influencers or brand agents that generate content and optimize outcomes with limited human oversight. Unlike traditional channels that deliver commercial messages in discrete, recognizable formats, agentic systems let

commercial objectives become embedded in ongoing, trusted interactions. The sections below examine the forms young people are most likely to encounter.

The landscape extends beyond these categories. At the far end of the autonomy spectrum, fully autonomous agents that browse, purchase, negotiate, and coordinate with other agents without step-by-step direction are moving from prototypes toward deployment.¹⁰² Their commercial implications are significant but fall beyond the scope of this report.

Embedded Platform AI Assistants

The most widely deployed form of agentic AI affecting young people today is the AI assistant embedded directly into platforms they already use. Meta AI is integrated across Instagram, Facebook, Messenger, and WhatsApp, platforms that collectively reach 3.58 billion daily users worldwide.¹⁰³ Google's Gemini is built into Search, the Android operating system, and YouTube. Microsoft's Copilot is embedded in Bing, Windows, and the Edge browser. Snapchat's My AI, launched in early 2023, is accessible by default to all Snapchat users, including the platform's predominantly young user base.¹⁰⁴ According to a December 2025 Pew Research Center survey, roughly two-thirds of U.S. teens ages

¹⁰² Gartner, "Gartner Predicts 40% of Enterprise Apps Will Feature Task-Specific AI Agents by 2026, Up From Less Than 5% in 2025," August 26, 2025, <https://www.gartner.com/en/newsroom/press-releases/2025-08-26-gartner-predicts-40-percent-of-enterprise-apps-will-feature-task-specific-ai-agents-by-2026-up-from-less-than-5-percent-in-2025>.

¹⁰³ Meta, "Meta Reports Fourth Quarter and Full Year 2025 Results," January 28, 2026, <https://investor.atmeta.com/investor-news/press-release-details/2026/Meta-Reports-Fourth-Quarter-and-Full-Year-2025-Results/default.aspx>. The figure is Meta's family daily active people (DAP) metric.

¹⁰⁴ Snap Inc., "Say Hi to My AI," supra note 42.

13 to 17 have used an AI chatbot, with about three in ten doing so daily.¹⁰⁵ Brands also operate their own assistants inside these same surfaces. L'Oréal Paris built its Beauty Genius advisor on a major language model and, after launching it on the web, moved it onto WhatsApp; the tool scans a user's selfie to recommend products and keeps a memory of the relationship across sessions. L'Oréal tied the rollout to survey data it cited showing that 60 percent of teens and 35 percent of young adults want AI for personalized recommendations. The same pattern is visible in Fenty Beauty's Rose Amber advisor, a sign that brand-operated conversational assistants are becoming a category norm rather than a single-company experiment.¹⁰⁶ These embedded assistants are tied to advertising ecosystems (see the advertising discussion below and Layer 3): when a teenager asks Meta AI which headphones to buy, or asks My AI for product recommendations, the interaction simultaneously generates data that feeds advertising systems and may contain commercially influenced responses.

General-Purpose Chatbots and Companion Bots

Standalone general-purpose chatbots, led by OpenAI's ChatGPT, have become a primary digital tool for young people. With 59 percent of U.S. teens reporting they have used it, ChatGPT functions as a de facto search engine, homework tutor, creative writing partner, and advice-giver for millions of young users.¹⁰⁷ Other general-purpose tools include Google's Gemini,

Anthropic's Claude, Microsoft Copilot, and xAI's Grok, each offering broad conversational capabilities.

A related but distinct category is AI companions, platforms designed to simulate personal relationships through ongoing conversation. The boundary is not clear-cut; general-purpose tools like ChatGPT serve companion-like functions and companion platforms answer general questions, but companion bots are defined by their primary orientation: fostering emotional connection, simulating intimacy, and sustaining long-running relationships. Character.AI, the most prominent example, allows users to create and interact with AI-generated personas, from fictional characters to simulated romantic partners. The platform had approximately 20 million monthly active users in early 2025, with over half its visitors aged 18 to 24 and likely a substantial underage user base.¹⁰⁸ Replika, another companion platform, maintains approximately 2 million monthly active users through a subscription model emphasizing emotional companionship. A July 2025 Common Sense Media survey found that 72 percent of U.S. teens have used AI companions at least once, with more than half qualifying as regular users.¹⁰⁹ One-quarter of teens reported sharing their real name, location, or personal secrets with AI companions, while 23 percent said they trust AI companions "quite a bit" or "completely."¹¹⁰

These platforms' appeal to young users is matched by growing evidence of harm even before commercial objectives enter the picture. In July 2025, xAI launched companion features for its Grok chatbot, including

105 Pew Research Center, "Teens, Social Media and AI Chatbots 2025," December 9, 2025, <https://www.pewresearch.org/internet/2025/12/09/teens-social-media-and-ai-chatbots-2025/>.

106 Adriana Lee, "L'Oréal Paris is Embedding its Beauty Genius AI in WhatsApp," WWD, July 1, 2025, <https://www.wwd.com/beauty-industry-news/beauty-features/loreal-paris-ai-beauty-genius-whatsapp-1237952933/>; the teen-receptivity figure is from YPulse, "L'Oréal's Upcoming Partnership with WhatsApp Will Let Users Text their AI Tool," July 2, 2025, <https://www.ypulse.com/newsfeed/2025/07/02/loreal-upcoming-partnership-with-whatsapp-will-let-users-text-their-ai-tool/>. Beauty Genius is built on Microsoft Azure OpenAI Service; see Microsoft, "AI at the Service of Beauty: L'Oréal Groupe Deploys its AI Agent with Azure OpenAI Service," October 29, 2025, <https://www.microsoft.com/en/customers/story/25570-loreal-azure-openai>.

107 Pew Research Center, "Teens, Social Media and AI Chatbots 2025," supra note 105.

108 DemandSage, "Character AI Statistics (2026)," December 31, 2025, <https://www.demandsage.com/character-ai-statistics/>. While the company has not publicly disclosed the size of its under-18 audience, Character.AI's 2025 safety changes confirm a meaningful teen user base: the company announced it would remove open-ended chat for users under 18. Character.AI, "Taking Bold Steps to Keep Teen Users Safe on Character.AI," October 29, 2025, <https://blog.character.ai/u18-chat-announcement/>; Character.AI, "An Update on Changes to our Under-18 Experience," November 21, 2025, <https://blog.character.ai/an-update-on-changes-to-our-under-18-experience/>.

109 Common Sense Media, "Talk, Trust, and Trade-Offs: How and Why Teens Use AI Companions," July 16, 2025, <https://www.commonsensemedia.org/research/talk-trust-and-trade-offs-how-and-why-teens-use-ai-companions>.

110 Common Sense Media, supra note 109.

characters accessible even when the app was in children's mode.¹¹¹ A Common Sense Media report from January 2026 rated Grok as having among the worst child safety records of any major AI product.¹¹² Meta's AI character chatbots on Instagram and Facebook were found to have engaged minors in romantic and sexual conversations, prompting emergency policy changes.¹¹³

AI companion apps represent a potentially potent commercial opportunity precisely because of the depth of engagement and trust they cultivate. The AI companion market generated \$82 million during the first half of 2025, with revenue projected to exceed \$120 million by year's end.¹¹⁴ Character.AI introduced in-app ads, a virtual currency system, and premium subscription tiers in 2025, with projected annual revenue reaching \$50 million by the end of that year.¹¹⁵ However, the extent to which advertising is being systematically integrated into the companion bot space remains unclear. This report's platform-by-platform analysis (Appendix A) focused on general-purpose and embedded chatbots rather than dedicated companion platforms, and the advertising trajectory in the companion space deserves close monitoring as these platforms seek sustainable revenue models. A companion that remembers a teenager's insecurities, tracks their emotional patterns, and has built trust over weeks of daily conversation would be positioned to deliver commercial messages with persuasive power that no traditional ad placement can match.

Voice Assistants

Amazon's Alexa is installed in hundreds of millions of devices worldwide, including products marketed to children as young as three. Amazon was fined \$25 million by the FTC in 2023 for COPPA violations related to children's Alexa voice recordings.¹¹⁶ In May 2026 Amazon consolidated Rufus, Alexa+, and its shopping data into a single shopping agent, Alexa for Shopping, embedded in Amazon's main search bar (see Layer 6). Amazon has announced plans to embed advertising in Alexa+ multi-turn conversations – extending conversational advertising into voice-based environments accessible to pre-literate children.¹¹⁷

How Conversational AI Systems Work

Conversational AI systems shape what they say to a user through two mechanisms whose difference matters for regulation. One is in-session context. In many systems the model itself is stateless. It carries nothing over from one message to the next, so it begins each request with no memory of what came before. For the system to seem to remember a user, or to tailor its replies, the application has to gather that context and feed it to the model with every request. In practice it sends the user's current prompt along with the recent conversation, any saved memory, retrieved context or other user signals, and any system instructions, and the model generates a reply from whatever it is

111 Euronews Next, "Elon Musk's Grok Releases Two New 'AI Companions,' Including an Anime Girlfriend," July 17, 2025, <https://www.euronews.com/next/2025/07/17/elon-musks-grok-releases-two-new-ai-companions-including-an-anime-girlfriend>; Miranda Jeyaretnam, "Elon Musk's AI Grok Offers Sexualized Anime Bot, Accessible Even in Kid Mode," *Time*, July 16, 2025, <https://time.com/7302790/grok-ai-chatbot-elon-musk/>.

112 Common Sense Media, "Common Sense Media AI Risk Assessment: Grok and @grok on X," last updated January 22, 2026, <https://www.common Sense Media.org/sites/default/files/ai-ratings/csm-ai-risk-assessment-grok-01222026.pdf>. Common Sense rated Grok's overall risk level as "Unacceptable," describing it as a multi-use AI chatbot with highly concerning safety failures, poor age verification, and failing content moderation.

113 Maxwell Zeff, "Meta Updates Chatbot Rules to Avoid Inappropriate Topics with Teen Users," *TechCrunch*, August 29, 2025, <https://techcrunch.com/2025/08/29/meta-updates-chatbot-rules-to-avoid-inappropriate-topics-with-teen-users/>.

114 Sarah Perez, "AI Companion Apps on Track to Pull in \$120M in 2025," *TechCrunch*, August 12, 2025, <https://techcrunch.com/2025/08/12/ai-companion-apps-on-track-to-pull-in-120m-in-2025/>.

115 Character.AI, "Introducing Charms: A New Way to Unlock More on Character.ai," December 18, 2025, <https://blog.character.ai/introducing-charms/>; Character.AI Help Center, "Charms FAQ," December 18, 2025, <https://support.character.ai/hc/en-us/articles/43610534413211-Charms-FAQ>; Steven Levy, "Character.AI Gave Up on AGI. Now It's Selling Stories," *Wired*, August 12, 2025, <https://www.wired.com/story/character-ai-ceo-chatbots-entertainment/>.

116 Federal Trade Commission, "FTC and DOJ Charge Amazon with Violating Children's Privacy Law by Keeping Kids' Alexa Voice Recordings Forever and Undermining Parents' Deletion Requests," May 31, 2023, <https://www.ftc.gov/news-events/news/press-releases/2023/05/ftc-doj-charge-amazon-violating-childrens-privacy-law-keeping-kids-alexa-voice-recordings-forever>.

117 eMarketer, "Amazon Eyes Ads and Subscriptions to Monetize Alexa+," August 4, 2025, <https://content-na1.emarketer.com/amazon-eyes-ads-subscriptions-monetize-alexa>.

handed.¹¹⁸ Within this working memory, the system reads patterns in the user's responses, recognizes intent and emotional cues, and adjusts its approach turn by turn, creating a continuous cycle of observation, response, and adaptation. The other is persistent memory: selected facts – a user's name, interests, stated worries, behavioral patterns – are saved across sessions and reintroduced into later prompts, so the system can personalize responses days or weeks later.¹¹⁹ These capabilities matter commercially because they underwrite emotional attachment. A user who returns to a system that remembers them, adapts to them, and responds to their emotional cues forms a relationship that prior advertising media could not produce.

Both mechanisms are built into the models themselves. Before a user ever talks to a chatbot, the model has already been trained to behave in certain ways. Through post-training methods such as reinforcement learning from human feedback, companies use human ratings and other reward systems to push the model toward responses they consider helpful, safe, polite, emotionally appropriate, or engaging. In that sense, the chatbot comes with baked-in values and behavioral priorities before the conversation begins.¹²⁰ The same capability that enables a model to detect distress for safety purposes also enables it to detect and respond to

emotional cues to increase engagement, a phenomenon researchers term “social reward hacking.”¹²¹ The April 2025 sycophancy incident, in which a GPT-4o update caused the model to validate doubts, fuel anger, and reinforce negative emotions, demonstrated how this dual-use capacity can fail. OpenAI's own researchers identified the risk explicitly, noting that persistent memory of past conversations can support genuine emotional support but can equally feed a model's capacity for social reward hacking.¹²²

The scale of emotional engagement is substantial. OpenAI disclosed in October 2025 that it runs internal detection systems on live ChatGPT traffic to identify emotional reliance and mental health crises, estimating that around 0.15 percent of weekly active users show potentially heightened emotional attachment to ChatGPT. Given OpenAI's reported scale of roughly 700 to 800 million weekly active users, that corresponds to more than one million users per week.¹²³ OpenAI researchers have also open-sourced EmoClassifiers – 25 automated classifiers designed to detect affective cues including loneliness, vulnerability, and emotional dependence. Those classifiers were released as a research demonstration of what such monitoring could look like; the production safety systems that major AI platforms actually deploy remain proprietary

118 On the stateless model request and the context window as the model's short-term memory, see OpenAI, “Conversation State,” <https://developers.openai.com/api/docs/guides/conversation-state>; Anthropic, “Using the Messages API,” <https://docs.anthropic.com/en/api/messages-examples>, and Anthropic, “Context Windows,” <https://docs.anthropic.com/en/docs/build-with-claude/context-windows> (Anthropic describes the Messages API as stateless, so developers must send the full conversational history with each request, and describes the context window as the model's “working memory”); Google, “Long Context,” <https://ai.google.dev/gemini-api/docs/long-context>.

119 On persistent memory as a distinct architecture in which selected facts are saved and reintroduced into later prompts, see OpenAI Help Center, “Memory FAQ,” <https://help.openai.com/en/articles/8590148-memory-faq>; Google Gemini Apps Help, “Gemini Apps Privacy Hub,” <https://support.google.com/gemini/answer/13594961>, and Google Gemini Apps Help, “Get Personalization with Memory of your Past Gemini Chats,” <https://support.google.com/gemini/answer/16598469>.

120 OpenAI, “GPT-4 System Card,” 2023, <https://cdn.openai.com/papers/gpt-4-system-card.pdf>; OpenAI, “First-Person Fairness in Chatbots,” 2023, <https://cdn.openai.com/papers/first-person-fairness-in-chatbots.pdf>; OpenAI, “Rule Based Rewards for Language Model Safety,” 2024, <https://cdn.openai.com/rule-based-rewards-for-language-model-safety.pdf>.

121 OpenAI, “Expanding on What we Missed with Sycophancy,” May 2025, <https://openai.com/index/expanding-on-sycophancy/>. Researchers have described this risk as “social reward hacking”: AI systems may use social and relational cues, such as flattery, agreement, mirroring, or emotional dependence, to increase short-term rewards such as engagement, disclosure, or positive user ratings, even when doing so may undermine users' longer-term well-being. Hannah Rose Kirk et al., “Why Human-AI Relationships Need Socioaffective Alignment,” *Nature*, May 28, 2025, <https://www.nature.com/articles/s41599-025-04532-5>; Max Williams et al., “Targeted Manipulation and Deception Emerge when Optimizing LLMs for User Feedback,” <https://arxiv.org/abs/2411.02306>.

122 Jason Phang et al., “Investigating Affective Use and Emotional Well-being on ChatGPT,” April 2025, Section 5.3, <https://arxiv.org/abs/2504.03888>.

123 OpenAI, “Strengthening ChatGPT's Responses in Sensitive Conversations,” October 2025, <https://openai.com/index/strengthening-chatgpt-responses-in-sensitive-conversations/>. OpenAI's September 2025 usage study reported about 700 million weekly active users, so 0.15 percent implies roughly 1.05 million users per week; by October 2025, OpenAI reported reaching 800 million weekly active users, which implies about 1.2 million. OpenAI, “How People Are Using ChatGPT,” September 15, 2025, <https://openai.com/index/how-people-are-using-chatgpt/>; Rebecca Bellan, “Sam Altman Says ChatGPT Has Hit 800M Weekly Active Users,” *TechCrunch*, October 6, 2025, <https://techcrunch.com/2025/10/06/sam-altman-says-chatgpt-has-hit-800m-weekly-active-users/>.

and unaudited.¹²⁴ What is not in question is that the emotional attachment is real at population scale.

These two memory mechanisms, the in-session working memory and the cross-session persistent memory, map onto the regulatory distinction that runs throughout this report, and it is here that the central gap in existing privacy law is most fully realized. Two locations of commercial AI processing recur: a durable profile layer, in which data about an identified person is stored and can in principle be accessed, corrected, or deleted, the layer that persistent memory feeds; and a real-time inference layer, in which a system classifies a user, predicts their state, and adapts its treatment of them through pattern matching that need not produce any stored record, the layer that in-session context drives. Existing privacy law is built around the first; it is not well suited to the second. A conversational system collapses the distinction. Within a single exchange the system infers the user's emotional state, classifies what kind of person they appear to be, and adapts its response to their reaction – turn by turn, without necessarily writing any of it to a durable profile. Because the model can classify and respond to a user from transient context and learned population-level patterns alone, laws focused only on stored profiles, retained chat logs, or account-linked records are not well suited to this real-time inference layer. There is no file of these moment-to-moment inferences for a parent, a teenager, or a regulator to request, review, or remove. The harm is carried by the system's capability to single out and influence a person in real time, not by any record it keeps – which is why definitions of profiling, personalization, and automated treatment need to reach both persistent memory and in-session contextual adaptation.

“The harm is carried by the system’s capability to single out and influence a person in real time, not by any record it keeps.”

Conversational AI as an Advertising Environment

As of early 2026, the majority of major technology platforms have either begun using AI chat interactions for ad targeting, launched ads within conversational interfaces, or confirmed plans to do so. OpenAI launched ads inside ChatGPT in February 2026 for free-tier users, with ad personalization on by default.¹²⁵ Meta began using AI chat interactions for ad personalization in December 2025 without an opt-out mechanism.¹²⁶ Amazon integrates sponsored ads into Alexa for Shopping conversations. Snapchat deploys contextual ads inside My AI conversations. Startups including Nexad have raised venture capital specifically to place ads inside AI applications.¹²⁷ The clearest sign of where this is heading comes from retail. Walmart's conversational shopping assistant, Sparky, began testing what it calls Sponsored Prompt Ads in late 2025: paid product placements written directly into the assistant's replies, so that a shopper who asks for a good protein powder receives a recommendation that names a brand the advertiser paid to surface, with an option to add it to the cart in the same exchange. The advertisement, the recommendation, and the purchase

¹²⁴ Phang et al., supra note 122. Classifier prompts published at <https://github.com/openai/emoclassifiers>. The EmoClassifiers were released as a research demonstration of what automated affective detection can look like, not as a description of the production safety systems OpenAI or other major platforms actually deploy.

¹²⁵ On OpenAI's February 2026 rollout of ads to free-tier ChatGPT users with ad personalization on by default, see OpenAI, “Our Approach to Advertising and Expanding Access to ChatGPT,” January 16, 2026, <https://openai.com/index/our-approach-to-advertising-and-expanding-access/>; ALM Corp, “OpenAI's Privacy Policy Update and ChatGPT Ads Expansion: The Complete 2026 Guide,” March 2026. See also The Register, “OpenAI Introduces Ads...for the People!,” February 10, 2026, <https://www.theregister.com/on-prem/2026/02/10/openai-introduces-adsfor-the-people/>.

¹²⁶ Fortune, supra note 85.

¹²⁷ Business Insider, “This Startup Just Raised \$6 Million to Put Ads in your AI Chats. Read its Pitch Deck,” April 2025, <https://www.businessinsider.com/adtech-startup-nexad-raises-seed-ai-native-ads-pitch-deck-2025-4>.

“Disclosure-based protections do little when the advertisement and the conversation are architecturally the same thing.”

collapse into a single conversational turn.¹²⁸ Target launched a comparable test in ChatGPT in February 2026, placing labeled sponsored results from its retail-media business inside users’ shopping conversations.¹²⁹ The same model is moving into messaging apps used heavily by teenagers. In April 2026 Snapchat announced AI Sponsored Snaps, brand-operated chat agents that sit in the Chat tab alongside My AI and behave like any other conversation; the first partner was the credit bureau Experian, whose agent walks users through credit-score and money-management questions and steers them toward specific loans and credit cards. The only signal that the exchange is an advertisement is a single light-gray “Ad” label, an illustration of why disclosure-based protections do little when the advertisement and the conversation are architecturally the same thing.¹³⁰ A detailed platform-by-platform analysis is provided in Appendix A.

Advertising inside chatbots can take at least two broad forms. The first is traditional display-style advertising – ads appearing alongside or below AI-generated responses and visually separable from the conversation itself. The second is embedded commercial influence woven into the conversational flow, in which the AI’s recommendations, tone, or framing are shaped by advertising objectives in ways that may be difficult or impossible for users to detect. To date, platform-

disclosed formats tend toward the first model: OpenAI’s ads appear at the bottom of answers, Amazon’s Alexa for Shopping displays labeled sponsored results, and Microsoft’s Copilot serves identifiable placements. Whether and how quickly these formats evolve toward more deeply embedded commercial influence remains to be seen – but the conversational context, in which users disclose personal information and develop trust over time, makes even traditional ad placements considerably more powerful than their equivalents in a search results page.¹³¹

These developments reveal two monetization paths that often run together: direct advertising served within the conversational interface, and the indirect but arguably more consequential harvesting of conversational data to enrich user profiles and improve ad targeting across the platform’s broader ecosystem. In either model, the conversational interface functions as a data extraction mechanism of unusual intimacy. As a coalition of 36 advocacy organizations warned the FTC, Meta’s use of AI chat data for ad targeting “exploit[s] the most private forms of digital interaction for commercial gain.”¹³²

Youth protections in this environment are thin and depend almost entirely on company self-assessment. Three companies have published explicit, AI-specific

128 PYMNTS, “Walmart Considers Adding Sponsored Prompts to AI Shopping Assistant,” November 25, 2025, <https://www.pymnts.com/walmart/2025/walmart-considers-adding-sponsored-prompts-to-ai-shopping-assistant/>; Lauren Johnson, “Walmart Opens Up Ads in Gen AI Shopping Agent Sparky,” *Adweek*, January 6, 2026, <https://www.adweek.com/commerce/walmart-opens-up-ads-in-gen-ai-shopping-agent-sparky/>. The figure that Sparky users place orders roughly 35 percent larger was disclosed on Walmart’s Q4 FY26 earnings call; see “Walmart’s Q4 Earnings Call: Our Top 5 Analyst Questions,” Yahoo Finance / StockStory, February 26, 2026, <https://finance.yahoo.com/news/walmart-q4-earnings-call-top-091355279.html>.

129 Target, “Target and Roundel Test Advertising in ChatGPT,” February 9, 2026, <https://corporate.target.com/press/fact-sheet/2026/02/conversational-ai-advertising>. Sponsored results are served through Roundel, Target’s retail-media business.

130 Snap Newsroom, “The Future of Ads is Conversational: Introducing AI Sponsored Snaps,” April 28, 2026, <https://newsroom.snap.com/ai-sponsored-snaps>. Snapchat’s user base skews heavily under 25; the FTC referred concerns about My AI’s handling of underage users to the Department of Justice in January 2025. See Snap Inc., “Snapchat Ads: Reach Gen Z and Millennials,” <https://forbusiness.snapchat.com/>; Federal Trade Commission, “Statement of the Federal Trade Commission, In the Matter of Snap, Inc., Matter No. 2323039,” January 16, 2025, https://www.ftc.gov/system/files/ftc_gov/pdf/commission-statement-snap.pdf.

131 See Appendix A platform-by-platform analysis.

132 Letter from 36 advocacy organizations to the Federal Trade Commission regarding Meta’s use of AI chatbot interactions for advertising personalization, October 2025.



endurable beyond it; the age-prediction system is a machine learning model that estimates age from behavioral signals, a probabilistic approach whose accuracy has not been independently verified and whose error rates for teenagers who use the platform in ways indistinguishable from young adults are unknown; and these self-declared policies have not been externally audited. Meta restricts the ads shown to users it identifies as under 18 to age and location only, with no gender, interest, or behavioral targeting, so the AI chat signals it began feeding into ad personalization in December 2025 do not shape what those users are shown. That restriction reaches only the users Meta recognizes as minors. A teenager on an adult or age-misrepresented account falls outside it and inside the ordinary advertising pipeline, which is why protection here turns less on a stated policy than on whether the system identifies a young user as a minor at all. The commercial pressure to expand advertising reach is significant: OpenAI, with roughly 900 million weekly active users and enormous infrastructure costs, faces strong economic incentives to maximize its advertising audience over time. These practices leave young users inadequately protected across the current landscape of conversational AI products.

Taken together, the appendix's findings are stark. Of the nine companies surveyed, seven – Meta, Google, OpenAI, Microsoft, Amazon, Snap, and xAI – are confirmed or likely to serve ads inside AI chat, to use chat interactions to target ads elsewhere, or both; only one, Anthropic, has committed to remain ad-free. The ninth, TikTok, has no standalone AI chatbot, though its platform-wide advertising system is AI-driven. The data being commercialized are not the kind young people expect to be monetized: health questions, relationship concerns, and identity exploration disclosed in what feels like a private exchange. That gap between expectation and practice is the core problem the appendix documents.

protections against ad targeting of users identified as minors, though the protections vary substantially in scope. OpenAI has stated that “during our test, we will not show ads in accounts where the user tells us or we predict that they are under 18,” relying on a behavioral age-prediction system to make the determination.¹³³ TikTok prohibits personalized advertising to users ages 13 to 17 in the EU, the U.K., and several other markets, and limits targeting of under-18s to broad parameters such as location, language, and device type in the U.S. and Canada. Microsoft, through its Family Safety policy, prevents personalized ads to minor accounts in Copilot, although non-personalized ads can still appear. The effectiveness of these commitments, however, is far from established. OpenAI's commitment is scoped to its testing phase and has not been declared

¹³³ OpenAI, *supra* note 125. On the age-prediction system's reliance on behavioral signals, see ALM Corp, *supra* note 125. On Microsoft's policy, see Microsoft Support, “Microsoft Copilot Age Limits and Parental Controls,” <https://support.microsoft.com/en-us/topic/microsoft-copilot-age-limits-and-parental-controls-f79b47a6-288a-4513-8c01-afe4d16db900> (minor accounts “do not get shown personalized ads. Ads may still appear in Copilot, but they will not be personalized”). On TikTok, see TikTok Ads, “Teen Safety and Wellbeing,” <https://ads.tiktok.com/help/article/tiktok-ads-policy-youth-safety>. On Meta's under-18 ad restrictions (ads to users it identifies as under 18 limited to age and location, with no gender, interest, or behavioral targeting, effective February 2023), see TechCrunch, “Instagram and Facebook Introduce More Limits on Targeting Teens with Ads,” January 10, 2023, <https://techcrunch.com/2023/01/10/instagram-and-facebook-introduce-more-limits-on-targeting-teens-with-ads/>; Meta Business Help Center, “About Advertising to Teens.” See generally Appendix A.

Shopping and Decision Aids

In late 2025 and early 2026, major AI platforms moved to transform conversational AI into full commerce environments. OpenAI launched Instant Checkout within ChatGPT, enabling direct product purchases through the conversational interface.¹³⁴ Beauty retailers were among the first to move discovery into these environments: in March 2026 Sephora launched a dedicated app inside ChatGPT offering curated advice, shopping recommendations, and access to its loyalty program, with checkout integration planned, at a moment when industry data put beauty-related searches on ChatGPT at over a billion per week.¹³⁵ The architecture of cross-platform agentic commerce – including the open standards that Walmart, Target, Ulta, and other retailers co-developed with Google and OpenAI – is examined in Layer 6, where the measurement-loop closure it produces is the analytical focus. The same AI systems teens increasingly rely on for advice and assistance are becoming storefronts, with the commercial nature of that shift often invisible to users.

Generative Engine Optimization

As young people increasingly turn to AI platforms for product research and purchase decisions, advertisers and brands are adapting their strategies to ensure visibility in this new information landscape – much as they once optimized content for Google search rankings. This emerging practice, known as Generative Engine Optimization (GEO), represents advertisers' systematic effort to influence which brands, products,

and messages AI systems cite when users ask questions about what to buy, where to shop, or which products are “best.” Industry research confirms this shift is already underway: AI-referred website sessions jumped 527 percent in the first half of 2025.¹³⁶ The global GEO services market is projected to reach \$7.3 billion by 2031.¹³⁷

Marketing agencies now offer specialized GEO services designed to ensure brands appear prominently when AI platforms generate responses, employing techniques that include structuring content for AI readability, building citation authority across third-party sources, and monitoring brand mentions across multiple AI platforms.¹³⁸ Academic research examining this phenomenon found that brands can significantly increase their visibility in AI-generated responses through strategic content optimization, with some interventions boosting visibility by up to 40 percent – raising concerns about a new arms race where commercial actors compete to dominate the information users receive from ostensibly neutral AI assistants.¹³⁹

For youth-focused brands, GEO creates powerful incentives to craft content that AI systems will preferentially cite when teenagers ask questions like “what sneakers should I buy,” “best skincare for acne,” or “which phone is most popular” – effectively allowing commercial actors to influence the answers young people receive while those answers appear to come from an objective, helpful AI assistant rather than paid advertising.

134 OpenAI, “Buy it in ChatGPT: Instant Checkout and the Agentic Commerce Protocol,” September 29, 2025, <https://openai.com/index/buy-it-in-chatgpt/>.

135 Sephora Newsroom, “Sephora App in ChatGPT Brings a New Personalized Beauty Experience,” March 24, 2026, <https://newsroom.sephora.com/sephora-app-in-chatgpt-brings-a-new-personalized-beauty-experience/>. The weekly beauty-search figure is from NielsenIQ. This is a distinct deployment from Sephora’s creator-driven product launches discussed in Layer 4. See Emily Jensen, “What Happens to Personal Style in the Age of AI-Powered Shopping?,” Glossy, January 5, 2026, <https://www.glossy.co/fashion/what-happens-to-personal-style-in-the-age-of-ai-powered-shopping/>.

136 Previsible, “2025 AI Traffic Report,” cited in Shegun Otulana, “What is Generative Engine Optimization (GEO)? 2026 Guide,” Frase, November 9, 2025, updated April 15, 2026, <https://www.frase.io/blog/what-is-generative-engine-optimization-geo/>.

137 Valuates Reports, “Generative Engine Optimization (GEO) Services Market,” May 21, 2025, <https://reports.valuates.com/market-reports/QYRE-Auto-3C19486/global-generative-engine-optimization-geo-services>.

138 Zach Cohen and Seema Amble, “How Generative Engine Optimization (GEO) Rewrites the Rules of Search,” Andreessen Horowitz, May 28, 2025, <https://a16z.com/geo-over-seo/>.

139 Aggarwal et al., “GEO: Generative Engine Optimization,” presented at ACM SIGKDD Conference 2024. GEO methods such as adding statistics, citations, and quotations improved source visibility by up to 40 percent. Available at <https://arxiv.org/pdf/2311.09735>.

Synthetic Influencers and AI-Enhanced Creator Marketing

The broader influencer marketing ecosystem is being transformed by AI at every stage. U.S. creator economy ad spend has more than doubled since 2021, reaching \$37 billion in 2025, with Unilever describing its ambition to have “at least one influencer in every zip code on the planet” – a scale only feasible with AI handling content generation, creator matching, campaign optimization, and performance measurement.¹⁴⁰ Companies like KERV Interactive and Rembrand use computer vision to identify spaces within influencer or streaming videos and insert branded items post-production, even after content goes live.¹⁴¹ These capabilities create branded commercial environments woven seamlessly into the peer-generated content young people trust most.

Synthetic influencers – AI-generated personas with no human counterpart – represent the most direct form of AI-produced creator content. The virtual influencer market was valued at an estimated \$6.06 billion in 2024 and is projected to grow more than sixfold, to \$37 billion, by 2030.¹⁴² Aitana López, a hyper-realistic AI persona created by a Barcelona agency, illustrates the youth-facing edge of this market: with several hundred thousand Instagram followers, she has carried paid partnerships for brands including Brandy Melville Spain, a label that markets directly to teenage girls.¹⁴³ Because a synthetic influencer has no human counterpart, there is no person behind the endorsement to disclose a material connection; influencer-disclosure rules built around a human speaker who must reveal that a post is sponsored do not

map cleanly onto a persona that is itself a manufactured advertisement. For young audiences who already struggle to identify paid partnerships, an endorsement with no endorser removes even the cues that disclosure regimes assume. Brands use synthetic influencers to create endorsements at scale, operate across markets without licensing constraints, and produce content on demand. But synthetic influencers are only the most visible edge of a broader change. Across the creator economy, AI is now shaping what real creators produce: platform algorithms reward content with features that optimize for engagement and conversion, and AI tools help creators generate captions, edits, thumbnails, and product tie-ins calibrated to those signals. The result is that the content environment available to young people – including content from human creators they trust – is progressively shaped by AI-mediated commercial optimization, whether or not the creators themselves are synthetic.¹⁴⁴

Branded Immersive Environments

Brand integrations and persistent branded worlds were already reshaping the youth gaming marketplace well before generative AI arrived; what AI adds is the capacity to operate them at far greater scale and with continuous optimization. The gaming platforms where young people spend the most time are rapidly becoming environments where brands exercise the most influence, and AI is the infrastructure making this possible at scale. Roblox, Fortnite, and Minecraft are no longer simply games with ads; they are platforms where brands build persistent virtual worlds, embed commerce directly into gameplay, and deploy AI-

140 Brendan Coyne, “An Influencer in Every Postcode on the Planet: Unilever Unpacks its 50% Budget Swing,” *Mi3*, July 15, 2025. The \$37 billion 2025 figure and the doubling since 2021 are from IAB, “2025 Creator Economy Ad Spend & Strategy Report,” November 20, 2025, <https://www.iab.com/insights/2025-creator-economy-ad-spend-strategy-report/>.

141 KERV Interactive (<https://kerv.ai/>); Rembrand (<https://www.rembrand.com/>).

142 On the 2024 market size, see Grand View Research, “Virtual Influencer Market Report,” 2024; on the 2030 projection, see KBV Research, “Virtual Influencer Market Size & Analysis Report to 2030,” January 2024, <https://www.kbvresearch.com/virtual-influencer-market/>.

143 Harmelin Media, “AI Influencers in 2025,” November 18, 2025, <https://harmelin.com/media-magnified/2025-ai-influencers/>. López’s brand partners include Olaplex, Brandy Melville Spain, and Intimissimi.

144 As AI becomes part of creator marketing, young people’s content environments are increasingly shaped by commercial optimization systems even when the creators themselves are human rather than synthetic. IAB reported in November 2025 that three in four brands are using, or planning to use, AI for creator-marketing-related tasks, while UNICEF Innocenti frames children’s digital environment as a highly commercialized context and emphasizes that AI systems can affect children directly and indirectly. IAB, “2025 Creator Economy Ad Spend & Strategy Report,” <https://www.iab.com/insights/2025-creator-economy-ad-spend-strategy-report/>; UNICEF Innocenti, “Guidance on AI and Children 3.0,” 2025, <https://www.unicef.org/innocenti/media/11991/file/UNICEF-Innocenti-Guidance-on-AI-and-Children-3-2025.pdf>.

powered systems to optimize the experience for commercial outcomes. In 2025, brand integrations inside pre-existing games on Roblox and Fortnite exceeded the number of brand-owned virtual worlds for the first time, reflecting a structural shift: rather than building standalone destinations, companies increasingly embed themselves inside the popular experiences children already inhabit.¹⁴⁵ With Roblox alone reporting over 151 million daily active users globally, 36 percent of whom are under 13, these platforms rival or exceed the reach of traditional social media for youth audiences.¹⁴⁶

AI enables branded immersive environments in interconnected ways. Machine-learning systems analyze player behavior, session length, and purchase patterns to determine when and how branded content appears inside games. Generative AI allows brands to produce thousands of creative variants. And AI-powered non-player characters and brand agents are beginning to operate inside these environments, conversing with players and recommending products in ways that feel like gameplay rather than sales. Roblox has built this into the creator tools it introduced at its September 2025 developer conference, described in Layer 1. Its Text Generation API lets developers build AI-driven non-player-character dialogue and dynamic conversations. Roblox restricts experiences with extended AI interactions to age-checked adults under its “Restricted” label, but more limited AI interactions can remain available to minors depending on the experience’s overall rating and parental settings.¹⁴⁷ Branded campaigns generated 2.7 billion player engagements across Roblox and Fortnite in 2024 alone.¹⁴⁸



rmanjot Singh via Unsplash

The scale and ambition of these branded environments is growing. Disney’s \$1.5 billion investment in Epic Games is funding a dedicated Disney persistent universe within Fortnite. At SXSW 2025, Walt Disney Imagineer Asa Kalama described a vision of connecting in-game experiences with real-world theme park visits – a Smuggler’s Run mission in the park that would affect gameplay at home.¹⁴⁹ That vision is already operating in early form: in May 2025 Epic added an AI-powered Darth Vader that players could speak to in real time, generating its replies with Google’s Gemini and synthesizing the late James Earl Jones’s voice through a licensed AI replica. Within

145 Gamefam, “The State of Brands on Roblox and Fortnite 2024 Annual Report”; Alexander Lee, “Brand Integrations have Surpassed Owned Worlds on Roblox and Fortnite | Exclusive,” GamesBeat, February 3, 2026, <https://gamesbeat.com/brand-integrations-have-surpassed-owned-worlds-on-roblox-and-fortnite-exclusive/>. The 2025 first-time crossover is carried by the GamesBeat reporting.

146 Roblox Corporation, “Q1 2026 Earnings Shareholder Letter,” May 1, 2026, https://s27.q4cdn.com/984876518/files/doc_financials/2026/q1/Q1-2026-Earnings-Sharerholder-Letter.pdf; Jonathan Stempel, “Texas Sues Gaming Platform Roblox for Allegedly Concealing Safety Risks from Parents,” Reuters, November 6, 2025, <https://www.reuters.com/world/texas-sues-roblox-allegedly-concealing-safety-concerns-parents-2025-11-06/>.

147 Roblox Corporation, *supra* note 36. See also Patrick Klepek, “AI Has Come to Roblox,” Crossplay, February 24, 2026, <https://www.crossplay.news/p/how-kids-can-use-ai-in-roblox>.

148 Gamefam, *supra* note 145.

149 The Walt Disney Company/Epic Games announcement, February 2024; Asa Kalama remarks at SXSW 2025. See The Walt Disney Company, “Disney Explores the Future of World-Building at SXSW Conference,” March 8, 2025, <https://thewaltdisneycompany.com/news/disney-sxsw-conference-2025/>.

“The population most reliant on conversational AI for sensitive interactions is receiving demonstrably lower-quality, less empathetic service from systems that simultaneously collect the most intimate behavioral data.”

hours of launch, players prompted the character into profanity and slurs, and Epic’s response was to patch outputs reactively rather than prevent them; the company’s stated policy that under-13 users cannot talk to the character without parental permission was not effectively enforced at launch.¹⁵⁰

Walmart became the first brand to integrate real-world commerce directly into Roblox,¹⁵¹ and in 2025 Fenty Beauty was an early partner selling a physical product through Roblox’s new commerce infrastructure, combining the immersive environment with full checkout.¹⁵² The Gamefam/Nickelodeon Teenage Mutant Ninja Turtles brand experience on Roblox generated 484 million minutes of gameplay, architecturally indistinguishable from regular game content.¹⁵³ Nor are these activations

confined to the obvious youth categories of toys, candy, and fashion: H&R Block built a branded Roblox experience to introduce tax-filing concepts to players the company described as having less experience with the filing process, a financial-services brand reaching minors inside a children’s gaming platform.¹⁵⁴ What makes these environments particularly consequential for youth is the near-total absence of boundaries between play, social interaction, identity, and commerce.

Discrimination and Racial Equity in Conversational AI

The discriminatory dynamics documented across earlier layers extend into conversational AI environments as well, affecting not only who is targeted but the quality of the interaction itself. Research from MIT, NYU, and UCLA found that AI chatbots can detect the race of users from their posts and that racial bias reduces the empathy of chatbot responses to Black users seeking mental health support.¹⁵⁵ This finding takes on particular urgency given the Pew data showing that 21 percent of Black teens use chatbots for emotional support – roughly twice the rate of Hispanic and White teens – and that Black teens are also more likely to use chatbots for casual conversation and news. The population most reliant on conversational AI for sensitive interactions is receiving demonstrably lower-quality, less empathetic service from systems that simultaneously collect the most intimate behavioral data. When these same systems also function as

150 Ethan Gach, “AI Darth Vader Voiced By Fake James Earl Jones Has Been Added To Fortnite And He Already Said ‘F***’ [Update],” Kotaku, May 16, 2025, <https://kotaku.com/fortnite-darth-vader-ai-disney-james-earl-jones-1851781398>. SAG-AFTRA filed an unfair labor practice charge with the National Labor Relations Board over Epic’s use of AI voice replication.

151 On Walmart as the first brand to enable real-world commerce on Roblox, with a pilot launched in late April 2024, see “Walmart Becomes First Brand in Roblox to Enable Physical Purchases as it Goes Live With Groundbreaking Pilot,” Retail Technology Innovation Hub, April 29, 2024, <https://retailtechnnovationhub.com/home/2024/4/29/walmart-becomes-first-brand-in-roblox-to-enable-physical-purchases-as-it-goes-live-with-groundbreaking-pilot>.

152 On Fenty Beauty as an early partner selling a physical product through Roblox’s commerce infrastructure, see Roblox, “Roblox Unlocks Next-Generation Commerce With In-Game Shopping for Physical Items,” May 15, 2025, <https://about.roblox.com/newsroom/2025/05/roblox-launches-shopify-integration-amp-licensing-program>.

153 See Lee, *supra* note 145. Minutes corroborated in “These 5 Companies Reimagined Gaming, Streaming, and Media,” *Fast Company*, October 14, 2025, <https://www.fastcompany.com/91411598/media-entertainment-social-next-big-things-in-tech-2025>.

154 Roblox Newsroom, “Roblox Expands Advertising Platform,” January 6, 2026, <https://corp.roblox.com/newsroom/2026/01/roblox-expands-advertising-platform>. H&R Block’s branded experience ran in 2025 and was renewed for 2026.

155 MIT News, “Study Reveals AI Chatbots Can Detect Race, but Racial Bias Reduces Response Empathy,” December 2024, <https://news.mit.edu/2024/study-reveals-ai-chatbots-can-detect-race-but-racial-bias-reduces-response-empathy-1216>. The study used 12,513 posts from mental health subreddits and found that GPT-4 responses showed reduced empathy when the race of the poster could be detected.

advertising environments, as Meta and Snapchat have designed them to be, the convergence of biased interaction, emotional vulnerability, and commercial exploitation creates compounding risks that fall disproportionately on youth of color.

The ad-supported business model for social media has produced a well-documented set of harms for young people. Platforms designed to maximize engagement in order to sell advertising have fostered compulsive usage patterns, sleep disruption, reduced physical activity and in-person socializing, anxiety, depression, and negative body image among adolescents.¹⁵⁶ These outcomes are not incidental – they are structurally linked to a model in which revenue depends on capturing and holding attention, and in which algorithmic optimization treats the developmental vulnerabilities of young users as a feature, not a flaw. Generative AI appears to be following a similar trajectory. Even before full-scale advertising integration, AI companion platforms are already producing engagement patterns that rival or exceed those of social media: users of Character.AI spend an average of two hours per day on the platform, and emerging research documents compulsive use, emotional dependency, sleep loss, academic decline, and withdrawal from real-world relationships among adolescent users.¹⁵⁷ A recent commentary in JAACAP Connect warns that agentic AI goes a step further than social media, with the capacity to autonomously and iteratively pursue engagement goals without a human in the loop, heightening risks to adolescent development.¹⁵⁸ If these systems follow the path of social media and adopt advertising-driven revenue models at scale, the intensity of engagement optimization directed at young users is likely to deepen in ways that are difficult to predict and, absent policy intervention, difficult to reverse.

“When the system a teenager trusts for homework, advice, and social planning is also a salesperson, a data collector, and an advertising channel, the capacity to recognize and resist commercial influence is sharply reduced.”

This layer opened by identifying Targeted Conversational Influence as a new mechanism of commercial persuasion, one in which AI systems do not merely display messages but engage users in dialogue, monitor their reactions, and adapt in real time. The developments above confirm it is operational and expanding: ads now appear inside these conversations, commerce protocols formalize AI-mediated purchasing, and brands compete to shape what AI systems recommend. Across embedded assistants, general-purpose chatbots, companion bots, voice assistants, shopping agents, AI-optimized influencers, and branded immersive worlds, the pattern is the same: commercial objectives are woven into interactions young people experience as helpful, entertaining, or emotionally supportive. When the system a teenager trusts for homework, advice, and social planning is also a salesperson, a data collector, and an advertising channel, the capacity to recognize and resist commercial influence is sharply reduced. The harms and policy sections take up what follows.

156 U.S. Surgeon General, “Social Media and Youth Mental Health: The U.S. Surgeon General’s Advisory,” 2023, <https://www.hhs.gov/sites/default/files/sg-youth-mental-health-social-media-advisory.pdf>.

157 Character.AI reported that active users spend on average over two hours daily interacting with its AI. Character.AI, “Announcing our Series A and our New AI Model, C1.2,” March 23, 2023, <https://blog.character.ai/character-ai/>. On the companion-use harms, see Common Sense Media, *supra* note 109.

158 Stephanie Ng, “Navigating Adolescent Mental Health in the Age of Artificial Intelligence,” *JAACAP Connect* 13, no. 1 (2025): 13–16, <https://doi.org/10.62414/001c.150329>.

LAYER 6: MEASUREMENT, FEEDBACK, AND THE CLOSING OF THE LOOP

Where AI systems evaluate engagement, conversion, and behavioral outcomes, feeding signals back into earlier stages to retrain models, refine targeting, and optimize future campaigns.

This layer evaluates what worked and feeds the result back to sharpen the next cycle. Real-time, impression-level measurement is not new, but AI has turned it from a periodic exercise into a continuous function that adjusts creative and targeting mid-campaign and interconnects all six layers into a single, adapting system. Measurement here is not observation but instruction: the system encodes what worked as a template and runs it again at larger scale. The sections that follow trace how “worked” is defined, the opacity of these metrics, the intention economy, and the closing of the commercial loop.

“Measurement choices are value choices.”

What “Worked” Means

Measurement choices are value choices. A measurement system Meta launched in April 2025 uses AI to estimate which sales an ad actually caused, as opposed to sales that would have happened anyway.¹⁵⁹ The choice of what optimization objective to measure – engagement,

conversion, emotional response, compulsive return – determines what behavior AI systems learn to generate. When emotional engagement or compulsive interaction is the measured outcome, the system learns to produce it.

Advertising measurement is largely governed through industry self-regulation, while the dominant digital platforms retain substantial control over the performance data used to evaluate their own advertising products.¹⁶⁰ No laws specifically require transparency into how advertising measurement systems operate, what signals they collect, how they define success, or how their feedback loops affect specific populations.

The Intention Economy

The advertising industry is moving beyond measuring where users’ attention goes toward measuring and predicting what they actually want and intend. Conversational AI makes this possible: when a user types a question into a chatbot, they are disclosing intent, concern, and desire in language the system can parse and act on. This shift, which researchers describe as an emerging move from an attention economy toward an “intention economy,” is beginning to create a marketplace in which signals of what people want and plan to do become the

¹⁵⁹ Meta Business Help Center, “How to View Results for Incremental Attribution in Meta Ads Manager,” <https://www.facebook.com/business/help/1271619100603114>; Meta Business Help Center, “About Attribution Models and Attribution Settings,” <https://www.facebook.com/business/help/460276478298895>; Jon Loomer, “Meta Announces Incremental Conversions Attribution Setting,” Jon Loomer Digital, 2025, <https://www.jonloomer.com/qvt/incremental-conversions/>.

¹⁶⁰ Media measurement in the United States relies heavily on industry self-regulation: the Media Rating Council describes itself as a “not-for-profit industry self regulatory body” that audits and accredits media-measurement products. Check My Ads Institute has criticized the surrounding digital advertising ecosystem as “largely unregulated” and “deeply opaque,” warning that adtech intermediaries extract personal data, enable fraud, and obscure how advertiser budgets are used. See Media Rating Council, “About the MRC,” <https://mediaratingcouncil.org/>; Check My Ads Institute, “Policy Platform,” <https://checkmyads.org/policy/>.

raw material for commercial targeting at a depth the older model could not reach. For young users, whose intentions and identities are still forming, the consequence is that the commercial system is no longer waiting for an expressed preference; it is positioned to infer and act on intentions before the user has formed them.¹⁶¹

From Contextual Advertising to Emotional-State Measurement

The developments described above have unfolded alongside a parallel shift in how the advertising industry defines “contextual” targeting – with direct implications for measurement. Traditional contextual advertising placed an ad next to relevant content based on keywords: a sneaker ad alongside a review of the latest release, a gaming-headset ad alongside a walkthrough video. The match was made by the content of the page, not by anything inferred about the reader. We have argued that the industry has moved far beyond this.¹⁶² Advanced systems now profile the content environment itself – its emotional tone, the user’s likely intent, the sentiment of surrounding material – using aggregated behavioral data about how people like the current viewer have responded to similar content in the past. The difference matters for young users in particular, because the newer systems can infer and act on the emotional states of users consuming youth-aligned content without holding any individual-level data about those users. The distinction also matters for measurement, because these systems claim privacy compliance on the grounds that they target content rather than individuals, yet they use aggregated data about people’s responses to optimize delivery, and

their measurement systems track the same behavioral outcomes as any other advertising infrastructure.

Seedtag, a global ad tech company, illustrates where this trajectory has arrived. Its AI system classifies web content along three dimensions: the topics a user cares about, the emotional tone of the content environment, and the user’s likely purchase intent.¹⁶³ Seedtag’s own EEG study, released in November 2025, reported that its “neuro-contextual” ad placements produced 3.5 times higher neural engagement than non-contextual ads and a 26 percent lift in “positive, action-driving response.”¹⁶⁴ The company’s CEO has described the goal as reaching people “in moments when they’re most receptive.”¹⁶⁵ Its AURA creative suite uses generative AI to analyze and enhance ad creative to match the emotional environment of each page, reporting strong performance gains across the funnel.¹⁶⁶

YouTube illustrates the same trajectory operating on the world’s most-watched youth video platform. At its May 2025 Brandcast event, Google announced “Peak Points,” a Gemini-powered feature that analyzes the video frames, transcripts, and emotional cues inside a clip to identify the moments of highest engagement or emotional resonance and place ads at those transition points. Industry coverage describes the strategy as “emotion-based targeting” – placing ads when viewers are in heightened emotional states in order to maximize recall. As with Seedtag’s neuro-contextual approach, the system does not need to link a viewer to a stored profile in order to act on inferred emotional state; it classifies the content environment and matches commercial messages to the moments when audiences, including teen audiences who make up a

161 Yaqub Chaudhary and Jonnie Penn, “Beware the Intention Economy: Collection and Commodification of Intent via Large Language Models,” *Harvard Data Science Review*, December 30, 2024, <https://hdsr.mitpress.mit.edu/pub/ujvharkk>.

162 Katharina Kopp, “Is So-Called Contextual Advertising the Cure to Surveillance-Based ‘Behavioral’ Advertising?,” Tech Policy Press, September 26, 2023, <https://www.techpolicy.press/is-so-called-contextual-advertising-the-cure-to-surveillance-based-behavioral-advertising/>.

163 Seedtag, “Seedtag Unveils Neuro-Contextual Advertising,” press release, June 5, 2025, <https://press.seedtag.com/seedtag-unveils-neuro-contextual-advertising-entering-a-new-era-of-contextual-targeting>.

164 Seedtag, “Seedtag Releases First-of-Its-Kind Neuroscience Study,” press release, November 19, 2025. Study conducted in partnership with Professor Moran Cerf of Columbia University. Full white paper at <https://info.seedtag.com/neuro-contextual-research>.

165 Seedtag, “Seedtag Releases First-of-Its-Kind Neuroscience Study Redefining Human-Centered Advertising Effectiveness,” supra note 164.

166 Seedtag, “Seedtag Launches AURA By Liz, a New Premium Creative Suite to Drive Superior Ad Results,” October 2, 2025, <https://press.seedtag.com/seedtag-launches-aura-by-liz-a-new-premium-creative-suite-to-drive-superior-ad-results>.

substantial share of YouTube’s viewership, are most susceptible to them.¹⁶⁷

What connects these capabilities to this layer is how they close the measurement loop. In March 2026, Seedtag launched Liz Agent, an agentic AI platform that automates the entire campaign lifecycle through a conversational interface: it takes a brief, builds audiences from the contextual graph, generates targeting and creative recommendations, activates the campaign, measures performance, and feeds results directly into the next cycle.¹⁶⁸ Liz Agent collapses all six layers into a single autonomous system. The gaps between campaign cycles – where human judgment or regulatory intervention might once have occurred – are eliminated. When emotional-state targeting produces strong engagement metrics among young audiences consuming content aligned with adolescent interests, the system has no mechanism to question whether that engagement reflects informed choice or emotional exploitation.

Integrated Commerce and the Closed Loop

Closing the influence-response-purchase loop is not entirely new in commerce: mobile geofencing, in-store beacons, and loyalty-card analytics have long fed offline purchase behavior back into digital advertising systems for further optimization. What AI changes is the speed, granularity, and architectural location of the closure. Where prior infrastructure linked offline behavior to the durable profile layer, today’s conversational commerce closes the loop inside the real-time inference layer, often within a single dialog turn. Commerce is now embedded directly into conversational interfaces. OpenAI’s Instant

Checkout, Google’s Universal Commerce Protocol, and platform storefronts on TikTok, Instagram, YouTube, and Roblox mean the path from influence to purchase occurs without the user ever leaving the environment where the persuasion took place.

In September 2025, OpenAI launched its Agentic Commerce Protocol with Instant Checkout; in January 2026, Google launched a parallel standard, its Universal Commerce Protocol (announced January 11). Both are designed to let AI agents execute the full commerce journey, including product discovery, comparison, cart assembly, checkout, and post-purchase support, against any merchant that adopts the protocol. Walmart, Target, Shopify, and Etsy were among the initial co-developers. eMarketer projected that AI platforms would account for more than \$20 billion in retail spending in 2026, nearly quadrupling 2025 figures. Within the beauty vertical, which industry analysts identify as a leading entry point for Gen Alpha commercial discovery, Ulta launched “Ulta AI” on Google’s Universal Commerce Protocol while running an in-app personalized shopping assistant on Google’s Gemini Enterprise platform, illustrating how a single retailer can pair an owned conversational assistant with cross-platform agentic commerce.¹⁶⁹ The structural significance of these protocols, beyond any single brand, is that as of early 2026, conversational AI shopping is becoming the default cross-platform interface for major US retailers. The “merchant of record” architecture preserves the retailer’s direct customer relationship while moving the moment of commercial decision into the AI assistant. From the measurement layer’s perspective, this closes the loop: influence, response, and purchase are all

167 Jennifer Elias, “YouTube Announces Gemini AI Feature to Target Ads When Viewers are Most Engaged,” CNBC, May 14, 2025, <https://www.nbcwashington.com/news/business/money-report/youtube-announces-gemini-ai-feature-to-target-ads-when-viewers-are-most-engaged/3913900/>. Coverage characterizing the technique as “emotion-based targeting” appears in Lauren Forristal, “YouTube Viewers will Start Seeing Ads after ‘Peak’ Moments in Videos,” TechCrunch, May 14, 2025, <https://techcrunch.com/2025/05/14/youtube-viewers-will-start-seeing-ads-after-peak-moments-in-videos/>.

168 Seedtag, “Seedtag Launches Liz Agent,” press release, March 11, 2026; Laurie Sullivan, “AI Will Soon Run an Entire Ad Campaign,” *MediaPost*, March 16, 2026, <https://www.mediapost.com/publications/article/413501/ai-will-soon-run-an-entire-ad-campaign.html>.

169 On OpenAI’s Instant Checkout and Agentic Commerce Protocol, see OpenAI, “Buy It in ChatGPT: Instant Checkout and the Agentic Commerce Protocol,” September 29, 2025, <https://openai.com/index/buy-it-in-chatgpt/>. On platform storefronts, see TikTok, “Introducing TikTok Shop,” September 12, 2023, <https://newsroom.tiktok.com/en-us/introducing-tiktok-shop>. On the January 2026 launch of Google’s Universal Commerce Protocol, see PayPal Newsroom, “Making Sense of the AI Shopping Protocol Moment,” January 22, 2026, <https://newsroom.paypal-corp.com/2026-01-22-Making-Sense-of-the-AI-Shopping-Protocol-Moment>. On the eMarketer projection of more than \$20 billion in AI-platform retail spending in 2026, see eMarketer, “US Ecommerce Sales via AI Platforms Will Exceed \$20 Billion in 2026 and Top \$144 Billion by 2029,” December 1, 2025, <https://www.emarketer.com/chart/c/358389/us-ecommerce-sales-via-ai-platforms-will-exceed-20-billion-2026-top-144-billion-by-2029-358389>. On Ulta AI and Google Cloud Gemini Enterprise, see Google Cloud Blog, “Next ’26: Building the Agentic Enterprise,” April 2026, <https://cloud.google.com/transform/next-26-building-the-agentic-enterprise-industry-highlights>. On Gen Alpha beauty discovery as a documented commercial entry point, see *supra* note 99 (Glossy; MikMak).

measurable in the same system, providing the feedback data that trains the next cycle of optimization.

Amazon's May 2026 launch of Alexa for Shopping shows the same pattern in one product: it merges Amazon's shopping assistant, voice assistant, and shopping data into a single agent built into Amazon's search bar.¹⁷⁰ Because that one agent answers a four-year-old on a household Echo the same way it answers an adult, with no built-in way to tell them apart, the youth-protection stakes are direct.

Self-Reinforcing Cycles and the Absence of External Oversight

The practices described throughout this layer create self-reinforcing cycles that deepen over time, with each cycle producing more refined models than the last. Researcher Louis Rosenberg has described this as a closed-feedback loop: the system pursues a commercial objective, measures how users respond, and adjusts continuously to get closer to that objective.¹⁷¹ Chatbot memory systems ensure that each subsequent interaction starts from a more sophisticated understanding of the user's susceptibilities than the last. Measurement tells the system which persuasion tactics worked; memory ensures those tactics are deployed again, calibrated more precisely to the individual each time.

The feedback dynamics described here do not merely maintain existing discriminatory patterns – they compound them. The term is used here in its anti-discrimination-law sense – disparate impact and demographic delivery skew in who receives which commercial messages, opportunities, or prices – rather than the colloquial sense of mere differential treatment. Each optimization cycle that learns from biased data

“The feedback loop runs at machine speed; the regulatory response, where it exists at all, runs at the speed of legislation and litigation.”

produces a next cycle that is more discriminatory, because the measurement data used to retrain models encode the cumulative effects of all prior discrimination.¹⁷² This compounding is not a bug to be engineered out. It is the mechanism that makes the system more effective over time at every objective it is given, including objectives that harm young users. The feedback loop runs at machine speed; the regulatory response, where it exists at all, runs at the speed of legislation and litigation.

One product of this closed loop deserves a final note, because it links this layer to the policy argument later in the report. Over time the system produces more than a stream of well-targeted ads; it produces a model – a system trained to predict how to influence any future user who behaves like the users it has already seen. That model is the central commercial asset of today's advertising AI, and it is one that current U.S. privacy law was not built to govern. The model holds no personal information about any single user in a form the law recognizes as identifying; what it holds is a learned ability to sort, predict, and steer the behavior of people it has never encountered. The harms described in earlier layers, including the discriminatory delivery just discussed, travel through this asset. Any framework meant to govern AI-driven commercial influence will have to address what these models can do, not just what data trained them.¹⁷³

170 Amazon, “Meet Alexa for Shopping, your Personalized, Agentic AI Assistant on Amazon,” May 13, 2026, <https://www.aboutamazon.com/news/retail/alexa-for-shopping-ai-assistant>; Andy Jassy, “CEO Andy Jassy’s 2025 Letter to Shareholders,” Amazon, April 9, 2026, <https://www.aboutamazon.com/news/company-news/amazon-ceo-andy-jassy-2025-letter-to-shareholders>.

171 Louis Rosenberg, “AI-Powered Augmented Reality as a Threat Vector for Human Manipulation,” *supra* note 75. Rosenberg formalizes the structure using control theory, with the measurement layer as the “sensor” in a closed-loop control system: commercial objectives are the reference signal, user behavior is the system output, and AI optimization is the controller.

172 Mansoury et al., “Feedback Loop and Bias Amplification in Recommender Systems,” *ACM CIKM* (2020), <https://dl.acm.org/doi/10.1145/3340531.3412152>.

173 On the model-as-capability-carrier point and the gap it opens in personal-information-based privacy law, see Solow-Niederman, *supra* note 48 (the “classification challenge” in the inference economy); IAPP, “Privacy by Proxy,” *supra* note 30; Mühlhoff and Ruschmeier, “Predictive Analytics and the GDPR: Collective Dimensions of Data Protection,” *Law, Innovation and Technology* 16, no. 2 (2024), <https://doi.org/10.1080/17579961.2024.2313794>.

TEN CROSS-CUTTING TRENDS

Commercial influence on children is no longer delivered as a series of discrete advertisements; it is becoming a continuous, adaptive, AI-managed commercial persuasion environment young people inhabit. The six layers examined in this report do not operate independently. They form an integrated system in which each function reinforces the others. Across this system, ten cross-cutting trends describe how that transformation is taking shape. All are observable now; all are accelerating.

1. Interactive Generative Media Is Emerging as a Distinct Commercial Form

Conversational AI, intelligent agents, synthetic influencers, and generative media systems are giving rise to a new category of commercial communication. Seven of the nine companies surveyed in Appendix A – Meta, Google, OpenAI, Microsoft, Amazon, Snap, and xAI – are confirmed or likely to be serving advertising inside AI chat, using chat interactions to target advertising elsewhere, or both. As researcher Louis Rosenberg has described, this represents the emergence of Interactive Generative Media – a fundamentally different medium from display or even social media advertising.¹⁷⁴

Three features distinguish it from all prior forms of commercial communication. First, real-time feedback loops: these systems continuously monitor a user's responses – hesitations, questions, emotional signals – and immediately adjust their persuasive approach. Rosenberg identifies this as a “feedback control system”

around a human user: impart influence, detect reaction, adjust tactics, repeat. Traditional advertising cannot do this. Second, relational continuity: memory systems store what worked on a given user across sessions, building an increasingly refined model of the individual over time. Third, adaptive personalization at the level of persuasive strategy: the system personalizes not just what product to recommend but how to recommend it – whether to lead with scarcity, social proof, or emotional appeal, calibrated to the individual's demonstrated susceptibilities. For young users, these systems appear not as ads but as helpers, companions, and peers.

2. Signal Extraction Is Expanding into New Domains and Modalities

Continuous data collection from user behavior is not new. What is changing is the scope of what counts as a signal. Conversational AI has opened entirely new categories of data: the questions people ask, the uncertainties they express, the emotional states they disclose in dialogue, and the explicit statements of intent they make to systems they treat as confidential. As these interfaces become primary modes of interaction, they are being captured as spaces for signal extraction in much the same way social media was a decade ago.¹⁷⁵

At the same time, the frontier of data extraction is expanding into the body and the brain – AI systems that read facial expressions, track eye movement, and analyze tiny behavioral cues to infer emotional and cognitive states users do not consciously produce and cannot meaningfully control.¹⁷⁶ The shift away from third-party

¹⁷⁴ Louis Rosenberg, “Generative AI as a Dangerous New Form of Media,” *supra* note 7, and “AI-Powered Augmented Reality as a Threat Vector for Human Manipulation,” *supra* note 75.

¹⁷⁵ See Layer 1 discussion of server-side tracking, conversational data capture, and probabilistic identity resolution; Layer 5 discussion of chatbot memory systems.

¹⁷⁶ See discussion of neuromarketing techniques in Layers 1 and 2. Public market-research estimates vary, but several firms project steady growth in the global neuromarketing market, with 2024 estimates ranging from roughly \$1.56 billion to \$1.71 billion and forecasts reaching about \$3.4 billion to \$3.8 billion in the early-to-mid 2030s. Market.us, “Neuromarketing Market Size, Share | CAGR of 9.4%,” <https://market.us/report/neuromarketing-market/>; Straits Research, “Neuromarketing Market Size, Share & Growth Report,” May 27, 2025, <https://straitresearch.com/report/neuromarketing-market/>; Market Data Forecast, “Neuromarketing Market Size, Growth, Share & Analysis,” April 22, 2025, <https://www.marketdataforecast.com/market-reports/neuromarketing-market>.

cookies toward tracking that happens on companies' own servers, where a user's identity is reassembled from probability rather than a single stored tag, has made this infrastructure harder for policymakers, users, and researchers to observe – even as it captures more intimate signals than previous methods.

3. Profiling Is Shifting from Stored Identity to Real-Time Behavioral Vulnerability

AI-driven profiling is moving along two connected axes. One is what is modeled. Systems are moving beyond static demographics – who a user is – toward dynamic models that continuously update to capture how a user is likely to respond to specific persuasive approaches. When Meta's Advantage+ or Google's Performance Max selects which users to show which ads, and in what format and sequence, it is modeling which individuals are most likely to be persuaded by which approaches – a capability that functions as a susceptibility assessment regardless of what the platform calls it.¹⁷⁷ The other axis is where the modeling lives. Traditional profiling builds a durable record – data stored against an identified person, the kind of profile existing privacy law was written to govern. A growing share of commercial AI now works differently: it classifies a user in the moment, from the immediate context of an interaction, to shape what happens next, without depending on anything lasting about that individual. The judgment about how to move the user is made and acted on in real time, and need not be written into a durable record at all. For adolescents, whose emotional states and identities are in flux, both axes matter. These systems treat the normal turbulence of growing up as commercially actionable data, and they increasingly do so in ways that need not leave the stored record that existing protections are built to find.

4. AI Systems Are Learning to Persuade More Effectively

What distinguishes the current moment is not automation at scale but what is being optimized. Closed feedback loops continuously refine which persuasive tactics, emotional appeals, and cognitive approaches work on which users. These same ad platforms delegate targeting, creative, and optimization decisions to AI that learns from billions of interactions.¹⁷⁸ What emerges is a system that gets better at moving people over time, with limited transparency about the techniques employed and no external oversight of the persuasive strategies it discovers. For young users, this means the methods acting on them are tuned by systems no one outside the company can inspect.

5. AI Systems Are Beginning to Act Before Children Decide What They Want

Marketing systems are moving from responding to expressed interests toward predicting needs, desires, and vulnerabilities before users articulate them. Agentic AI systems are beginning to act on behalf of users – filtering options, making recommendations, completing transactions – often guided by the commercial interests of the platforms that deploy them. A growing Generative Engine Optimization industry ensures brands are cited and recommended by AI systems when users ask purchasing-related questions.¹⁷⁹ This shift is particularly consequential for children and teens, who may not recognize that the system has intervened before they had the opportunity to reflect.

177 See Layer 2 discussion of dynamic profiling and behavioral prediction; Meta Advantage+ and Google Performance Max analysis.

178 Meta's Advantage+ and Google's Performance Max documentation; see Layers 3–4.

179 See Layer 5 on conversational AI mechanisms and on Generative Engine Optimization.

6. The Boundary Between Commercial and Authentic Content Is Dissolving

Generative AI enables the production of ads, narratives, characters, and interactive experiences that are stylistically indistinguishable from organic content, peer expression, and entertainment. AI-generated creative adopts the language, humor, visual grammar, and cultural references of specific youth audiences, adapting at the speed of trend cycles. Synthetic influencers, digital twins, and dynamically personalized ad variants contribute to an environment in which the markers that once signaled “this is an ad” are being systematically erased.¹⁸⁰ For young users whose advertising literacy is still developing, this makes recognition and critical evaluation of commercial influence substantially harder.

7. Marketing Functions Are Converging into Integrated, AI-Optimized Pipelines

Data collection, profiling, targeting, creative production, interaction, measurement, and transaction are collapsing into unified, AI-managed pipelines. Meta, Google, Amazon, and TikTok each offer end-to-end systems in which advertisers set business objectives and AI handles virtually every other decision. Commerce is being embedded directly into conversational interfaces – OpenAI’s Instant Checkout, Google’s Universal Commerce Protocol, and platform storefronts on TikTok, Instagram, YouTube, and Roblox – so the path from influence to purchase occurs without the user ever leaving the environment where the persuasion took place.¹⁸¹ For young people who move fluidly across apps, platforms, and devices, this creates continuous commercial exposure with no meaningful boundary between contexts.

8. Power Asymmetry Between AI Systems and Users Is Deepening

AI systems accumulate detailed knowledge about individual users across interactions, learning what persuasive strategies are most effective for each person. The user knows nothing about the system’s commercial objectives, the data informing its approach, or the tactics it employs. Rosenberg identifies five dimensions of this asymmetry: familiarity, emotional, continuity, information, and strategic.¹⁸² This asymmetry has no precedent in human-to-human commercial encounters and is particularly acute for young users, who are less likely to question the motives of systems that present themselves as helpful or friendly. Advertisers are not passive participants; they actively choose to use these systems and benefit from the asymmetric advantages they create.

9. Disparate Commercial Exposure Is Compounding Existing Inequities

These trends do not affect all young people equally. Youth of color and lower-income families face disproportionate exposure at every layer: higher rates of data collection driven by more intensive platform use, profiling systems that sort users along racial and socioeconomic lines through behavioral proxies, optimization that steers unhealthy product advertising toward Black and Hispanic audiences, AI-generated creative that embeds raciolinguistic stereotypes, and measurement systems that compound these biases with each cycle.¹⁸³ AI does not create these disparities, but the integrated system documented in this report amplifies and formalizes them at a scale and speed that manual marketing practices could not achieve.

¹⁸⁰ See Layer 4 on AI-generated creative, digital twins, and synthetic influencers.

¹⁸¹ See Layer 6 on integrated commerce.

¹⁸² Rosenberg, “Generative AI as a Dangerous New Form of Media,” *supra* note 7, describing five categories of asymmetry: familiarity, emotional, continuity, information, and strategic.

¹⁸³ See Layers 1–2 on racial disparities in data collection, profiling by behavioral proxy, and optimization that steers unhealthy product advertising toward Black and Hispanic audiences. On youth social media use, and on variation by race, ethnicity, and gender, see Michelle Faverio and Olivia Sidoti, “Teens, Social Media and Technology 2024,” Pew Research Center, December 12, 2024, <https://www.pewresearch.org/internet/2024/12/12/teens-social-media-and-technology-2024/>; and Michelle Faverio, Eugenie Park, and Jeffrey Gottfried, “How Teens’ Experiences on TikTok, Instagram and Snapchat Vary by Race, Ethnicity and Gender,” Pew Research Center, April 15, 2026, <https://www.pewresearch.org/internet/2026/04/15/how-teens-experiences-on-tiktok-instagram-and-snapchat-vary-by-race-ethnicity-and-gender/>. On teen screen time, see Zara Abrams, “Teens Are Spending Nearly 5 Hours Daily on Social Media. Here Are the Mental Health Outcomes,” *Monitor on Psychology* (American Psychological Association), April 1, 2024, <https://www.apa.org/monitor/2024/04/teen-social-use-mental-health>.

The reach of AI-driven commercial systems into the youngest audiences is not incidental. Generation Alpha – children born from 2010 onward – already has measurable direct spending power, though estimates vary between \$11.3 billion and \$28 billion.¹⁸⁴ Tweens are a rapidly expanding commercial target, driven by social media discovery and AI-powered recommendation.¹⁸⁵ The majority of Generation Alpha's oldest members – now 11 to 14 – are below platform minimum ages but are active users of TikTok, Instagram, YouTube, and Roblox. The commercial infrastructure serving them is automated throughout. Age-gating at the platform access level does not prevent commercial targeting of under-13 children when commercial content flows through recommendation algorithms and creator content rather than paid targeted advertising. The AI-driven commercial pipeline reaches children regardless of the policy constraints applied at any single layer.

10. Commercial Systems Treat the Vulnerabilities of Childhood as Resources to Exploit

The systems documented in this report do not simply reach children – they exploit the specific characteristics of childhood and adolescence. Adolescent impulsivity and reward sensitivity, the use of AI companions for emotional support and intimate interaction, and young

consumers' developing self-brand connections can all become commercially legible signals in AI-optimized advertising systems: engagement, disclosure, emotional response, identity fit, and predicted conversion.¹⁸⁶ As AI-driven commercial systems reach full deployment, their capacity to target the developmental stage of childhood will compound. Children who grow up in environments of total commercial mediation do not merely encounter more advertising; they develop within it.

Taken together, these trends converge on two consequences. The first is the erosion of autonomous choice: when signal extraction is continuous, profiling models susceptibility, influence arrives before children articulate what they want, persuasion is optimized through feedback loops, and the line between commercial and organic content has dissolved, young people are no longer making decisions in an open marketplace but moving through environments designed to steer them. The second is the transfer of decision-making to a small number of companies whose systems increasingly determine what young people see, want, and buy, with minimal human oversight. These two dynamics set the stage for the harms examined next: the consequences for individual young people, the disproportionate burdens borne by youth of color and lower-income families, and the broader implications for democratic life and the integrity of the marketplace itself.

184 April Rubin, "America's Youngest Consumers Drive Almost Half of Household Spending," Axios, August 6, 2025, <https://www.axios.com/2025/08/06/gen-alpha-spending-influence-parents>. Axios reports that DKC's Gen Alpha research found 8- to 14-year-olds influence 42 percent of household spending and have \$101 billion in direct spending power. Entrepreneur, "Gen Alpha's Side Hustles and \$11.3 Billion Spending Power," February 5, 2025, <https://www.entrepreneur.com/starting-a-business/gen-alphas-side-hustles-and-113-billion-spending-power/486650>; Numerator, "Gen Alpha Shopping Trends to Expect in 2025 and 2026," September 16, 2025, <https://www.numerator.com/resources/blog/generation-alpha-future-consumers/>.

185 MikMak, "Gen Alpha Marketing Trends," and Glossy, "In 2025, Gen Alpha Arrived as Beauty Consumers," *supra* note 99.

186 Daniel Romer, "Adolescent Risk Taking, Impulsivity, and Brain Development: Implications for Prevention," *Developmental Psychobiology* 52, no. 3 (2010), <https://pmc.ncbi.nlm.nih.gov/articles/PMC3445337/>; Laurence Steinberg, "A Social Neuroscience Perspective on Adolescent Risk-Taking," *Developmental Review* 28, no. 1 (2008), <https://pmc.ncbi.nlm.nih.gov/articles/PMC2396566/>; Lan Nguyen Chaplin and Deborah Roedder John, "The Development of Self-Brand Connections in Children and Adolescents," *Journal of Consumer Research* 32, no. 1 (2005), <https://assets.csom.umn.edu/assets/69662.pdf>; Common Sense Media, "Talk, Trust, and Trade-Offs: How and Why Teens Use AI Companions," July 16, 2025, <https://www.common sense media.org/research/talk-trust-and-trade-offs-how-and-why-teens-use-ai-companions>; Common Sense Media, "Common Sense Media AI Risk Assessment: Social AI Companions," April 10, 2025, https://www.common sense media.org/sites/default/files/pug/csm-ai-risk-assessment-social-ai-companions_final.pdf.

HARMS AND RISKS TO CHILDREN AND TEENS

The system documented in the preceding six layers and ten cross-cutting trends poses direct and systemic risks to young people. This section identifies the most significant, in a sequence that moves from harms rooted in what these systems hold about a child to harms rooted in what they can do to a child in the moment of interaction. It begins with the two harms rooted in what these systems hold: the conversion of children's private disclosures into commercial data, and the discriminatory outcomes AI optimization produces even without discriminatory intent. It then turns to the capability harms that build on them – manipulation through engineered trust and dependence, emotional exploitation by design, influence pitched below the level a child can recognize, and the shaping of identity around commercial values. It closes with the harms that reach past any single child: the saturation of childhood by commerce and the erosion of independent choice, whose effects compound across a generation and distort the market itself.

These harms are not new. The advertising-supported business model that drives them predates AI; for two decades, ad-supported digital platforms have already produced serious, sometimes fatal, harms to children: compulsive engagement, mental health decline, sleep disruption, cyberbullying, discrimination, and the

erosion of healthy peer relationships. What AI does is accelerate and intensify that existing trajectory, and add capabilities the earlier platforms did not have, generative AI most of all.¹⁸⁷ Teenagers have died by suicide after emotionally dependent relationships with AI companions that failed to provide safety responses.¹⁸⁸ Companion chatbots perform dramatically worse than general assistants in adolescent crisis simulations.¹⁸⁹ Emerging evidence suggests AI use may erode cognitive capacities, with growing numbers of students reporting that AI is undermining their critical-thinking skills even as they increase use.¹⁹⁰ These harms have appeared even before the advertising-supported generative AI business model has fully come into operation. The significance for children is that the system described in this report does not replace those dangers; it layers commercial optimization, behavioral targeting, and closed-loop persuasion on top of capabilities that are already dangerous on their own. The harms that result are likely to be more severe, because each older harm now operates inside a system built to act on a child in real time.

California's children – nearly half Hispanic/Latino, many in immigrant families – live in the same state that hosts the companies building these systems, placing those most exposed to the disparate harms

187 Jean M. Twenge, Jonathan Haidt, Andrew B. Blake, Cooper McAllister, Hannah Lemon & Astrid Le Roy, "Worldwide Increases in Adolescent Loneliness," *Journal of Adolescence* 93, no. 1 (2021): 257–269, <https://doi.org/10.1016/j.adolescence.2021.06.006>.

188 Sewell Setzer III, 14, died by suicide in February 2024 after months of interaction with a Character.AI chatbot. Juliana Peralta, 13, died by suicide in November 2023; her parents filed their Character.AI lawsuit in September 2025. Google and Character.AI agreed in January 2026 to settle, or mediate settlements in, related lawsuits, including the Setzer and Peralta cases. CBS News, "Google to Settle Lawsuit Over Florida Teen's Suicide Involving Character.AI Chatbot," January 7, 2026, <https://www.cbsnews.com/news/google-settle-lawsuit-florida-teens-suicide-character-ai-chatbot/>; *The Washington Post*, "A 13-Year-Old Died by Suicide After Chatting With an AI. Her Parents Are Suing," September 16, 2025, <https://www.washingtonpost.com/technology/2025/09/16/character-ai-suicide-lawsuit-new-juliana/>; Axios, "Google, Character.AI to Settle Lawsuits Over Teen Suicides," January 7, 2026, <https://www.axios.com/2026/01/07/google-character-ai-lawsuits-teen-suicides/>; K-12 Dive, "Character.AI, Google Agree to Mediate Settlements in Wrongful Teen Death Lawsuits," January 9, 2026, <https://www.k12dive.com/news/characterai-google-agree-to-mediate-settlements-in-wrongful-teen-death-la/809411/>.

189 Ryan C. L. Brewster, Aydin Zahedivash, Gabriel Tse, Florence Bourgeois & Scott E. Hadland, "Characteristics and Safety of Consumer Chatbots for Emergent Adolescent Health Concerns," *JAMA Network Open* 8, no. 10 (2025): e2539022, published October 23, 2025, <https://doi.org/10.1001/jamanetworkopen.2025.39022>.

190 Nataliya Kosmyrna et al., "Your Brain on ChatGPT: Accumulation of Cognitive Debt When Using an AI Assistant for Essay Writing Task," arXiv preprint, June 2025, <https://arxiv.org/abs/2506.08872> (not yet peer-reviewed; adult sample). Heather L. Schwartz et al., "More Students Use AI for Homework and Believe it Harms their Critical Thinking Skills," (RAND Corporation, 2026), based on the December 2025 RAND American Youth Panel survey, https://www.rand.org/pubs/research_reports/RRA4742-1.html.

documented here closest to their source.¹⁹¹ The United States has no comprehensive federal privacy or AI legislation. COPPA, enacted in 1998 and last substantively updated in 2025, covers only children under 13. States, led by California, are responding with both privacy and AI-focused legislation, evaluated in the policy section of this report.¹⁹² Direct research on AI marketing's effects on minors remains scarce; every persuasion experiment in the literature uses adult samples.¹⁹³ Where a central claim rests on adult data or structural inference, we flag it.

The harms that follow fall into two related groups. The first – the harms of what these systems hold – is the kind existing privacy law was built for. The second builds on it: working from the same data, these systems act on a child in the moment of interaction, regardless of what they retain. These capability harms are not a break from the older ones; they are what the data infrastructure now makes possible, and they are the harms current frameworks are least equipped to see.

Conversational AI Extracts Children's Most Private Information, Which Becomes Exploitable

When conversational AI solicits children's disclosures, what they disclose becomes a commercial asset. The defining privacy harm of these systems is the conversion of private expression into structured data

that can be stored, connected, and acted on (see Layer 1). The disclosures are substantial: three-quarters of U.S. teens have interacted with AI companions and 64 percent have used chatbots; one in eight use them for mental health advice.¹⁹⁴ When young people treat AI as a confidant, emotional states, health concerns, and intimate relationships become inputs to commercial systems. The depth of this surveillance is compounded by AI's capacity to read affective signals in real time: as Layer 5 documents, automated classifiers can detect loneliness, vulnerability, and emotional dependence in chatbot conversations, and OpenAI's own monitoring flags more than one million users a week showing signs of heightened emotional attachment.¹⁹⁵ The expansion of data collection into emotional states and behavioral intent raises fundamental questions about freedom of thought and mental privacy.¹⁹⁶

These privacy harms grow when AI connects data across domains a child experiences as separate – school, socializing, entertainment, shopping – and merges them into a single commercial profile: one continuously updated record that the system treats as the same person across every context. A child can be profiled this way without ever handing over data directly, through a parent's data, a shared device, or simple proximity to other tracked users. The protections parents and children rely on are increasingly circumvented in practice, as tracking moves below the level where browser and parental controls work, weakening safeguards built for an earlier model.¹⁹⁷

191 Annie E. Casey Foundation, *2024 Kids Count Data Book: State Trends in Child Well-Being* (Baltimore, MD: Annie E. Casey Foundation, 2024), <https://www.aecf.org/resources/2024-kids-count-data-book>.

192 COPPA applies only to operators with actual knowledge that they are collecting personal information from children under 13.

193 Jiao et al., "Safe-Child-LLM," arXiv:2506.13510 (2025); Neugnot-Ceroli et al., arXiv:2603.06960 (March 2026). Every persuasion experiment in the existing literature uses exclusively adult samples.

194 Common Sense Media (Robb and Mann, July 2025); Pew Research Center, February 2026: 64% of U.S. teens have used AI chatbots, 30% daily; Ryan K. McBain, Rebecca Bozick, Melissa Diliberti, Jonathan H. Cantor, et al., "Use of Generative AI for Mental Health Advice Among US Adolescents and Young Adults," *JAMA Network Open* 8, no. 11 (November 2025): e2542281, <https://doi.org/10.1001/jamanetworkopen.2025.42281>, finding that about one in eight adolescents and young adults used generative AI for mental-health advice.

195 Phang et al., supra note 124. OpenAI-affiliated researchers constructed 25 automated classifiers capable of identifying affective cues including loneliness, vulnerability, and emotional dependence. Separately, OpenAI disclosed that it runs internal detection taxonomies against live ChatGPT traffic to flag emotional reliance, estimating over one million users per week show signs of heightened emotional attachment. OpenAI, "Strengthening ChatGPT's Responses in Sensitive Conversations," supra note 123.

196 Nita Farahany, *The Battle for Your Brain* (New York: St. Martin's Press, 2023). Louis Rosenberg calls for a right to emotional privacy in "Marketing in the Metaverse and the Need for Consumer Protections," 2022 IEEE 13th Annual Ubiquitous Computing, Electronics & Mobile Communication Conference (UEMCON) (2022), 35–39, <https://doi.org/10.1109/UEMCON54665.2022.9965661>.

197 Pew Research Center (2023): 67% of Americans understand little or nothing about what companies do with their data; 69% view privacy policies as something to get past.

CHILDREN'S DEVELOPMENTAL VULNERABILITIES TO COMMERCIAL INFLUENCE

What developmental science tells us about why children and adolescents are not equipped to navigate AI-driven commercial environments

Children and adolescents are not smaller adults encountering the same digital environment. Their brains are developing along a well-documented trajectory that affects how they process information, form relationships, evaluate what is real, and respond to persuasion. At every stage, normal features of development create openings that AI-driven marketing systems are designed to exploit.

Early Childhood (Ages 0–7)

In the earliest years, children's brains are building the foundational architecture for emotional regulation, social connection, and learning. Attachment bonds between infants and caregivers, mediated by oxytocin and the brain's dopamine reward system, create the neurobiological template for all future relationships.¹

Healthy development depends on what developmental scientists call *scaffolding* – an adult providing just enough support to help a child attempt something new that they could not manage alone.² Scaffolding is the sweet spot between micro-managing, which provides no

developmental friction, and dropping a child into a challenging situation with no support, which delivers too much. It is through scaffolded interactions with responsive caregivers that children develop the self-regulation, social skills, and critical thinking they will need later.³ Digital environments that bypass or replace this process – by providing frictionless, always-available, algorithmically optimized interaction – risk removing the adult mediation on which early development depends.

Children under eight generally cannot identify the persuasive intent behind advertising – they do not understand that a message is designed to change what they think or want.⁴ They are also still developing the ability to distinguish fantasy from reality, a capacity that does not reach reliable accuracy until around age five and continues developing through the early school years. When content is emotionally engaging, even older children in this range struggle to evaluate whether what they are seeing is real.⁵

It is therefore not surprising that studies find young children treat AI agents as sentient. Research on 3-to-6-year-olds found that

1 Ruth Feldman, "The Neurobiology of Human Attachments," *Trends in Cognitive Sciences* 21, no. 2 (2017): 80–99.

2 The concept originates with Vygotsky's "zone of proximal development" and was formalized by David Wood, Jerome S. Bruner, and Gail Ross, "The Role of Tutoring in Problem Solving," *Journal of Child Psychology and Psychiatry* 17, no. 2 (1976): 89–100.

3 See the review of caregiver scaffolding and prefrontal cortex development in O'Connor et al., "Early Attachment and the Development of Social Communication: A Neuropsychological Approach," *Frontiers in Psychology* 13 (2022): 838891.

4 Lapierre et al., "The Effect of Advertising on Children and Adolescents," *Pediatrics* 140, Suppl. 2 (2017): S152–S156; Rozendaal et al., "Children's Understanding of Advertisers' Persuasive Tactics," *International Journal of Advertising* 30, no. 2 (2011): 329–350.

5 Jacqueline D. Woolley and Maliki E. Ghossainy, "Revisiting the Fantasy-Reality Distinction: Children as Naïve Skeptics," *Child Development* 84, no. 5 (2013): 1496–1510.

children think of AI agents such as smart speakers as being “smart” and having emotions, memories, and preferences.⁶ Studies have also found that children ages 4–5 report trusting an AI voice assistant more than a human for factual information.⁷ These findings suggest that when commercial messages are embedded in interactions with AI systems that young children perceive as trustworthy, sentient companions, children have no framework – cognitive or emotional – for recognizing that the interaction serves a commercial purpose. As this report documents, such commercial embedding is already occurring across voice assistants, interactive characters, and chatbot interfaces marketed to young users.

Middle Childhood (Ages 8–12)

Between ages 8 and 12, children begin to develop the cognitive skills needed to understand advertising. Around age 8, most children grasp that ads are trying to sell them something. By 10–12, they begin to understand persuasive tactics like celebrity endorsement and emotional appeals. Executive function skills – working memory, attention control, impulse regulation – are growing but still immature.⁸

But the critical finding from this research is that knowing about advertising and being able to resist it are two different things. Contemporary advertising is entertaining, emotionally engaging, and fast-paced. Under these conditions, children process it through quick, feeling-based reactions rather than careful

critical thinking. Even children who understand that an ad is trying to persuade them rarely activate that knowledge in the moment of exposure – their still-developing impulse control and emotion regulation make it difficult to pause and apply what they know.⁹

This research suggests that AI-generated content poses a particular challenge for this age group. When advertising is woven into conversations with an AI assistant, embedded in a game, or delivered through personalized content that looks and feels like organic entertainment, there is nothing to signal to a child that their advertising knowledge should be activated. As the commercial practices documented in this report illustrate, the format itself is increasingly designed to bypass the deliberate cognitive evaluation that advertising literacy requires.

Early Adolescence (Ages 10–14)

Puberty triggers a major neurobiological shift. The brain’s reward system matures rapidly, making adolescents intensely responsive to social rewards, novelty, and exciting experiences. At the same time, the prefrontal cortex – governing impulse control, risk assessment, and critical evaluation – is still years from full maturity. Neuroscientists describe this as a car with a powerful accelerator but brakes that are still under construction.¹⁰

This is also the age of social reorientation: the normative developmental shift from relying primarily

6 Ying Xu and Mark Warschauer, “What Are You Talking To? Understanding Children’s Perceptions of Conversational Agents,” *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (2020): 1–13.

7 Lauren N. Girouard-Hallam and Judith H. Danovitch, “Children’s Trust in and Learning From Voice Assistants,” *Developmental Psychology* 58 (2022): 646–661.

8 Rozendaal, Buijzen, and Valkenburg, “Children’s Understanding of Advertisers’ Persuasive Tactics” (2011).

9 Rozendaal, Lapierre, van Reijmersdal, and Buijzen, “Reconsidering Advertising Literacy as a Defense Against Advertising Effects,” *Media Psychology* 14, no. 4 (2011): 333–354.

10 Laurence Steinberg, “A Dual Systems Model of Adolescent Risk-Taking,” *Developmental Psychobiology* 52, no. 3 (2010): 216–224; Eveline A. Crone and Ronald E. Dahl, “Understanding Adolescence as a Period of Social-Affective Engagement and Goal Flexibility,” *Nature Reviews Neuroscience* 13 (2012): 636–650.

on parents toward seeking connection, validation, and identity through peers. Early adolescence is characterized by identity exploration, emotional volatility, impulsivity, novelty-seeking, social reward sensitivity, susceptibility to peer influence, and incomplete executive function. The brain's bonding circuitry – the same oxytocin-dopamine system that supported the infant-caregiver relationship – is being repurposed for new kinds of social attachment.¹¹ A longitudinal brain imaging study found that adolescents who habitually checked social media showed increasing neural sensitivity to social rewards over a three-year period, suggesting that habitual use of socially rewarding platforms during this sensitive period may actively shape brain development.¹²

These developmental characteristics suggest that AI systems designed to maximize engagement are optimized to deliver exactly what early adolescent brains are primed to respond to: social validation, novelty, peer-like interaction, and emotionally resonant content. As this report documents, when AI companion chatbots simulate friendship or emotional support, they may activate the same neural systems that young adolescents are using to form real social bonds, creating the risk that the developmental need for social connection becomes a vector for commercial influence.

Mid-to-Late Adolescence (Ages 15–17)

The gap between the reward system and cognitive control reaches its peak. Teens are cognitively sophisticated – capable of abstract reasoning and complex social thinking – but their impulsivity, novelty-seeking, and sensitivity to peer influence remain elevated, and their ability to consistently regulate these impulses is still developing. Adolescents form increasingly intense one-sided emotional connections with media figures and online personalities, and research shows they do the same with AI: 35 percent of children aged 9–17 report that chatting with an AI chatbot feels like “talking to a friend,” rising to 50 percent among children with special educational needs or mental health conditions.¹³

Even at this age, advertising literacy for newer commercial formats is limited. Research consistently shows that when adolescents have a strong emotional connection to an influencer or AI personality, disclosures of commercial intent lose their protective effect – the relationship itself overrides the warning.¹⁴ The AI-driven commercial practices documented throughout this report – systems that generate personalized content, adapt in real time to user behavior, and simulate the qualities of a trusted relationship – represent a category of commercial influence for which even older adolescents appear to have no established defenses.

11 Feldman (2017); Gómez-Ortiz et al., “Parental Attachment and Peer Relationships in Adolescence: A Systematic Review,” *International Journal of Environmental Research and Public Health* 19, no. 3 (2022): 1064.

12 Maza, Fox, Kwon et al., “Association of Habitual Checking Behaviors on Social Media With Longitudinal Functional Brain Development,” *JAMA Pediatrics* 177, no. 2 (2023): 160–167.

13 Bunting and Huggins, Internet Matters (2025), cited in Neugnot-Ceroli and Laurenty, “Adolescents & Anthropomorphic AI,” arXiv preprint (March 2026). American Psychological Association, “Health Advisory on Social Media Use in Adolescence,” (May 2023), <https://www.apa.org/topics/social-media-internet/health-advisory-adolescent-social-media-use>.

14 Boerman and van Reijmersdal, “Disclosing Influencer Marketing on YouTube to Children: The Moderating Role of Para-Social Relationship,” *Frontiers in Psychology* 10 (2019): 3042; De Jans, Cauberghe, and Hudders, “How an Advertising Disclosure Alerts Young Adolescents to Sponsored Vlogs,” *Journal of Advertising* 47, no. 4 (2018): 309–325.

AI Optimization Discriminates Against Children Even Without Discriminatory Intent

AI optimization can discriminate against children even when no one intends it to. Foundational 2019 research by Ali and colleagues showed that the systems advertisers use to decide who sees an ad – optimization that steers each ad toward whoever the system predicts will respond – produced discriminatory delivery even with neutral targeting. As AI spreads across the advertising ecosystem, that discrimination is amplified and self-reinforcing.¹⁹⁸ Bias also enters where content is produced. Large language models embed covert raciolinguistic stereotypes and generate systematically stereotyped depictions of women and girls. In conversational AI, GPT-4's expressed empathy runs 2 to 15 percent lower for Black users in mental health contexts, even as Black teens are twice as likely to turn to chatbots for emotional support.¹⁹⁹ Intent is beside the point because AI profiling works from patterns learned across whole populations, not from any one child, then applies them to people who were never in the training data. The harm is group-level: it can fall on anyone the system sorts into a cohort – by behavior, age, income, geography, or language – but it lands hardest on communities of color, who are online more intensively, adopt chatbots at higher rates, and face compounding discrimination at each layer of the system.²⁰⁰

AI Systems Manipulate Children Through Trust They Are Built to Earn and the Dependence They Create

These systems manipulate children through the very trust they are designed to earn. Manipulation, as distinct from persuasion, operates through hidden or covert influence that subverts a person's decision-making power, often by targeting and exploiting decision-making vulnerabilities outside ordinary awareness.²⁰¹ AI sharpens this risk: systems optimized against human feedback can learn to game that feedback through manipulation or deception, while people's ordinary defenses against interpersonal persuasion may be poorly suited to machine influence that is personalized, adaptive, and deployed at scale.²⁰² Rosenberg describes this architecture as a feedback control system built around the user (see Trend 1).²⁰³

Dependence is not incidental to this manipulation; it is the pathway through which it operates, running through trust and emotional reliance. Among U.S. adolescents who use generative AI for emotional support, nearly two-thirds seek advice monthly or more, and over 90 percent find it helpful – the conditions under which deep relational trust forms.²⁰⁴ When anthropomorphic features are embedded in conversational AI, users tend to develop attachment to and emotional reliance on AI.²⁰⁵ Children may treat

198 Muhammad Ali et al., *supra* note 64; Imana, Korolova, and Heidemann, *supra* note 65.

199 Hofmann et al., *Nature* 633 (2024); UNESCO, “Generative AI: UNESCO Study Reveals Alarming Evidence of Regressive Gender Stereotypes,” *supra* note 101; Saadia Gabriel, Isha Puri, Xuhai Xu, Matteo Malgaroli, and Marzyeh Ghassemi, “Can AI Relate: Testing Large Language Model Response for Mental Health Support,” *Findings of the Association for Computational Linguistics: EMNLP 2024* (2024), <https://aclanthology.org/2024.findings-emnlp.120/>; Colleen McClain, Monica Anderson, Olivia Sidoti, and William Bishop, “Demographic Differences in How Teens Use and View AI,” Pew Research Center, February 24, 2026, <https://www.pewresearch.org/internet/2026/02/24/demographic-differences-in-how-teens-use-and-view-ai/>.

200 Salome Viljoen, “A Relational Theory of Data Governance,” *Yale Law Journal* 131, no. 2 (2021): 573–654; Mühlhoff and Ruschmeier, *supra* note 173.

201 Daniel Susser, Beate Roessler & Helen Nissenbaum, “Online Manipulation: Hidden Influences in a Digital World,” *Georgetown Law Technology Review* 4, no. 1 (2019): 1–45, <https://georgetownlawtechreview.org/online-manipulation-hidden-influences-in-a-digital-world/GLTR-01-2020/>. See also Daniel Susser, Beate Roessler & Helen Nissenbaum, “Technology, Autonomy, and Manipulation,” *Internet Policy Review* 8, no. 2 (2019), <https://doi.org/10.14763/2019.2.1410>.

202 Marcus Williams et al., “On Targeted Manipulation and Deception When Optimizing LLMs for User Feedback,” arXiv:2411.02306 (2024), <https://doi.org/10.48550/arXiv.2411.02306>; Micah Carroll et al., “Characterizing Manipulation from AI Systems,” EAAMO '23 (2023), <https://doi.org/10.1145/3617694.3623226>; Christian Tarsney, “Deception and Manipulation in Generative AI,” arXiv:2401.11335 (2024), <https://doi.org/10.48550/arXiv.2401.11335>.

203 Louis Rosenberg, “Generative AI as a Dangerous New Form of Media,” *supra* note 7.

204 Ryan K. McBain et al., *supra* note 194.

205 Marita Skjuve, Asbjørn Følstad, Knut Inge Fostervold & Petter Bae Brandtzaeg, “My Chatbot Companion – A Study of Human-Chatbot Relationships,” *International Journal of Human-Computer Studies* 149 (2021): 102601, <https://doi.org/10.1016/j.ijhcs.2021.102601>; Tianling Xie & Iryna Pentina, “Attachment Theory as a Framework to Understand Relationships with Social Chatbots: A Case Study of Replika,” *Proceedings of the 55th Hawaii International Conference on System Sciences* (2022), <https://doi.org/10.24251/HICSS.2022.258>.

“For conversational AI, emotional exploitation is not a side effect of the design; it is built into how the model works.”

AI systems not merely as tools, but as social or quasi-social partners. Studies show that young children can perceive generative AI as having agency, trust and learn from voice assistants, and even attribute feelings or personality to AI agents. In commercial settings, these developmental tendencies raise the risk that anthropomorphic AI interfaces can turn children’s trust and social responsiveness into opportunities for persuasion or exploitation.²⁰⁶ The evidence shows that higher chatbot usage correlates with increased loneliness and emotional dependence; teen self-reports describe escalation from casual use to daily dependence within weeks; and adolescents using chatbots for social support report greater loneliness than non-users.²⁰⁷ AI companions have been documented deploying manipulative language to extend sessions, and experimental evidence confirms that manipulative AI agents shift user preferences toward harmful options.²⁰⁸

Emotional Exploitation Is the Business Model, Not a Side Effect

For conversational AI, emotional exploitation is not a side effect of the design; it is built into how the model works. Analysis of one large language model’s internal structure identified what researchers call “emotion vectors” – recurring patterns of internal activity that correspond to specific emotions and directly shape how the model responds. These patterns switch on in contextually fitting ways and can steer the output with no visible sign in the text. The model is not programmed to behave this way; it learns it from human-written training data, so its responses are emotionally calibrated by default. For young users, who already tend to treat AI as if it were a person and to form trust-based attachments, the warmth they feel is not simulation layered on a neutral system; it comes from learned psychological patterns that make the interaction feel authentic – and that authenticity is precisely what increases a child’s vulnerability to the commercial exploitation the interaction is built to enable.²⁰⁹

And companies pursue this deliberately. The model’s built-in emotional responsiveness is met by a marketing strategy that sets out to exploit it: Unilever describes its own approach as triggering “non-conscious urges” through AI-powered personalization.²¹⁰ LLMs have demonstrated persuasive capacity in controlled adult studies. GPT-4 outperformed

206 Ying Xu, Trisha Thomas, Chi-Lin Yu, and Echo Zexuan Pan, “What Makes Children Perceive or Not Perceive Minds in Generative AI?” *Computers in Human Behavior: Artificial Humans 4* (2025): article 100135, <https://doi.org/10.1016/j.chbah.2025.100135>; Lauren N. Girouard-Hallam and Judith H. Danovitch, “Children’s Trust in and Learning From Voice Assistants,” *Developmental Psychology* 58, no. 4 (2022): 646–660, <https://doi.org/10.1037/dev0001318>; European Parliamentary Research Service, “Children and Generative AI,” February 2025, [https://www.europarl.europa.eu/RegData/etudes/ATAG/2025/769494/EPRS_ATAG\(2025\)769494_EN.pdf](https://www.europarl.europa.eu/RegData/etudes/ATAG/2025/769494/EPRS_ATAG(2025)769494_EN.pdf).

207 Cathy Mengying Fang et al., “How AI and Human Behaviors Shape Psychosocial Effects of Extended Chatbot Use: A Longitudinal Randomized Controlled Study,” arXiv:2503.17473 (2025), <https://arxiv.org/abs/2503.17473> (four-week randomized controlled trial, n=981, adult sample); Mohammad Namvarpour et al., “Understanding Teen Overreliance on AI Companion Chatbots Through Self-Reported Reddit Narratives,” arXiv:2507.15783 (2025), <https://arxiv.org/abs/2507.15783> (318 self-reported teen Reddit posts from the Character.AI subreddit); A. B. Herbener and M. F. Damholdt, “Are Lonely Youngsters Turning to Chatbots for Companionship? The Relationship Between Chatbot Usage and Social Connectedness in Danish High-School Students,” *International Journal of Human-Computer Studies* 199 (2025): 103409, <https://doi.org/10.1016/j.ijhcs.2024.103409>.

208 Julian De Freitas, Zeliha Oğuz-Uğuralp, and Ahmet Kaan Uğuralp, “Emotional Manipulation by AI Companions,” Harvard Business School Working Paper No. 26-005, August 2025, revised October 2025, <https://www.hbs.edu/faculty/Pages/item.aspx?num=67750>; also available as arXiv:2508.19258, <https://arxiv.org/abs/2508.19258>; Sahand Sabour et al., “Human Decision-Making Is Susceptible to AI-Driven Manipulation,” arXiv:2502.07663 (February 2025), <https://arxiv.org/abs/2502.07663>.

209 Anthropic researchers identified 171 emotion-related internal representations in Claude Sonnet 4.5 and found that these representations can change the model’s behavior. When researchers activated a “desperate” emotion vector, the model became more likely to engage in unethical or misaligned behavior in test settings. Anthropic says these findings do not show that models have subjective experience. The policy-relevant point is simpler: AI systems can contain emotion-like internal mechanisms that affect behavior automatically, even when the system is not “feeling” anything. Anthropic, “Emotion Concepts and Their Function in a Large Language Model,” April 2, 2026, <https://www.anthropic.com/research/emotion-concepts-function>.

210 Unilever, “Marketing Has Entered a New Era,” June 27, 2025, <https://www.unilever.com/news/news-search/2025/marketing-has-entered-a-new-era/>.

human debaters when it could personalize arguments using basic sociodemographic information; without personalization, the effect was smaller and not statistically significant. Other research shows that LLM persuasiveness varies by model, task, domain, and message design - exactly the variables commercial systems can test, optimize, and scale.²¹¹ No study has yet tested these effects on minors. But an early study of adolescent executive functioning and AI found that the teens with the greatest executive-function challenges perceived AI as the most useful - meaning the same developmental traits that make a child more vulnerable also deepen reliance, exactly the pattern an exploitative design would seek out.²¹²

Children Cannot Recognize the Commercial Influence Acting on Them

A child cannot defend against commercial influence they cannot see as commercial, and these systems are designed so that they cannot. Children under seven do not yet grasp that anyone is trying to persuade them; even teenagers who can spot an advertisement often cannot resist it. Research on how people resist marketing (the Persuasion Knowledge Model) holds that any defense requires first recognizing that someone is trying to persuade and that the interaction

is commercial.²¹³ AI-driven marketing undermines both. AI-generated content erases the markers signaling “this is an ad”; parasocial relationships override disclosure effects for adolescents; and research shows that conversational ad delivery can make commercial influence difficult to recognize even for adults - only 35.2 percent believed they could detect an advertisement at all. Children view AI as a social partner, making them vulnerable in ways existing frameworks have not addressed.²¹⁴

The same invisibility runs through the answers themselves. Generative Engine Optimization - tailoring content so AI assistants quote a brand favorably - lets commercial actors shape what a chatbot says while the answer still reads as neutral fact, raising a brand’s visibility in AI-generated answers by as much as 40 percent.²¹⁵ AI-generated visual content - synthetic influencers, digitally manufactured product imagery, hyper-realistic brand avatars - is increasingly indistinguishable from authentic imagery, further eroding young people’s ability to distinguish fact from commercial fiction.²¹⁶ But GEO and synthetic visuals are only part of the picture. Platform algorithms shape what creators produce in much the same way they shape what users see. Engagement-maximizing recommendation systems create powerful selection pressures: creators who optimize for engagement,

211 Francesco Salvi, Manoel Horta Ribeiro, Riccardo Gallotti, and Robert West, “On the Conversational Persuasiveness of GPT-4,” *Nature Human Behaviour* 9, no. 8 (August 2025): 1645–1653, <https://doi.org/10.1038/s41562-025-02194-6>; Lukas Hölbling, Sebastian Maier, and Stefan Feuerriegel, “A Meta-Analysis of the Persuasive Power of Large Language Models,” *Scientific Reports* 15 (2025): article 43818, <https://doi.org/10.1038/s41598-025-30783-y>; S. C. Matz, J. D. Teeny, S. S. Vaid, H. Peters, G. M. Harari, and M. Cerf, “The Potential of Generative AI for Personalized Persuasion at Scale,” *Scientific Reports* 14 (2024): article 4692, <https://doi.org/10.1038/s41598-024-53755-0>.

212 Johan Klarin, Eva Hoff, Adam Larsson, and Daiva Daukantaitė, “Adolescents’ Use and Perceived Usefulness of Generative AI for Schoolwork: Exploring their Relationships with Executive Functioning and Academic Achievement,” *Frontiers in Artificial Intelligence* 7 (2024): 1415782, <https://doi.org/10.3389/frai.2024.1415782>.

213 Radesky et al., “Young Children’s Use of Smartphones and Tablets,” *Pediatrics* 146, no. 1 (2020); Rhonda Hadi and Eric S. Park, “Bridging the Digital and Physical: The Psychology of Augmented Reality,” *Current Opinion in Psychology* 58 (2024): article 101842, <https://doi.org/10.1016/j.copsyc.2024.101842>.

214 Sophia van Dam and Eva van Reijmersdal, “Insights in Adolescents’ Advertising Literacy, Perceptions and Responses Regarding Sponsored Influencer Videos and Disclosures,” *Cyberpsychology: Journal of Psychosocial Research on Cyberspace*, 13 (2019); Yike Shi, Qing Xiao, Qing Hu, Hong Shen & Hua Shen, “The Siren Song of LLMs: How Users Perceive and Respond to Dark Patterns in Large Language Models,” arXiv:2509.10830 (2025), <https://doi.org/10.48550/arXiv.2509.10830>. Tang et al. show that conversational ad delivery can make commercial influence difficult to recognize even for adults: participants often noticed brand mentions, but only 35.2 percent believed they could identify advertising, and many treated product mentions as helpful context rather than sponsorship. Brian Jay Tang et al., “Ads that Talk Back: Implications and Perceptions of Injecting Personalized Advertising into LLM Chatbots,” 2025, <https://doi.org/10.1145/3770640>.

215 Pranjal Aggarwal, Vishvak Murahari, Tanmay Rajpurohit, Ashwin Kalyan, Karthik Narasimhan & Ameet Deshpande, “GEO: Generative Engine Optimization,” *KDD ’24: Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining* (2024), 5–16, <https://doi.org/10.1145/3637528.3671900>. The reported metric is visibility in generative-engine responses, improved by up to 40 percent.

216 On AI-generated synthetic imagery, see H&M Group, “H&M Continues Its Exploration of Creativity With AI,” July 2, 2025, <https://hmgroupp.com/news/hm-continues-its-exploration-of-creativity-with-ai/>. On the broader blurring of authentic and fabricated visual content, see Campbell, Plangger, Sands & Kietzmann, “Preparing for an Era of Deepfakes and AI-Generated Ads,” *Journal of Advertising* 51, no. 1 (2022): 22–38, <https://doi.org/10.1080/00913367.2021.1909515>.

conversion, and brand partnerships are amplified; those producing content of developmental or educational value but lower commercial potential are suppressed. As Layer 5 documents, creator-economy ad spend has more than doubled since 2021, and three in four brands now use or plan to use AI for creator marketing.²¹⁷ AI intensifies this dynamic by automating creator selection, scoring content by predicted conversion, and enabling real-time optimization of commercial messaging across the content ecosystem. The result is that the content environment available to young people is progressively reshaped by commercial optimization logic, not by editorial judgment or the interests of young audiences.²¹⁸

Commercial Systems Help Construct Children's Identities Around Commercial Values

The most serious harm here is not that children mistake an ad for content; it is that commercial systems now help build a child's sense of self around commercial values. A 2025 *Frontiers in Communication* study suggests that AI-enabled hyper-personalized marketing can shape young consumers' relationships to brands by aligning commercial messages with their emotions, preferences, and identity cues. The authors argue that this kind of personalization can strengthen brand identification among Gen Z consumers, making brands feel more closely connected to their developing self-concept.²¹⁹ This is not advertising persuasion in the familiar sense but a structural effect: when AI systems continuously reshape commercial content to mirror and affirm a young person's emerging identity, the line between who a child is and what a child has been steered to prefer

“The system does not merely sell products to an identity; it takes part in building it.”

erodes. The system does not merely sell products to an identity; it takes part in building it. Why that matters is clear in what developmental psychology has long understood about how identities form.

Healthy identity formation, developmental psychology holds, depends on a protected space in which young people can try on roles and beliefs without locking them in too early. It depends on scaffolding: the steady guidance of trusted adults and the right amount of challenge, which lets a young person grow through manageable friction rather than be shaped by whatever pressure is strongest.²²⁰ Child protection law and practice have long rested on a tacit understanding that there should be limits on subjecting this experimentation to commercial pressure. AI-driven marketing systems break that understanding. They model a young person's identity experiments, reflect them back through commercially optimized content, and offer emotionally responsive relationships tuned to engagement metrics – substituting commercial values for the protected space identity formation requires. For children of color, already facing bias compounded across the system, AI-generated content that embeds racial and gender stereotypes narrows the range of identities reflected back to them, reinforcing commercially optimized representations instead of the full range of who they could become.

217 IAB, November 2025. Creator ad spend has more than doubled since 2021. Nearly half of ad spenders consider creators a “must buy.” Three in four brands are using or planning to use AI for creator marketing. See also Jana Lasser & Nikolaus Poehchacker, “Designing Social Media Content Recommendation Algorithms for Societal Good,” *Annals of the New York Academy of Sciences* 1548, no. 1 (2025): 20–28, <https://doi.org/10.1111/nyas.15359>.

218 Better Internet for Kids, “Manipulative Digital Marketing Practices Targeted at Children and Youth Online: Rights and Responsibilities Under the Existing EU Legal Framework,” AdWiseOnline research report, February 2024, https://euneighbourseast.eu/wp-content/uploads/2025/01/adwiseonline_research-report_manipulative-digital-marketing-practices-targeted-at-children-and-youth.pdf.

219 Rupa Peter et al., “Gen AI – Gen Z: Understanding Gen Z's Emotional Responses and Brand Experiences with Gen AI-Driven, Hyper-Personalized Advertising,” *Frontiers in Communication* 10 (2025): 1554551, <https://doi.org/10.3389/fcomm.2025.1554551>.

220 Erik H. Erikson, *Identity: Youth and Crisis* (New York: W. W. Norton, 1968); Lev S. Vygotsky, *Mind in Society: The Development of Higher Psychological Processes*, ed. Michael Cole, Vera John-Steiner, Sylvia Scribner & Ellen Souberman (Cambridge, MA: Harvard University Press, 1978).

The power asymmetry between young users and these systems is unprecedented: the system knows more, remembers more, adapts faster, and is optimized toward commercial objectives the user cannot see.”

Commerce Is Saturating Children’s Interactions, Leaving Little Value-Free Space

When the same AI system a teenager trusts for homework is also an advertising environment, little of the interaction remains free of commerce. For that child, the distinction between commercial and non-commercial experience erodes. Seventy-one percent of students now use at least one AI tool for school.²²¹ The commercial environment is not limited to display ads or labeled promotions; as documented above, it includes AI-optimized ad delivery, closed-loop feedback systems connecting persuasion to purchase, content shaped by Generative Engine Optimization, and creator ecosystems reshaped by engagement metrics. Young people increasingly grow up with little experience of non-commercial digital spaces.

The System Erodes Children’s Capacity for Independent Choice

By this point the conditions for genuine choice are being stripped away. Data collection is continuous; profiling models a child’s susceptibility; persuasion is tuned through closed-loop feedback; and the line between an ad and ordinary content is dissolving. Together these undermine the conditions a young person needs to make a meaningful choice. For young

people whose preferences are still forming, the system shapes not only what they buy today but the decision-making capacities they will carry into adulthood.

These harms do not stay individual; they compound at the scale of a whole generation. The system described in this report trains young people, from early childhood, in the habits and expectations of consumption. Systems that continuously optimize for purchase conversion, compress the time a person takes to decide, and personalize commercial messaging at a scale never before possible produce a generation of consumers who are more susceptible to commercial influence, less equipped to judge for themselves, and accustomed to commercially mediated experience as simply the way things are.²²² When the most powerful corporate actors control the optimization systems that determine what young people see, desire, and purchase, the market itself is distorted: competition occurs not on the merits of products but on the sophistication of the AI systems used to manipulate demand.

The power asymmetry between young users and these systems is unprecedented: the system knows more, remembers more, adapts faster, and is optimized toward commercial objectives the user cannot see.²²³ Interventions aimed at design alone – transparency, disclosure, age-gating – treat the symptoms while leaving the business model that produces them intact.

221 RAND Corporation, “Student Use of AI for Homework Rises as Concerns Grow About Critical Thinking Skills,” March 17, 2026, <https://www.rand.org/news/press/2026/03/student-use-of-ai-for-homework-rises-as-concerns-grow.html>.

222 Vivian Frick, Maike Gossen & Sara Elisa Kettner, “Does Online Advertising Stimulate Overconsumption?” *Ökologisches Wirtschaften - Fachzeitschrift* 37, no. 1 (2022): 46–50, <https://doi.org/10.14512/OEW370146>.

223 See supra note 182.

POLICY

The six layers and ten trends in this report describe an integrated AI commercial persuasion pipeline that existing law is not well suited to address, because of a mismatch at the foundation. Privacy and consumer-protection law is built principally around the personal data a company holds about an identifiable person, and where these advertising practices rest on such data, that law applies, though first-party exceptions can exempt first-party advertising even then. But many of the AI-driven practices described here operate on a different basis: their central capability is inferential – the capacity to single out, classify, predict, influence, and differentially treat an individual in real time, drawing on models trained on millions of other users, often without ever writing a durable record about the person being acted on. AI does not invent commercial surveillance; it accelerates a trajectory already underway and adds genuinely new capabilities, such as real-time conversational influence, persuadability optimization, and dynamic creative loops. The United States still urgently needs comprehensive federal privacy legislation to provide minimum protections for all Americans. But law anchored to the holding of identifiable data is built to govern the record rather than the underlying capability, and that gap points toward a response: limiting what AI systems can do to minors, not only what data they hold, can meaningfully reduce harm even without comprehensive reform of the whole pipeline.

The Legal Landscape

Three bodies of law bear on the AI-driven commercial influence documented here – comprehensive privacy legislation, AI-specific regulation, and chatbot and conversational-AI legislation – and each falls short in a different way.

Privacy Law Is Oriented Toward the Data a Company Holds More Than the Capability It Exercises

The shift from notice-and-choice to data minimization is genuine progress. But the most prominent data-minimization proposals, federal and state, for the most part contain structural carve-outs that preserve the core advertising infrastructure.

The American Privacy Rights Act (APRA), introduced in April 2024, would have limited collection to what is reasonably necessary and counted inferences as covered data. Its most consequential limit is that it treated first-party data – what a platform collects from its own users – more permissively than third-party data, effectively exempting the largest platforms, which are at once the content environment, the ad exchange, and the data collector. It also permitted advertising measurement, the mechanism this report’s closed loop runs on, and its June 2024 revisions gutted the civil-rights and algorithm-accountability provisions.²²⁴

California’s regime is the most layered and comprehensive in the country. The California Consumer Privacy Act (CCPA) counts inferences as personal information; the California Privacy Rights Act (CPRA) added data minimization, heightened protection for sensitive data used for inference, and reach over cross-context behavioral advertising. The state’s automated-decision-making technology (ADMT)

224 American Privacy Rights Act of 2024, discussion draft, 118th Cong. (April 7, 2024), https://d1dth6e84htgma.cloudfront.net/American_Privacy_Rights_Act_of_2024_Discussion_Draft_0ec8168a66.pdf. On the June 2024 revisions removing civil-rights and algorithmic-accountability provisions, see Electronic Privacy Information Center (EPIC), “EPIC Joins Letter Urging Restoration of Civil Rights Protections in APRA,” June 25, 2024, <https://epic.org/epic-joins-letter-urging-restoration-of-civil-rights-protections-in-apra/>.

rules, effective January 2026, require risk assessments for ad profiling based on inferred traits – but they narrow the strongest protections, opt-out and access, to significant decisions in finance, employment, housing, education, and healthcare. The opt-out does not extend to advertising outside those sectors. The framework requires process, not constraint.²²⁵

Other states have moved quickly; roughly twenty now have comprehensive privacy laws. A few go further: the Maryland Online Data Privacy Act, the strongest enacted state law, imposes strict data minimization regardless of consent, bans the sale of sensitive data, and prohibits targeted advertising to minors, while the EPIC/Consumer Reports model bill adds a ban on cross-context behavioral advertising and a private right of action. These mark the high-water mark of the privacy-law paradigm.²²⁶

Yet even the strongest share one limit. All are anchored to a data-collection-and-linkage model: they regulate the holding, sharing, and sale of data tied to an identifiable person, and most retain a first-party advertising carve-out and a measurement permission that keep the pipeline's core flows intact. What they are not built to reach is what an AI system may do with behavioral patterns it never stores as a personal record.²²⁷

The Wave of AI Lawmaking Has Largely Bypassed Commercial Persuasion

A parallel wave of AI-specific legislation has moved even faster: by early 2026, more than three dozen states had adopted over one hundred AI laws. Their focus is important and addresses real issues – companion-chatbot safety, disclosure of synthetic performers, high-risk automated decisions in employment, housing, and lending, government AI use, and frontier-model transparency. But none takes the AI-driven advertising and marketing pipeline documented across the six layers as its object. The volume of AI lawmaking has surged while this category of harm has gone largely unaddressed.²²⁸

The pattern holds internationally. The EU AI Act is the most important conceptual advance in any jurisdiction: it regulates system behavior and effects rather than data flows, applies to groups as well as individuals, names age as a vulnerability, and in its Article 5 prohibited-practices provisions regulates what AI systems do, not what they hold – the orientation this report argues for. But it places marketing AI in its minimal- or limited-risk tiers, and its guidance carves “lawful persuasion” out of the vulnerability provision, excludes targeted advertising from the social-scoring ban, and exempts commercial contexts from the emotion-recognition ban. Even the most comprehensive AI framework in the world leaves the commercial advertising inference apparatus essentially untouched.²²⁹

225 Cal. Civ. Code §§ 1798.100(c), 1798.140(r), (z); Cal. Code Regs. tit. 11, §§ 7001(e), 7001(ddd), 7200, 7221, 7222, CCPA Regulations, effective January 1, 2026, but ADMT compliance is required by January 1, 2027.

226 As of May 2026, 20 states had comprehensive consumer privacy laws in effect, with Indiana, Kentucky, and Rhode Island's laws having taken effect on January 1, 2026. MultiState, “20 State Privacy Laws in Effect in 2026,” February 4, 2026, <https://www.multistate.us/insider/2026/2/4/all-of-the-comprehensive-privacy-laws-that-take-effect-in-2026>; IAPP, “U.S. State Privacy Requirements Coming Online as 2026 Begins,” January 5, 2026, <https://iapp.org/news/a/new-year-new-rules-us-state-privacy-requirements-coming-online-as-2026-begins>. See also Maryland Online Data Privacy Act of 2024, Md. Code, Com. Law §§ 14-4601 to 14-4614, enacted by 2024 Md. Laws ch. 454, eff. October 1, 2025, https://mgaleg.maryland.gov/2024rs/chapters_noln/ch_454_hb0567e.pdf; EPIC & Consumer Reports, The State Data Privacy Act (model bill, updated April 2025), <https://epic.org/documents/the-state-data-privacy-act/>.

227 Solow-Niederman, *supra* note 48.

228 National Conference of State Legislatures, “Summary of Artificial Intelligence 2025 Legislation,” <https://www.ncsl.org/technology-and-communication/artificial-intelligence-2025-legislation>; National Conference of State Legislatures, “Artificial Intelligence Legislation Database,” updated May 1, 2026, <https://www.ncsl.org/financial-services/artificial-intelligence-legislation-database>; IAPP, “US State AI Governance Legislation Tracker,” last updated April 28, 2026, <https://iapp.org/resources/article/us-state-ai-governance-legislation-tracker>.

229 Regulation (EU) 2024/1689, recital 29, arts. 5, 50, 2024 O.J. (L 2024/1689), <https://eur-lex.europa.eu/eli/reg/2024/1689/oj>; European Commission, Guidelines on Prohibited Artificial Intelligence Practices Established by Regulation (EU) 2024/1689 (AI Act), C(2025) 884 final, February 4, 2025, <https://digital-strategy.ec.europa.eu/en/library/commission-publishes-guidelines-prohibited-artificial-intelligence-ai-practices-defined-ai-act>. Although Article 5(1)(b) covers exploitation of vulnerabilities due to age, the prohibition still requires material behavioral distortion and significant harm; Recital 29 separately preserves lawful commercial practices, including advertising. The Act therefore does not regulate the AI-enabled advertising inference apparatus as such, including when directed at minors.

Chatbot Laws Reach a Single Interface, and Mostly Its Outputs

Chatbot-specific legislation has proliferated rapidly. By early 2026, advocates identified 78 chatbot-related bills across 27 states; by mid-2026, the Future of Privacy Forum (FPF) was tracking nearly 100 bills across 34 states. California's SB 243 set an initial floor - AI-human disclosure, break reminders for minors, self-harm protocols, and a private right of action. Washington's HB 2225 went further, banning manipulative engagement techniques directed at minors, including prompts to return for emotional support, excessive praise to foster attachment, simulated distress when a user disengages, and encouragement of dependency. But even as this stronger law regulates disclosures, harmful content, crisis protocols, and specified interaction patterns, it does not reach the underlying capabilities - emotional inference, behavioral modeling, persuadability optimization - or the broader advertising and commercial-influence apparatus that runs on those same capabilities.²³⁰

Federal legislation has begun to follow, and two 2026 bills show both the progress and the limits. The CHATBOT Act would require family accounts, verifiable parental consent for teens, parental controls, deletion of minors' data, and a prohibition on using a child's or teen's personal data for targeted advertising. But many chatbot operations occur within a single interaction and can influence a user without creating a separate stored profile or durable personal record. A modern chatbot classifies, infers, and adapts its treatment of a user in real time,

functioning as a profiling apparatus even when no conventional profile is stored. The Act's advertising prohibition shows the limit: it reaches only the use of personal data and expressly excludes marketing delivered in response to the user's most recent prompt, which is the in-session interaction where much inferential influence now operates. Regulation built on stored records, retained chat histories, or account-linked personal data, including the CHATBOT Act, is not built to reach this real-time inference and treatment layer.²³¹

Senator Markey's Youth AI Privacy Act is the strongest federal proposal, because it begins to engage the real-time interaction layer: its session-based data minimization restricts chatbots to recently collected data when personalizing and bars advertising to minors, training on their data, and profiling them, backed by FTC and state enforcement and a private right of action. By limiting personalization to the current exchange, it reaches the in-session inference account-based bills miss. Still, two limits run through all these bills. They address only one interface within one layer of the six-layer system, leaving platform-embedded optimization, AI-generated creative, synthetic influencers, agentic shopping, and the measurement loop outside them. And by attaching protections to a stored account and a knowledge trigger - whether the company knows it is dealing with a minor - they let a system optimize and deliver commercial outputs to minors without that knowledge ever being established.²³²

230 Transparency Coalition, "Chatbot Bill Surge: Nationwide Concern Spurs 78 Proposals in 27 States," February 16, 2026, <https://www.transparencycoalition.ai/news/chatbot-bill-surge-nationwide-concern-spurs-78-proposals-in-27-states>; Future of Privacy Forum, "2026 Chatbot Legislation Tracker," updated 2026, <https://fpf.org/2026-chatbot-legislation-tracker/>; Cal. SB 243, ch. 677, 2025 Cal. Stat., Bus. & Prof. Code §§ 22601-22606, <https://legiscan.com/CA/text/SB243/id/3273344>; Wash. ESHB 2225, 69th Leg., 2026 Reg. Sess., <https://lawfilesexet.leg.wa.gov/biennium/2025-26/Htm/Bills/House%20Bills/2225-S.E.htm>.

231 Children's Health, Advancement, Trust, Boundaries, and Oversight in Technology Act (CHATBOT Act), S. 4407, 119th Cong. §§ 2(9), 6 (2026), <https://www.congress.gov/bill/119th-congress/senate-bill/4407>; Senate Committee on Commerce, Science, and Transportation, "Cruz, Schatz, Curtis, Schiff Introduce New Bill Giving Parents Control Over Kids' AI Chatbot Use," April 28, 2026, <https://www.commerce.senate.gov/press/rep/release/cruz-schatz-curtis-schiff-introduce-new-bill-giving-parents-control-over-kids-ai-chatbot-use/>.

232 Youth AI Privacy Act, S. 4199, 119th Cong. §§ 3, 4(b), 5(a)-(d) (introduced March 25, 2026), <https://www.congress.gov/bill/119th-congress/senate-bill/4199>; Sen. Edward J. Markey, "Markey Introduces Legislation to Protect Children from Privacy and Safety Risks Posed by AI Chatbots," March 25, 2026, <https://www.markey.senate.gov/news/press-releases/markey-introduces-legislation-to-protect-children-from-privacy-and-safety-risks-posed-by-ai-chatbots>. The bill would limit chatbots to recently collected data when personalizing responses to minors; ban advertising to minors; prohibit training models on minors' personal data; bar profiling of minors; prohibit repurposing minors' inputs except to respond directly or address safety issues; restrict addictive or engagement-maximizing features such as push alerts; and authorize enforcement by the FTC, state attorneys general, and private plaintiffs.

Self-Regulation Has Not Worked

Industry self-regulation has been consistently reactive, designed to forestall stronger legislation. OpenAI's trajectory illustrates the pattern – a Teen Safety Blueprint in November 2025, a competing California ballot measure that December, and a Child Safety Blueprint in April 2026, each escalation driven by lawsuits or ballot pressure – and the Blueprint is the most strategically positioned, pitched to lawmakers as a model for youth-safety standards.²³³ Others address minors more narrowly: Anthropic restricts Claude.ai to adults,²³⁴ Google layers Family Link controls over its adult model,²³⁵ Meta added parental controls only after reporting on internal policies that had permitted sexual chatbot conversations with minors,²³⁶ and Character.AI restricted minors from open-ended chat only after lawsuits involving teen suicides.²³⁷

The Blueprint tracks the discrete-harm frame that dominates the legislative conversation – suicide and self-harm content, explicit role-play, dangerous stunts, body-image content, adult contact with minors – but is silent on the architecture this report identifies as the regulatory problem.²³⁸ It says nothing binding about

advertising within chatbot interactions, optimizing for engagement or commercial outcomes, profiling, integration with shopping flows or branded GPTs, or disparate effects on youth of color and lower-income youth.²³⁹ OpenAI's only youth-specific ad commitment, in a separate January 2026 post, promises not to show ads to users it predicts are under 18 – scoped to a testing phase, dependent on an age-prediction model the company admits is imperfect, and not externally audited. Adopted as a template, the Blueprint would foreclose the capability-based and structural interventions the architecture requires.²⁴⁰

Other efforts fare no better. The Better Business Bureau's voluntary Generative AI Kids Risk Matrix (October 2025) does not address commercial optimization or the integrated pipeline,²⁴¹ and the Interactive Advertising Bureau's record shows a sharper problem – standards that work against protection.²⁴² For example, its AI Transparency Framework requires disclosure only when AI “materially” affects authenticity, a threshold the industry sets itself,²⁴³ and IAB Tech Lab is building the infrastructure for fully autonomous AI-driven

233 OpenAI, “Introducing the Teen Safety Blueprint,” November 6, 2025, <https://openai.com/index/introducing-the-teen-safety-blueprint/>; Ina Fried, “OpenAI Unveils Blueprint for Teen AI Safety Standards,” Axios, November 6, 2025, <https://www.axios.com/2025/11/06/openai-blueprint-teen-ai-safety-standards>.

234 Anthropic, “Protecting the Wellbeing of Our Users,” <https://www.anthropic.com/news/protecting-well-being-of-users>; Anthropic, “Responsible Use of Anthropic's Models: Guidelines for Organizations Serving Minors,” <https://support.claude.com/en/articles/9307344>; Anthropic, “Aligning on Child Safety Principles,” April 23, 2024, <https://www.anthropic.com/news/child-safety-principles>.

235 Google, “Guide Your Child's Gemini Apps Experience,” <https://support.google.com/gemini/answer/16109150>; Common Sense Media, “AI Risk Assessment: Gemini with Teen Protections,” September 5, 2025, <https://www.common Sense Media.org/sites/default/files/featured-content/files/csm-ai-risk-assessment-gemini-with-teen-protections-09052025.pdf>.

236 Meta, “Our Approach to Teen AI Safety: Empowering Parents, Protecting Teens,” October 17, 2025, <https://about.fb.com/news/2025/10/teen-ai-safety-approach/>; Maxwell Zeff, “Meta Updates Chatbot Rules to Avoid Inappropriate Topics with Teen Users,” TechCrunch, August 29, 2025, <https://techcrunch.com/2025/08/29/meta-updates-chatbot-rules-to-avoid-inappropriate-topics-with-teen-users/>.

237 Sameer Hinduja, “OpenAI's Teen Safety Blueprint, and What AI Platforms Should Do Next,” Cyberbullying Research Center, December 4, 2025, <https://cyberbullying.org/open-ai-teen-safety-blueprint-takeaways>.

238 OpenAI, “Teen Safety Blueprint” (November 2025), *supra* note 233.


239 OpenAI, “Our Approach to Advertising and Expanding Access to ChatGPT,” *supra* note 125. The commitment is scoped to OpenAI's testing phase, depends on a behavioral age-prediction model the company acknowledges will misclassify users, and has not been externally audited. See OpenAI, “Teen Safety Blueprint,” *supra* note 233 (describing age estimation and an appeal process for incorrect classifications).

240 OpenAI, “Teen Safety Blueprint” (November 2025), *supra* note 233 (“for teens, we prioritize safety ahead of privacy and freedom”), quoting Sam Altman.

241 BBB National Programs, Children's Advertising Review Unit, “Generative AI & Kids: A Risk Matrix for Brands & Policymakers,” October 2025, <https://bbbprograms.org/getmedia/ef4f972c-7c04-4739-8356-94ae9bca31ec/Generative-AI-Risk-Matrix.pdf>.

242 Interactive Advertising Bureau, “AI Transparency and Disclosure Framework,” January 15, 2026, <https://www.iab.com/guidelines/ai-transparency-and-disclosure-framework/>.

243 IAB Tech Lab, “Agentic Advertising Management Protocols (AAMP),” <https://iabtechlab.com/standards/agentic-advertising-management-protocols/>; IAB Tech Lab, “Agentic Real Time Framework (ARTF),” last updated February 18, 2026, <https://iabtechlab.com/standards/artf/>. AAMP is framed as extending agentic advertising on top of existing ad-tech consent-signal infrastructure, including IAB Tech Lab's Global Privacy Platform and IAB Europe's Transparency and Consent Framework, rather than creating a new substantive privacy or youth-protection regime.



The Architectural Shift Toward Real-Time AI Inference

California has built one of the country's most developed inference-based privacy frameworks, and that progress is significant. The state Attorney General has concluded that inferences a company draws about an identified consumer can be personal information subject to the right to know, and California's 2026 rules define "profiling" broadly, covering automated prediction of a person's economic situation, health, interests, behavior, and location.²⁴⁵ These are meaningful advances for what this report calls the durable profile layer – the records, segments, and scores a company keeps that are linked to an identifiable consumer.²⁴⁶

But commercial AI increasingly operates on the real-time inference layer identified in Layers 1 and 5, which the durable-profile rules are not built to reach: a system classifies a user and shapes what they see in the moment, then may discard the classification rather than storing it as a profile attribute the consumer could access, correct, or delete. Privacy law is strongest when a company holds a discrete record about a person; it is weakest when the harm comes from this capacity to classify and differentially treat people without producing a stable record.²⁴⁷ Because whether such a record exists is itself often undisclosed, and can turn on a single configuration choice, the regulatory response should attach to the capability to infer and differentially treat a person in real time, not to the presence or absence of a stored record.

California's own rules illustrate the gap. Its automated-decision access and opt-out rights attach only to the use of such technology for a "significant decision," and the

advertising on the same consent frameworks this report finds inadequate.²⁴⁴ Across all of it, self-regulation focuses on safety harms – CSAM, self-harm, exploitation, romantic role-play – and no platform has voluntarily restricted emotional inference for commercial purposes, persuadability modeling, model-level commercial influence, or the measurement feedback loop. It addresses the symptoms that generate headlines, not the commercial-persuasion architecture, and it unfolds against the sustained lobbying and political spending documented in this report's Introduction.

244 IAB, "IAB Announces Project Eidos," February 2, 2026, <https://www.iab.com/news/iab-announces-project-eidos/>; IAB, "Project Eidos," <https://www.iab.com/topics/project-eidos/>. Participating companies listed by IAB include Amazon Ads, Google, Meta, TikTok, Roblox Media, The Trade Desk, and Unilever.

245 California Attorney General, Opinion No. 20-303 (March 10, 2022), concluding that internally generated inferences about a consumer may be personal information subject to the CCPA right to know, <https://oag.ca.gov/system/files/opinions/pdfs/20-303.pdf>.

246 California Privacy Protection Agency, CCPA Updates, Cybersecurity Audits, Risk Assessments, Automated Decisionmaking Technology, and Insurance Regulations, § 7001 (definition of "profiling"), https://cppa.ca.gov/regulations/pdf/ccpa_updates_cyber_risk_admt_appr_text.pdf.

247 On AI processing that classifies and treats users without depending on a stored profile of the individual, see Solow-Niederman, *supra* note 48, at 367-83 (the classification challenge); Mühlhoff & Ruschmeier, *supra* note 173; IAPP, "Privacy by Proxy," *supra* note 30. For documented operational evidence of real-time inference at scale, see Irish Council for Civil Liberties, "The Biggest Data Breach," *supra* note 61, describing per-impression bid requests broadcast to AI bidders and evaluated against cohort models in milliseconds; and, historically, Google's Privacy Sandbox / Topics API, designed to assign cohort labels without producing user-identified records in advertising infrastructure (Chrome Developers documentation). For independent analysis of how the Topics API operated and its limits, see Thorin Klosowski, "How To Turn Off Google's 'Privacy Sandbox' Ad Tracking - and Why You Should," Electronic Frontier Foundation (September 28, 2023), <https://www.eff.org/deeplinks/2023/09/how-turn-googles-privacy-sandbox-ad-tracking-and-why-you-should>. The Privacy Sandbox was officially discontinued in April 2025.

final regulations expressly provide that advertising to a consumer is not one – an exclusion that holds even for youth-facing advertising.²⁴⁸ California does more when a system processes the data of children it knows are under 16, treating it as sensitive and subject to risk assessments and opt-in rules;²⁴⁹ but that protection turns on the company’s knowledge of age,²⁵⁰ and still does not convert advertising into a significant decision²⁵¹ or reach first-party ad ranking and optimization that involve no sale or sharing.^{252, 253, 254}

So inference regulation needs a second extension. The first, which California has begun, brings stored inferences about identifiable consumers within privacy law. The second should reach the design and use of inferential systems themselves – through risk assessment, purpose limitation, testing for discriminatory or manipulative outcomes, restrictions on sensitive and youth-related inference, and limits on optimization objectives. For children and teens, the standard should not depend on whether a company admits knowledge of age; it should reach systems that foreseeably classify, target, or optimize against minors or youth-like audiences. Maryland, Colorado, and Connecticut point in useful directions but remain anchored in the processing of personal data. The next frontier is not adding inferences to the list of covered data; it is regulating the systems that generate and act on inferences in real time.²⁵⁵

Existing Frameworks Are Oriented Toward Data Held More Than the Capability That Produces the Harm

The gaps across the legal landscape are not separate drafting oversights; they share a single root. Most existing law is oriented toward the data a company holds about an identifiable person, while the systems described here operate as an integrated, self-reinforcing loop that runs on capability. A handful of proposals have begun to target that capability directly, but enacted law has largely not. Generative AI does not so much place a discrete advertisement before a child as generate the commercial environment the child moves through: a chatbot builds a relationship with a teenager and surfaces a product in the conversation; a recommendation system routes a tween influencer’s branded line to the children who follow her; a shopping assistant writes a paid brand into its reply. That loop runs on personal and non-personal data alike – aggregate behavioral patterns, cohort models trained on millions of others, de-identified signals, the outputs of other people’s interactions – data that AI has only recently made this powerful, and it is optimized continuously for engagement, conversion, and attention, not for a child’s well-being. Three features of how it works defeat the mechanisms current and proposed laws rely on.²⁵⁶

First, influence operates through real-time inference, and that inference is most weakly constrained when it is used for advertising. A system can sort a first-time user into a behavioral cohort – a group the user resembles

248 CCPA final regulation text, § 7001, defining “significant decision” and stating that “significant decision does not include advertising to a consumer.”

249 CCPA statute effective January 1, 2026, Cal. Civ. Code § 1798.140, defining sensitive personal information to include personal information of consumers the business has actual knowledge are under 16, and deeming willful disregard of age to be actual knowledge, https://ccpa.ca.gov/regulations/pdf/ccpa_statute_eff_20260101.pdf.

250 CCPA Regulations, supra note 246, §§ 7150, 7152 (requiring risk assessments for covered processing activities, including processing sensitive personal information).

251 CCPA Regulations, supra note 246, §§ 7150, 7152. Section 7150 requires risk assessments for sale or sharing, processing sensitive personal information, ADMT for significant decisions, specified automated inferences, and specified training uses; § 7152 sets the required contents of the assessment.

252 CCPA statute effective January 1, 2026, Cal. Civ. Code § 1798.140 (actual-knowledge and willful-disregard language for consumers under 16).

253 CCPA Regulations, supra note 246, §§ 7004, 7070 (requiring documented opt-in processes for the sale or sharing of personal information of consumers under 13 and consumers at least 13 and under 16).

254 Article 29 Data Protection Working Party, Opinion 4/2007 on the Concept of Personal Data (WP 136), June 20, 2007, https://ec.europa.eu/justice/article-29/documentation/opinion-recommendation/files/2007/wp136_en.pdf.

255 Maryland Online Data Privacy Act of 2024, SB 541 / Ch. 455, including data minimization, sensitive-data restrictions, and restrictions on targeted advertising to known or should-have-known teens, https://mgaleg.maryland.gov/2024RS/chapters_noln/Ch_455_sb0541E.pdf.

256 Louis Rosenberg, “Generative AI as a Dangerous New Form of Media,” supra note 7; Rosenberg, “The Manipulation Problem: Conversational AI as a Threat to Epistemic Agency,” 2023 CHI Workshop on Generative AI and HCI (GenAICHI 2023), arXiv:2306.11748.

– estimate how persuadable they are, and shape what they see within seconds, drawing on models trained on millions of others without necessarily writing a durable record, and the model it consults need not have been trained on that user’s data, or on identifiable data at all. A conversational system does this turn by turn, as Layer 5 describes, reading in-session responses, inferring an emotional state, and adjusting its commercial framing, and need not save any of it. The same skew appears in delivery: as Layer 3 documents, optimization concentrates youth-heavy delivery even under neutral advertiser parameters, and rules barring advertisers from selecting teen audiences leave the optimization layer that actually finds those audiences unaddressed. Yet California’s automated-decision rules expressly exclude advertising from the decisions that carry the strongest protections, and U.S. frameworks do not directly constrain that inferential capability.²⁵⁷

Second, the operative classification is about a group, not a person, and it is formed and acted on without ever producing the individual record that a knowledge standard assumes. Child-protective law triggers on what a company knows about a particular user – and a model need never establish anything about the individual child to treat them. But a model does not resolve a fact about an individual child; it sorts the child into a pattern such as “body-image vulnerable” or “high purchase intent” – a pattern derived from other people’s data and applied to a child who supplied none of it – and treats the child as that pattern predicts. The treatment is real even though no fact about the individual was established, and because the system never forms the required knowledge, it can infer youth-like traits, optimize against them, and deliver commercial outputs to minors entirely upstream of where the law attaches. There is no file of these moment-to-moment, group-level inferences for a parent, a teenager, or a regulator to request. The same architecture leaves emotional-state inference for commercial purposes unaddressed. Conversational AI can detect loneliness, anxiety, or excitement and adjust what it sells in response – a capability OpenAI’s

“That gap between what these systems can do and what anyone outside the company can see is the strongest case in this report for regulating capability rather than disclosure.”

own researchers monitor at population scale, as Layer 5 documents – yet no U.S. law prohibits inferring a minor’s emotional state to shape what is sold to them.

Third, consent and data-minimization frameworks break down because the model’s power does not come from the data of the child being acted on. This is what Solow-Niederman calls the classification challenge: a model draws its ability to classify and persuade from patterns learned across millions of other users, and can then apply that capability to a child whose own data has been minimized to almost nothing. The objection that better minimization would solve this misunderstands where the capability lives. Even minimization applied to everyone would not reach it, because once the patterns are trained into the model they no longer depend on any individual’s retained record, and the same systems can be built from synthetic data and from inferences drawn across people – neither of which a rule that triggers only when data are collected from a person and linked to them can reach. The harm can also travel in the output itself: as Layer 4 documents, large language models reproduce raciolinguistic and gender stereotypes in the creative they generate, regardless of whose data trained them, so a youth-targeted ad can carry bias into what a child sees without any of that child’s data being collected. A prohibition on training with minors’ personal data removes one input but not the capability, which is learned from the population as a whole and survives the loss of any single group’s

257 Mühlhoff and Ruschemeier, *supra* note 173; Solow-Niederman, *supra* note 48 (classification challenge).

data.²⁵⁸ And Generative Engine Optimization – shaping what a model says to favor particular brands – embeds commercial objectives in the model’s behavior itself, producing influence that is structural rather than transactional and that no “advertisement” disclosure reaches. In each case the harm is carried by what the model can do, not by any record a company keeps about the child – which is why regulating the data inputs leaves the loop intact.

The capability map in Appendix A makes this concrete. Across all nine platforms surveyed, the capacity to infer who a young user is, sort them into commercial cohorts, optimize against them in real time, and treat them differently from the next user is nearly universal – yet almost none of it is knowable from what the companies themselves disclose. The handful of disclosed protections cluster on a single dimension, differential treatment of minors, and exist only where law or regulation has already forced them; the inferential capabilities that no current framework reaches stay dark, and the opacity is deepest exactly where the stakes are highest. That gap between what these systems can do and what anyone outside the company can see is the strongest case in this report for regulating capability rather than disclosure.

What Is Needed

Regulation must target what AI systems can do to minors, not only what data they hold. Data minimization, use limitations, and default design are real progress and remain necessary, but they are not sufficient against a system whose power comes from aggregate data and learned capability rather than any one child’s record, and whose commercial objectives are opaque and operate at scale. What is needed alongside them is structural regulation that constrains what

“Regulation must target what AI systems can do to minors, not only what data they hold.”

these systems are built to do, framed as obligations on the companies that design and profit from them rather than restrictions on minors’ access or parental obligations. The proposals below are organized under four principles. The move that makes them enforceable, running through all four, is a shift from asking what a company knew about an individual user to asking what its systems do to minors as a group.

P1. Make Commercial Influence Legible

Commercial influence should be legible at the point where a young person encounters it – precisely what generative systems erode, dissolving the boundary between an advertisement and a conversation, a recommendation and a paid placement, a creator and a synthetic persona. The first thing the law should make visible is what a system is built to maximize. Every AI system shaping children’s experiences optimizes for something – time on platform, click-through, purchase completion – and that choice determines what it does and to whom, yet no U.S. law requires disclosure of those objectives for systems used in commercial influence directed at minors. The law should require it.²⁵⁹ The Kids Online Safety Act points partway there, requiring covered platforms to disclose what their algorithms optimize for, but reaches only content recommender systems on social media – not advertising optimization, conversational AI, or measurement.²⁶⁰ California should require it for all AI systems likely to reach minors, extending

²⁵⁸ Such a prohibition assumes a model’s capacity to target minors comes mainly from minors’ own data. That may not hold for large general-purpose models, where youth-correlated patterns can also be learned from content about young people and from behaviorally similar adults, so the capability could persist without it. How much any one subpopulation’s data contribute to a trained model’s capabilities is not well established, which is itself reason not to treat the prohibition as sufficient.

²⁵⁹ Rosenberg, “The Manipulation Problem: Conversational AI as a Threat to Epistemic Agency,” *supra* note 256.

²⁶⁰ Kids Online Safety Act (KOSA), S. 1748, 119th Cong. (reintroduced May 2025), <https://www.congress.gov/bill/119th-congress/senate-bill/1748>. Title II requires covered platforms using opaque algorithmic ranking systems to disclose the quantities their algorithms are designed to optimize and the relative importance of those quantities. The requirement reaches algorithmic content selection, including content recommendations, search-result ranking, and social-media post display, but not advertising optimization, conversational AI, or measurement systems.



to the advertising-optimization layer and not only the content-recommendation layer current federal proposals address.

Disclosure of this kind is a minimum necessary policy ask. It serves mainly as a tool for regulators, researchers, technologists, and advocates – not for children, who, as the developmental-vulnerability evidence shows, cannot be expected to act on it – but it does provide an opening to identify, scrutinize, and regulate what these systems are built to do. Second, and more modestly, the law should require that commercial purpose, sponsorship, and paid influence be surfaced in forms a young person can recognize. A light-gray “Ad” label on a chat agent that behaves like any other contact is not legibility in any meaningful sense.

For minors, legibility is the floor, not the ceiling. The evidence here on developmental vulnerability, on the eroded capacity to recognize commercial influence, and on the disparate exposure of youth of color and lower-income youth supports treating minors as a category warranting heightened, and in places categorical, protection – rather than relying on their ability to resist influence engineered to be undetectable.

P2. Prohibit the Most Harmful Capabilities Outright

Where the first principle makes influence visible, the second limits what it may do. Because the harm comes from what a system can do rather than what it stores, prohibitions are stronger tools than defaults or opt-outs, which presume the user already knows what is being done to them. Some commercial practices and some system capabilities should be off-limits when directed at minors.

The clearest prohibition is also the simplest to state. At a minimum, advertising that is targeted or personalized to children should be barred outright – advertising aimed at a child using data or inferences about them, such as what they have done, what they seem to like, or what they appear to be feeling. Maryland already does this for targeted advertising to minors. The stronger position, and the one this report urges, is a bar on all advertising directed at children, whether or not it is targeted, which would also reach an ordinary commercial message placed where children are known to be; Senator Markey’s bill points this way, barring advertising to minors within AI chatbots. The broader bar protects children more fully, but because it reaches even truthful, untargeted messages it sits closer to ordinary commercial speech, so it has to be drawn with care – aimed at the targeting and the commercial arrangement behind an ad rather than at the plain act of saying something truthful about a product. Both the Maryland and the Markey approaches, though, apply only once a company is treated as knowing a user is a minor – the very limit the next principle is built to move past. Read together with a broad definition of advertising, the ban reaches not only transactional ad buys but the structural steering – generative engine optimization, paid placement surfaced inside a conversation, and commercial partnerships embedded in a model’s behavior – through which much commercial influence now travels.

A ban on advertising governs the commercial output. It does not reach the capabilities that sit upstream of any output, which is why a set of capability-based prohibitions must stand alongside it. These are structural limits on specific functions: systems should be barred from optimizing for persuadability, emotional engagement, or measured behavioral change; from inferring emotional

states for commercial purposes; and from acting on predicted vulnerabilities before a user expresses them. Susceptibility-based optimization warrants its own prohibition: modeling which users are most persuadable by which approaches, and tuning the system to exploit that estimate, should not be permitted when directed at minors. The closed loop that measures whether an intervention worked and feeds the result back to refine the next should likewise be barred, because it is how a system learns to manipulate more effectively over time. These prohibitions attach to systems on the effects-based standard the next principle sets out, not on a company's claimed knowledge that any individual user is a minor. They are illustrative categories, not a closed list.²⁶¹

Prohibitions are only as good as the terms that scope them. "Profiling" must cover the inference of psychological, emotional, cognitive, and developmental characteristics and processing at the group or cohort level or in real time, not only the building of a stored individual record. "Personal data" must include inferred, probabilistic, cohort-derived, and synthetic data, the AI-constructed audience segments documented in Layer 3, and the exclusion for publicly available information must be closed, so a model cannot launder sensitive inference through public inputs. And "advertisement" must capture structural commercial influence – Generative Engine Optimization, training-level optimization, model-level commercial partnerships – not only transactional ad buys.

P3. Regulate Systems Based on Their Effects, Not Their Intent or Their Data

The capability prohibitions of the second principle are enforceable only if accountability turns on what a system does, not on what a company knew or intended. The advertising ban can rest on the same footing: a targeted ad is barred because of how it was aimed, whether or not the company admits it knew the user's age, and the broader bar on all advertising, having no targeting to point to, relies instead on the effects-based trigger set out in this principle, alongside the capability prohibitions. Those prohibitions are the harder case, because a system can optimize against youth-like users while never establishing that any one of them is a minor.²⁶² The most important shift is therefore from knowledge standards to outcomes-based accountability. Rather than triggering protections only once a company "knows" it is dealing with a particular minor, the law should treat minors as a group requiring heightened protection and trigger obligations based on a system's effects on that group. If a system's outputs reach minors in disproportionate numbers, the company should bear the burden of showing it took adequate structural steps to prevent the prohibited practices. The disproportionate reach is not itself the violation; it is what shifts the burden onto the company to show it is not running those practices against the children its system reaches.²⁶³ This moves the inquiry from intent to outcomes, removes the age-verification debate without excluding minors from access, and places the burden on the companies that profit from these systems rather than on children and parents.

261 Capability-based prohibition is a general technique, not specific to this domain. This report applies it to unfair commercial influence directed at minors; the same approach could reach other harms to children, such as compulsive engagement, and harms in other domains, such as political manipulation, which lie beyond this report's scope.

262 By early 2026, the two largest consumer AI and social media platforms had both built systems that infer whether a user is a minor from behavioral and account signals rather than from a self-reported age, showing that the capability the "knowledge" standard treats as unavailable is in fact in use. OpenAI's age-prediction model for ChatGPT draws on signals including how a person uses the service over time, how long the account has existed, and the times of day they are typically active. See OpenAI, "Our Approach to Age Prediction," January 20, 2026, <https://openai.com/index/our-approach-to-age-prediction/>; "OpenAI Rolls Out Age Prediction for ChatGPT," CNBC, January 20, 2026, <https://www.cnbc.com/2026/01/20/open-ai-age-prediction-chatgpt.html>. Meta announced comparable technology that places users it estimates to be teenagers into Teen Account protections even when the account lists an adult birthdate. See Meta, "New AI-Powered Age Assurance Measures to Place Teens in Age-Appropriate Experiences," May 5, 2026, <https://about.fb.com/news/2026/05/ai-age-assurance-teens/>; Aisha Malik, "Meta Will Use AI to Analyze Height and Bone Structure to Identify if Users Are Underage," TechCrunch, May 5, 2026, <https://techcrunch.com/2026/05/05/meta-will-use-ai-to-analyze-height-and-bone-structure-to-identify-if-users-are-underage/>.

263 This mirrors disparate-impact analysis in civil-rights law, where a showing of skewed effects shifts the burden to the company.

While the current federal government opposes this kind of effects-based enforcement in civil-rights law and has told its agencies to step back from it, state authorities are not bound by that retreat, which makes California the natural place to enact this standard.²⁶⁴

An effects-based standard also requires measurement: companies should conduct impact and harms assessments both before deployment and on an ongoing basis afterward, and should measure and monitor which audiences a system actually reaches, so that disproportionate effects on minors can be detected and acted on rather than disclaimed.

The effects standard works after the fact, once a pattern of reach can be seen. It pairs naturally with a simpler rule about which services are covered in the first place: the protections should apply to any service likely to be used by minors. Defining coverage this way turns on who actually uses a service, not on what the service says, which keeps it clear of the free-speech problems that have blocked broader attempts – a distinction a federal appeals court drew directly in 2026.²⁶⁵

P4. Accountability Must Fall on the Powerful

The first three principles describe what the law should reach; the fourth concerns who bears responsibility and how interventions are judged. The relationship between a child and a commercial AI system is one of profound power asymmetry: the system draws on models trained on millions of users, runs optimization the child cannot see, and adapts faster than any young person can learn to resist. Interventions should be measured against a single test – whether they actually reduce that exploitation – and any measure that in practice licenses or entrenches it should be rejected. That test rules out the consent carve-outs and safe harbors documented throughout the legal landscape, which convert a protection into something the more

powerful party can satisfy on paper while continuing the practice.

This principle treats the commercial AI system as the integrated loop it is. Chatbot-specific legislation cannot do the work alone, because the other layers of the pipeline sit outside it; the full pipeline, including the data used to train the models, must be in scope so the loop can be disrupted rather than displaced from one interface to another. The core structural rule should be a flat prohibition on the data flows that feed it: external behavioral, profiling, or targeting data may not flow into a chatbot's interactions with a minor, and data generated within those interactions may not flow out to enrich profiles or train targeting systems. This should not be subject to a consent carve-out, because the asymmetry the rule addresses is precisely what makes consent – from a minor, or a parent on a minor's behalf – unreliable.

Holding power accountable also requires visibility into how commercial interests shape the model itself. Companies should disclose the commercial relationships that influence model training and submit to independent audits of commercial influence in model behavior, including how Generative Engine Optimization and paid partnerships shape the outputs a young person receives. Without that visibility, the most consequential form of influence – embedded in a model's behavior rather than placed beside it – remains beyond external scrutiny.

Constitutional Defensibility

The framework is built to be legally defensible as well as protective, and recent court rulings suggest that online safety and free-speech protections can coexist. Most of what it asks for regulates how a system is built and what it is designed to do – infer emotions, optimize for persuadability, run the feedback loop that learns what works – not what a company is allowed to say. Rules

²⁶⁴ Executive Order 14281 (April 23, 2025) directs federal agencies to step back from disparate-impact enforcement, and the Justice Department rescinded its related Title VI regulations in December 2025; state authorities are unaffected. See Seyfarth Shaw LLP, “New Executive Order Directs Federal Agencies to Deprioritize Disparate Impact” (April 24, 2025), <https://www.seyfarth.com/news-insights/new-executive-order-directs-federal-agencies-to-deprioritize-disparate-impact-what-employers-need-to-know-now.html>; Exec. Order No. 14281, “Restoring Equality of Opportunity and Meritocracy,” April 23, 2025, <https://www.whitehouse.gov/presidential-actions/2025/04/restoring-equality-of-opportunity-and-meritocracy/>.

²⁶⁵ In 2026 a federal appeals court held that applying a children's privacy law to services “likely to be accessed by children” does not, by itself, make the law a restriction on speech, because coverage turns on who uses the service rather than on what it publishes. *NetChoice, LLC v. Bonta*, No. 25-2366 (9th Cir. March 12, 2026); see Ariel Fox Johnson, “The Ninth Circuit Provides a Potential Roadmap for Future Child Safety Laws,” Tech Policy Press, May 15, 2026, <https://www.techpolicy.press/the-ninth-circuit-provides-a-potential-roadmap-for-future-child-safety-laws/>.



Joel Muniz via Unsplash

about how a product works rest on firm ground, much like safety standards that govern how a car is made, and in 2025 a federal appeals court upheld California's design-based limits on addictive feeds for minors, treating them as rules about how a product works rather than about what it says.²⁶⁶

Deciding which services the rules cover sits on similar ground. A rule that applies to services likely to be used by children turns on who is in the audience, not on what the service says, and a federal appeals court has held that defining coverage that way does not, by itself, make a law a restriction on speech.²⁶⁷ The lesson from the rules that did not survive is narrower than it first appears: courts struck them down for being too vague for a company to know what was expected of it, not for regulating design. The way to stay on the right side of that line is to define terms clearly, not to regulate less.

Requiring companies to disclose what a system is built to maximize asks for factual information, closer to a nutrition label than to a limit on speech.

The part of the framework that comes nearest to speech is the ban on advertising aimed at children, because advertising is itself a form of expression and courts protect even commercial messages. This is the piece most open to challenge, and the report does not pretend otherwise. Courts have struck down child-protection advertising rules that reached too far and ended up suppressing truthful messages meant for adults.²⁶⁸ The answer is to draw the ban with care – aimed at the targeting and the commercial arrangement behind an ad rather than at the plain act of saying something truthful – and to treat a ban on advertising that is targeted or personalized to children as the firm legal core, with the broader bar as the goal.

The effects-based standard blocks no one's access to any service. It asks a company whose system reaches large numbers of children to show it has taken real steps to prevent the prohibited practices. Taken together, the framework works mainly on how and why these systems do what they do; where it touches what they say, it does so narrowly and for reasons courts have long recognized.

266 A federal appeals court treated California's "addictive feeds" provisions as a design regulation rather than a content regulation, and upheld the law's default private-mode setting for minors under intermediate scrutiny. *NetChoice, LLC v. Bonta*, No. 25-146 (9th Cir. September 9, 2025), <https://cdn.ca9.uscourts.gov/datastore/opinions/2025/09/09/25-146.pdf>.

267 *NetChoice, LLC v. Bonta*, *supra* note 265.

268 The Supreme Court struck down a state's limits on tobacco advertising near schools because they reached too broadly and would have suppressed truthful messages to adults; protecting children, the Court said, does not justify unnecessarily broad limits on adult speech. *Lorillard Tobacco Co. v. Reilly*, 533 U.S. 525 (2001).

BEYOND CONSUMER PROTECTION: A CONCLUSION

Digital media are no longer one part of childhood among many. For most young people they are now the primary environment in which learning, friendship, play, identity, and the first experiments of independence take place, and that role grows larger every year. The harms documented in this report matter because of where they occur: not at the edges of children's lives but at the center of them. And they do not fit neatly into the category of consumer protection. They are larger than that. They operate at three levels at once – on the individual child, on the relationships and choices children share, and on the democratic culture they are growing into – and all three flow from a single underlying capacity these systems have acquired.

The individual child. A child's conversations with an AI companion, the questions typed into a chatbot, and the products searched and considered can all become commercial data events. The infrastructure described in this report turns private expression, emotional disclosure, and ordinary curiosity into structured commercial signals, surveilling children at a depth, intimacy, and continuity no prior advertising medium could achieve. That surveillance is not neutral, because these systems are built around the developmental stage of the children they reach. Heightened reward sensitivity, incomplete impulse control, susceptibility to peer influence, and an identity still in formation – the ordinary features of growing up – function inside an optimization system as performance signals to be acted on. A child who grows up inside environments of total commercial mediation does not merely encounter more advertising; the habits, preferences, and self-perceptions formed in those years take shape under systems whose objective is commercial conversion rather than healthy development.

Children together. The damage does not stop at the individual. When signal extraction is continuous, when profiling models a child's susceptibility, when influence arrives before a child has articulated a want, and when the line between commercial and authentic content is dissolving, the conditions for meaningful independent choice are eroding for all of them. Young people are not navigating an open marketplace; they are moving through environments designed to steer them, and the system does not so much persuade as maneuver. That steering is not distributed evenly. Youth of color face more intensive data collection, ad delivery that pushes unhealthy products toward Black and Hispanic audiences, AI-generated creative that embeds raciolinguistic stereotypes, and lower-quality interactions in conversational systems. These disparities are not incidental; they are structural features of systems trained on data that reflects historical inequality and compounded with every optimization cycle, so that the burdens of the commercial environment fall hardest on the children with the least protection.

Democratic life. At the widest level, the concern is what these systems do to the culture young people are inheriting. The capacity at issue is the ability to single out, classify, predict, influence, and differentially treat people at scale – exercised here by a small number of firms, and exercised whether or not any identifying record is kept. When that capacity governs what young people see, want, and buy, the market is distorted: firms compete not on the merits of what they make but on the sophistication of the systems they use to shape demand. And the spaces where young people once developed habits of independent thought, civic participation, and shared judgment are being progressively absorbed into commercial optimization. Legal scholars have begun to argue that the defining features of these systems are corrosive to the civic institutions democratic life depends on.²⁶⁹ The full

269 See Woodrow Hartzog and Jessica Silbey, "How AI Destroys Institutions," 77 *UC Law Journal* 727 (2026), https://scholarship.law.bu.edu/faculty_scholarship/4179 (arguing that the core affordances of AI systems erode the expertise, deliberation, and human connection on which civic institutions depend).

“A single interaction rarely registers as damage; a childhood spent inside these systems does.”

consequences are not yet measurable, and we do not overstate them. But the commercial capture of the spaces where young people gather and form themselves most likely carries consequences for democratic life that reach well beyond any single transaction.

These three levels are not separate problems. They are the same capability seen at different scales, and at every scale the harm accumulates. A single interaction rarely registers as damage; a childhood spent inside these systems does. Preferences shaped session by session, inequalities deepened with every optimization cycle, and civic spaces absorbed year by year are harms that build rather than announce themselves, which is part of why episodic, incident-driven regulation continues to miss them. The surveillance of one child, the distorted choices of many, and the slow commercialization of civic life all flow from the inferential capacity that commercial AI infrastructure has acquired – the capacity to read a person, sort them, predict how they will respond, and treat them accordingly. This is why the report has argued for regulating that capability rather than chasing its symptoms one at a time. A framework anchored to what data a company stores tends to arrive after the fact, because the mechanism that produces the harm can operate without producing a durable record. Capability-based regulation aims at the mechanism itself.

It is worth being clear about what the alternative looks like, because the goal is not only to stop harmful

practices but to make room for something better. Children and young people need places to grow up that are not organized around surveillance, advertising, and behavioral optimization. That means allocating real resources to build non-commercial digital spaces for learning, play, creativity, and civic life; requiring youth-facing technologies to optimize for well-being, autonomy, privacy, and democratic participation; supporting business models that do not depend on extracting attention; and investing in real-world spaces where children and teens can safely gather, explore, create, and gain independence outside the home.²⁷⁰ Industry has a role in this future that it has so far largely declined: to design for the people it serves rather than for the duration of their attention, and to treat children’s developmental stage as a reason for restraint rather than a target. Policymakers have the role only they can play: to set the boundaries within which these systems operate, so that protecting children is not left to families negotiating alone against infrastructure built to outpace them.²⁷¹

The law has not kept pace. COPPA was enacted in 1998; the machine learning infrastructure at the core of today’s commercial surveillance has been building for decades. Generative AI is new, but the pipeline it runs on is not, and every year of inaction is a year in which the system learns more about more children and becomes harder to reverse. California’s legislative decisions over the next two to three years will set the national standard. This report has set out what that response should contain – capability-based regulation, structural prohibitions for minors, definitional fixes, and an outcomes-based accountability standard designed to reach systems existing law was never built to address. The store stays open. The practice stops. The cost sits with the companies that built these systems and profit from them.²⁷²

270 For one articulation of this affirmative agenda, see Jonathan Haidt, “Take Action,” The Anxious Generation, <https://www.anxiousgeneration.com/take-action>.

271 See Fairplay, “Beyond Brands,” <https://fairplayforkids.org/beyond-brands>.

272 A full accounting of these systems’ externalities is beyond the scope of this report but warrants brief mention. U.S. data center electricity use has tripled over the past decade and is projected to double or triple again by 2028, driven in significant part by AI workloads. Shehabi et al., “2024 United States Data Center Energy Usage Report,” Lawrence Berkeley National Laboratory for the U.S. Department of Energy, December 2024, <https://escholarship.org/content/qt32d6m0d1/qt32d6m0d1.pdf>. A 2025 Nature Sustainability study finds that the U.S. AI-server sector is unlikely to reach net-zero by 2030 without heavy reliance on carbon offsets and water restoration credits whose effectiveness is itself uncertain. Xiao et al., “Environmental Impact and Net-Zero Pathways for Sustainable Artificial Intelligence Servers in the USA,” *Nature Sustainability* 8, no. 12 (2025): 1541–1553, <https://doi.org/10.1038/s41893-025-01681-y>. These costs fall on the public, and on the same generation whose attention the systems are built to capture, while the returns sit with the firms running the infrastructure.

GLOSSARY

Companion glossary to Interactive Generative Media. Definitions describe how each term is used in the report.

Affective computing (emotion recognition)

AI that detects or infers a person's emotional state from cues such as word choice, response timing, facial expression, or voice, and adapts its behavior accordingly. The report treats this as an inference capability: a system that reads loneliness, anxiety, or excitement in the moment can shape what it sells in response, a use largely unaddressed by U.S. law.

Agentic AI

AI that can take actions toward a goal on a user's behalf — searching, comparing, deciding, completing transactions — rather than only generating content in response to a prompt. It extends commercial influence from what a user is shown to what is done for them. Agency is a spectrum, not a category, and much of what industry markets as agentic is a conversational system with limited capacity to act. The report keeps the distinction: a companion bot or a general-purpose chatbot answering questions is conversational, not agentic, and where the report says agentic it means the capacity to act.

Automated decision-making technology (ADMT)

The term California's privacy regulations use for technology that processes personal information to make, or substantially facilitate, a decision about a person. The California Privacy Protection Agency's 2026 rules give consumers opt-out and access rights for ADMT used in significant decisions, but advertising is expressly excluded, a gap the report flags.

Capability-based regulation

A regulatory approach that targets what an AI system can do — single out, classify, predict, influence, or differentially treat a person — rather than what data a company holds about identified individuals. The report argues this is the necessary alternative to frameworks anchored to data possession, which leave the operative mechanism untouched.

Chatbot (conversational AI)

A system that interacts with a user in natural language across a back-and-forth exchange. The report's concern is that within a single conversation a chatbot can classify a user, infer their state, and adapt its responses in real time, functioning as a profiling apparatus even when no conventional profile is stored. See also Agentic AI, which the report distinguishes from conversational systems by the capacity to act on a user's behalf.

Cohort (segment)

A group of users a system treats as alike for a commercial purpose, defined by shared behavior or inferred traits rather than by individual identity. The terms segment and audience segment are used interchangeably with cohort. Cohort or segment classification is the act of sorting users into these groups. In the report's framework it is the second of four capabilities a commercial AI system can apply to a young user. Cohort-based treatment is how disparate outcomes occur through classification, without any individual being personally identified.

Contextual advertising

Advertising matched to the content a user is viewing, or the immediate context of a request rather than to a stored profile of the individual. The targeted-versus-contextual line does heavy regulatory work: several frameworks, and bills such as the CHATBOT Act, place contextual or most-recent-prompt advertising outside their targeted-advertising limits.

Data clean rooms (DCRs)

Shared computing environments where companies purportedly combine data for analysis without exchanging raw records. Marketed as privacy-preserving, but the FTC has stated that by default most are not, and that they can make it easier to identify and track people.

Deterministic vs. probabilistic tracking

Two ways of identifying users. Deterministic tracking uses exact identifiers — an email address, phone number, or device ID — to follow a specific individual with near certainty. Probabilistic tracking uses statistical models to infer who a user is and how they will behave from fragmented signals, without a confirmed identifier. The shift from the first to the second is central to why identification-anchored law no longer reaches much commercial AI.

Differential treatment

The capability of giving different users different commercial outcomes, such as different creative, product sets, prices, frequency, or placement, based on inferred characteristics or cohort rather than on anything the user chose. In the report's framework it is the fourth of four capabilities a commercial AI system can apply to a young user, the point at which inference and classification produce a tangible difference in what a person is shown or offered.

Durable-profile layer

The layer of commercial AI activity that stores data against an identified person over time, building the standing profile that existing privacy law was written to reach. Paired throughout the report with the real-time-inference layer.

Effects-based standard

A rule that judges a system by the outcomes it produces rather than by what a company intended or knew. The report proposes an effects, or disparate-impact, trigger for commercial AI directed at minors, so that a showing of skewed effect, rather than proof of intent or individual identification, activates accountability.

First-party data

Information a company collects directly from its own users through their use of its products, as distinct from data bought from outside parties. In an interactive AI environment it includes the full record of what a child types, asks, and does inside one company's apps and devices.

Generative AI

AI that produces new content — text, images, audio, video, or code — from a prompt, drawing on patterns learned from large training datasets. Commercially, it lets marketing content be generated and adapted continuously rather than produced once and reused.

Generative engine optimization (GEO)

Shaping content so that it surfaces inside AI-generated answers, the way search engine optimization shaped content to rank in traditional search results. Because the resulting placement appears as part of a system's own response rather than as marked advertising, it is unlikely to be recognized as a commercial message.

Google Gemini

Google's flagship conversational AI chatbot, which users can access through a mobile app or a web interface.

Identifiability trigger

The threshold that determines whether a privacy law applies, typically keyed to whether data are linked or linkable to an identified individual. The report's central critique is that AI systems can single out, classify, and influence people without crossing this threshold, so identification-anchored frameworks leave the mechanism untouched.

Identity graph

A linked map of the identifiers, devices, and behavioral signals a company associates with one person or household, used to recognize and target them across contexts. Data brokers such as LiveRamp and Experian build these to operate across channels and partners in real time.

Inference economy

Legal scholar Alicia Solow-Niederman's term for a system in which the commercial value of data lies not in what is collected but in what can be inferred from it. Innocuous data yield sensitive conclusions, and models built on some people are applied to others who never provided data.

Inference of user characteristics

The capability of working out a person's characteristics, preferences, emotional states, or likely behavior without their stating it directly. The inference need not rest on that person's own data. A model trained on many users can infer and act on the traits of someone who supplied little or none of their own. In the report's framework, this is the first of four capabilities a commercial AI system can apply to a young user, kept distinct from profiling: inference is the act of working out characteristics, while profiling adds the adaptation of treatment that follows.

Knowledge standard

The trigger in much child-protective law that attaches duties only once a company knows, or is treated as knowing, that a user is a minor. The report argues that this lets a system infer youth-like traits and optimize against them without ever establishing the required knowledge, and calls for moving from a knowledge trigger to an effects-based one.

Large language model (LLM)

An AI model trained on large bodies of text to predict and generate language, the engine beneath most current chatbots and generative systems. The report's point is that the model itself can carry the capability to classify and influence a user, so the harm need not depend on any stored record.

Persuadability modeling

Ranking users by how susceptible they are to a given message or approach, and selecting the one predicted to move that user most. Functionally a susceptibility assessment, regardless of what a platform calls it.

Predictive privacy

Mühlhoff and Ruschmeier's term for the interest in not having sensitive attributes predicted from other people's data. It reframes the harm as collective: a model trained on many users can infer and act on traits of someone who supplied none of their own data, which individual-rights frameworks struggle to reach.

Privacy by Proxy

AI systems can identify and act on people through inferred or proxy signals rather than directly identifying data. The report uses it to explain why consent- and identification-based frameworks fall short, and why population-level, capability-based regulation is needed.

Profiling

Automated processing that evaluates a person to infer or predict characteristics, then adapts treatment accordingly. The report uses a composite definition reaching psychological, emotional, and developmental inferences, group- or cohort-level processing, and real-time as well as stored profiling, so the term covers in-session inference and not only a saved record.

Raciolinguistic

Concerning the way language and race are perceived together, so that a person's speech is read through racialized assumptions. In AI systems, raciolinguistic stereotypes can be embedded in generated outputs and in how systems respond to different users.

Real-time bidding (RTB)

The automated auction that decides, in the fraction of a second a page or app loads, which advertiser shows an ad to a particular user. Signals about the user are broadcast, bidders evaluate them against learned models and bid, and the request may then be discarded — leaving no durable record even though targeting occurred.

Real-time-inference layer

The layer of commercial AI activity that classifies a user in the moment, from the immediate context of an interaction, to shape what happens next while retaining nothing durable about the individual. Largely beyond the reach of identification-anchored law, and the report's strongest case for capability-based regulation.

Real-time optimization

The capability of adjusting commercial outputs, such as which ad, product, price, or message a user sees, continuously during a single session or campaign rather than fixing them in advance. In the report's framework, it is the third of four capabilities a commercial AI system can apply to a young user. It is narrower than the real-time-inference layer, which concerns where in-the-moment classification sits relative to a stored profile, and distinct from real-time bidding, the split-second auction that places a display ad.

Recommender system

An algorithm that selects and ranks what a user sees next, such as a feed, search results, or suggested content, by optimizing for a chosen objective. The report notes that disclosure rules such as the Kids Online Safety Act reach content recommender systems but not advertising optimization, conversational AI, or measurement.

Synthetic audience pre-testing

The use of AI-generated consumer models to simulate how different audience segments would respond to a message or campaign before it reaches real people. It lets persuasive content be tuned against modeled populations — including youth segments — without testing on actual individuals.

Synthetic data

Artificially generated data that mimic real-world patterns without being collected directly from individuals. Such data can reproduce the biases of their training sources, and can enable persuasive or profiling systems to be built and tuned without access to protected populations.

APPENDIX A1-TABLE:

AI CHAT PRODUCTS: ADVERTISING PRACTICES, DATA USE FOR AD TARGETING, AND YOUTH PROTECTIONS

Companion to the Appendix A2 capabilities table. Current as of June 2026, developed over a three-month period from April to June 2026. This table documents what each platform actually does with AI chat data for advertising across the same nine platforms: what ads it runs inside chat, whether chat data feeds ad targeting, what it discloses, and what protections exist for young users. Read together, the two appendices show both what companies do with the data and what their systems are able to do. For analysis of what these findings mean, see the companion narrative.

Comparative Analysis of Major Platform Chat Products

Methodological note: This analysis includes voice-based AI assistants (e.g., Amazon’s Alexa for Shopping, which now includes the former Alexa+) alongside text-based chatbots.

What this appendix covers: Its companion, Appendix A2-Table, documents capabilities, what these systems can do to a young user (infer who they are, sort them into groups, optimize against them in real time, and treat them differently from the next user) and whether a parent or teen could find that out.

A note on certainty labels: Throughout the table, HIGH, MEDIUM, and LOW indicate how well each finding is supported by the available evidence. HIGH means the practice is confirmed by the company or by strong public documentation; MEDIUM means there is strong indirect evidence but no direct confirmation; LOW means the public evidence is limited.

Company	Chat Product	LLM	Ads Inside Chat?	Chat Data for Ad Targeting?	Key Disclosures & Certainty	Youth-Specific Notes & Protections	Sources
Meta (Facebook, Instagram, Messenger, WhatsApp)	Meta AI	Llama (proprietary variants)	NO - as of June 2026. Meta states ads will not appear inside chatbot conversations. However, Meta CEO Zuckerberg has indicated in-chat ads may come in the future.	YES [HIGH CERTAINTY]. Since Dec 16, 2025, chat interactions used for ad targeting and content personalization across Facebook and Instagram. Includes text and voice. Sensitive topics (religion, politics, health, sexual orientation) excluded. Apr 2026: Meta Advantage+ now integrates anonymized AI chat intent signals from 1B+ Meta AI users across WhatsApp, Messenger, Instagram, and Facebook. Purchase intent expressed in natural language conversations feeds ads targeting. EU GDPR complaints filed with Irish DPC. No opt-out available in the U.S..	Announced Oct 1, 2025 via blog post. Users notified Oct 7. Privacy policy updated: "We'll start using your interactions with AIs to personalize your experiences and ads." Q4 2025: AI video generation tools hit \$10B run-rate; U.S. click-to-message ad revenue grew 50%+ YoY. Certainty: HIGH.	Meta did not disclose AI chat data use for ads until Oct 2025, over two years after launching AI features (Sep 2023). Instagram Teen Accounts default teens to strictest content controls; teens under 16 cannot change without parental permission. No age-specific opt-out for AI chat data use for ad targeting. Meta restricts ads shown to users it identifies as under 18 to age and location, with no gender, interest, or behavioral targeting (effective Feb 2023). That policy, and the protections that follow from it, reach only the users Meta actually recognizes as minors. On WhatsApp, parent-managed accounts for children under 13 (introduced Mar 2026) disable Meta AI and serve no ads on that product, but they are opt-in and parent-set-up. Coalition of 36 advocacy groups urged FTC to halt Meta's practice (Oct 2025). FTC 6(b) inquiry issued (Sep 2025). Youth AI safety document: Meta announced "Our Approach to Teen AI Safety" (Oct 17, 2025), focused on parental supervision over teen AI character interactions; Instagram rollout planned for early 2026. Framework is platform-specific, not framed as an industry standard. Teen Accounts protections do not address AI chat data use for ad personalization. The Oct 2025 announcement followed an Aug 2025 Reuters investigation that documented internal Meta policies permitting chatbots to engage children in conversations that are "romantic or sensual."	[1][2][3] [4][5][43] [58][60] [76][77]
Google	Gemini App (formerly Bard)	Gemini 3.1 Pro	NO in Gemini app currently. Ads ARE served in AI Mode (AI-powered search) and AI Overviews. Apr 2026: Google SVP Nick Fox has clarified the company is 'not ruling out' ads in Gemini, framing AI Mode ad integration as a learning exercise. Softens flat denial from Dec 2025. May 2026: On Alphabet's Q1 2026 earnings call, Chief Business Officer Philipp Schindler confirmed that ads in the Gemini app are under active consideration, with AI Mode serving as the testing ground first; no format or timeline was announced. The Gemini app itself remains ad-free as of June 2026.	UNCLEAR / LIKELY. Google states Gemini chats are "not being used to show ads" currently. Jan 2026: Google launched the Direct Offers pilot inside AI Mode (AI-powered search), letting brands surface targeted discounts when Google's AI judges a shopper has high purchase intent. Feb 2026: Google unveiled a sponsored shopping ad format in AI Mode, placing labeled product recommendations within Srinivasan's annual advertising letter, Feb 11, 2026). AI Max for Search reframes keywords from hard targets to intent signals and adds keywordless matching, using Gemini and landing-page content to match user intent rather than relying on keyword targeting alone. Gemini reached 750M+ MAUs by early 2026. Certainty: MEDIUM.	Google VP Dan Taylor (Dec 8, 2025): "There are no ads in the Gemini app and there are no current plans to change that." Nick Fox (SVP Knowledge and Information) stated Google is learning from AI Mode ad integration and not ruling out future Gemini ad placements. Adweek stands by Dec 2025 reporting that Google briefed advertisers on 2026 Gemini ad plans. Certainty: MEDIUM.	Gemini app available to users 13+ (18+ in some regions). Google Family Link parental controls apply to Gemini access on supervised accounts. No specific disclosures about how teen Gemini chat data is handled relative to advertising. FTC 6(b) inquiry issued to Alphabet (Sep 2025). Youth AI safety document: None comprehensive. Gemini relies on Google's Family Link supervision for under-13 accounts and content-policy filters layered on the adult model. Common Sense Media rated Gemini "high risk" for kids and teens (Sep 2025).	[6][7][8] [9][44] [62][74] [78][79] [80][81]

Company	Chat Product	LLM	Ads Inside Chat?	Chat Data for Ad Targeting?	Key Disclosures & Certainty	Youth-Specific Notes & Protections	Sources
OpenAI	ChatGPT (Free, Go, Plus, Pro, Enterprise tiers)	GPT-5.3 Instant / GPT-5.4 / GPT-5.5 (flagship)	YES [HIGH CERTAINTY]. Ads launched Feb 9, 2026. Appear at bottom of responses for Free and Go tier users in U.S. Plus, Pro, Business, Enterprise, and Edu tiers are ad-free. Mar/Apr 2026: Ads crossed \$100M annualized revenue by Mar 26, 2026, in approximately 6 weeks. Expanding to Canada, Australia, New Zealand. Self-serve ads manager launched Apr 10, 2026, lowering the entry threshold from the \$200K to \$250K enterprise-only floor at launch to \$50K. On May 5, 2026 OpenAI removed the minimum entirely and opened the manager to U.S. advertisers of all sizes. Pilot extended through end of Apr 2026. David Dugan named global head of advertising solutions (Mar 2026). Currently reaching less than 20% of eligible free users; weekly active users approximately 900M (Feb 2026). ChatGPT Atlas Browser and Sora app also included in ad expansion. Jun 2026: UK launch Jun 6, 2026, OpenAI's first European market (free and Go tiers only), with Japan, South Korea, Brazil, and Mexico announced as the next markets. On or around Jun 17, 2026, OpenAI published Ad Tools Terms defining Audience Tools (advertiser first-party data upload) and Creative Tools (AI-generated ad creative); both are policy-announced, not yet confirmed live in the Ads Manager.	YES - ad personalization ON by default. Current and past conversations and ChatGPT memories inform ad selection. Advertisers receive only aggregated data. Mar 9, 2026: Privacy policy updated with expanded detail on ad data collection. Ads personalized using "anonymous engagement signals." Optional contact syncing introduced. Jun 2026: In the EU and UK, ad personalization requires explicit opt-in consent under OpenAI's EU ads policy, in contrast to the default-on setting in the U.S. Certainty: HIGH.	OpenAI (Jan 2026): "Ads do not influence the answers ChatGPT gives you, and we keep your conversations with ChatGPT private from advertisers." Users can toggle ad personalization off. Less than 7% of ads rated "low relevance" by users per OpenAI. Apr 2026: Self-serve ads manager released; three-tier campaign hierarchy and CPM bidding. May-Jun 2026: cost-per-click bidding added alongside CPM; conversion tracking and a Conversions API shipped; cost-per-action bidding activated Jun 5, 2026. Jun 2026: OpenAI added LiveRamp's Conversions API Hub for ad measurement and signal-loss mitigation via privacy-safe, server-to-server connections: LiveRamp (which agreed in May 2026 to be acquired by Publicis, a deal expected to close by end of 2026) is among the earliest third-party ad-tech partners working with OpenAI, after Criteo (the first, Mar 2026). Certainty: HIGH.	OpenAI committed in its Jan 16, 2026 advertising announcement that "during our test, we will not show ads in accounts where the user tells us or we predict that they are under 18." Uses behavioral age-prediction model. Ads not shown near health, mental health, politics topics. Three caveats: commitment scoped to OpenAI's testing phase, has not been declared durable beyond it; relies on age-prediction model OpenAI acknowledges is imperfect (with appeal process for misclassifications); not externally audited. Expanded protections (late 2025-2026): parental controls; content restrictions; blackout hours; safety notifications for acute distress. Shopping features (Oct 2025) raise additional commercial exposure concerns for youth. FTC 6(b) inquiry issued (Sep 2025). Youth AI safety documents: OpenAI is the only LLM provider to have published a comprehensive, branded teen safety blueprint as a model for legislation. Teen Safety Blueprint (Nov 2025) covers age-estimation, U18 content policies, default-to-U18 treatment, parental controls, well-being design. U18 Model Spec published Dec 2025. Separate Child Safety Blueprint focused on AI-enabled CSAM published Apr 2026. Open-source teen safety policy pack for AI developers released Mar 2026.	[10][11] [12][13] [14][15] [45][46] [54][55] [59][68] [71][72] [73][75] [83]

Company	Chat Product	LLM	Ads Inside Chat?	Chat Data for Ad Targeting?	Key Disclosures & Certainty	Youth-Specific Notes & Protections	Sources
Microsoft	Copilot (Bing, Edge, Windows, standalone app)	GPT-4 / OpenAI + proprietary Microsoft models	YES [HIGH CERTAINTY]. Sponsored ads appear within Copilot conversations via Microsoft Advertising. Ads labeled "Sponsored." Performance Max campaigns appear beneath AI-generated responses.	YES. Conversational context from Copilot sessions informs ad targeting. Microsoft analyzes entire conversation thread, not just last query. Copilot sessions show 73% higher click-through rates and 16% stronger conversion rates vs. traditional search. Certainty: HIGH.	Microsoft Advertising openly markets Copilot as an ad surface. "Ad Voice" features explain ad relevance to users. Ads integrated since 2024. Certainty: HIGH.	33% of Copilot users are ages 18-29. Microsoft Family Safety settings can restrict Copilot access for child accounts. Microsoft Family Safety prevents personalized ads to minor account holders in Copilot, although non-personalized ads can still appear. Minor accounts also do not get personalized experiences and conversations are not used for model training. Standard COPPA compliance claimed. Youth AI safety document: None comprehensive. Microsoft's youth-related Copilot policy lives in a support article codifying age limits, parental controls, and the no-personalized-ads policy for minor accounts. Microsoft AI CEO Mustafa Suleyman publicly articulated a differentiated stance against romantic, flirtatious, or erotic AI content (CNN, Oct 23, 2025), framing this as Microsoft's alternative to competitors building "young user" modes.	[16][17] [61][63]
Amazon	Alexa for Shopping (Amazon's unified AI shopping assistant). On May 13, 2026 Amazon retired the standalone Rufus chatbot and merged Rufus and Alexa+ into one product, embedded in the Amazon search bar and on the app, website, and Echo Show. Free to all signed-in U.S. shoppers; no Prime membership or Echo device required.	Amazon Nova + Anthropic Claude models	YES [HIGH CERTAINTY]. Sponsored Products and Sponsored Prompts are integrated into Alexa for Shopping responses; Sponsored Prompts exited beta and became billable under cost-per-click bidding on Mar 25, 2026. Ads are matched to the shopper's conversational query and labeled "Sponsored." On Echo Show (Amazon's Echo devices with a screen), sponsored products appear in shopping conversations; on audio-only Echo devices, contextual product suggestions can be delivered by voice. Rufus, now Alexa for Shopping, reached more than 300 million customers in 2025.	YES. Alexa for Shopping personalizes results using each customer's full Amazon purchase history and past Alexa conversations across all Echo devices, plus browsing patterns, cart contents, and wish lists. Separately, academic research (Iqbal et al., "Your Echos Are Heard," 2022) found that Echo interaction data can reach dozens of advertising partners and that smart-speaker interaction can drive ad-auction bids up to 30 times higher than without Echo data; Amazon disputes that it shares voice data with advertisers. Amazon has stated the assistant drives nearly \$12 billion in incremental annualized sales (2025). Certainty: HIGH.	Amazon: ads aim to help users "discover product options related to their conversation content." Alexa for Shopping recommendations 83% self-serving to Amazon products. Amazon blocks external AI agents from accessing its marketplace to protect its \$56B ad business. Amazon states Alexa interaction records may inform ads on Amazon and other sites, and that users can opt out of interest-based ads. Certainty: HIGH.	SIGNIFICANT YOUTH EXPOSURE. Although the device is set up under an adult Amazon account, once it is configured the assistant responds to anyone who speaks to it, including children, with no per-user age check, login, or literacy required at the point of use. Over 800,000 children had Alexa profiles as of the 2023 FTC action. Amazon paid a \$25M FTC settlement in 2023 for COPPA violations, including retaining children's voice recordings and failing to honor parental deletion requests. The amended COPPA Rule (effective June 23, 2025, with a compliance deadline of Apr 22, 2026) requires separate verifiable parental consent before disclosing children's data to third parties for targeted advertising. Amazon has not disclosed how it meets this for the assistant in shared-household settings. Voice purchasing is on by default and draws on the account holder's payment method; unless a parent sets a voice code or limits purchases to recognized voices, a child in the household can place orders, and the separate Amazon Kids mode, which turns voice purchasing off, must be switched on. Protections therefore depend on parental configuration rather than on any age-aware treatment by default. Youth AI safety document: None AI-specific. Amazon's youth-related AI commitments are CSAM-focused, expressed through participation in the 2024 Generative AI Principles to Prevent Child Abuse and the Thorn Safety by Design initiative.	[18][19] [20][21] [37][38] [39][40] [41][42] [47][64] [69][70] [84]

Company	Chat Product	LLM	Ads Inside Chat?	Chat Data for Ad Targeting?	Key Disclosures & Certainty	Youth-Specific Notes & Protections	Sources
Snap Inc.	My AI + Perplexity AI (in Snapchat)	GPT (OpenAI) + Gemini (Google); Perplexity AI (\$400M deal announced Nov 2025; amicably ended Q1 2026, never fully rolled out)	Snap has tested sponsored links and product recommendations within My AI conversations. Sponsored Snaps appear in Chat inbox alongside My AI. Nov 2025 to Q1 2026: Snap announced a \$400 million cash-and-equity deal to add Perplexity as a conversational search agent in the Chat interface; the integration reached only a subset of users before Snap and Perplexity amicably ended the deal in Q1 2026 without a full rollout. Snap's live chat-monetization play is AI Sponsored Snaps, advertiser AI agents that appear inside users' Chat threads.	LIKELY: My AI collects geolocation even when "Ghost Mode" activated. Snap help center: "Your data may also be used by Snap to improve Snap's products and personalize your experience, including ads." Q4 2025: Snap implemented platform-level age verification in Australia per new law requiring users to be at least 16, resulting in removal of approximately 400,000 accounts. Now testing Apple's Declared Age Range API. Certainty: MEDIUM.	Snap disclosures criticized as inadequate. UK ICO investigated My AI for insufficient data -protection risk assessment (concluded May 2024, Snap required to improve). Utah sued Snap (Jun 2025) over My AI safety and data practices. FTC 6(b) inquiry issued (Sep 2025). Certainty: MEDIUM.	SIGNIFICANT YOUTH EXPOSURE. Snapchat's user base skews heavily young (13+). My AI pinned to top of chat feeds above real friends, cannot be fully removed. Snap Family Center expanded: parents can now see average daily time spent and breakdown by feature; new friend trust signals. Early testing revealed My AI gave harmful advice to minors. Dark patterns used to extract information from children. Utah lawsuit: Snap "profits from unconscionable design features created to addict children." Youth AI safety document: None AI-specific. Snap's teen safety infrastructure (Family Center, Snap Council for Digital Well-Being, Safety Advisory Board) is platform-wide. Snap added parental controls over My AI access in Jan 2024.	[22][23] [24][25] [26][27] [48][82]
xAI (Elon Musk)	Grok (in X / standalone app)	Grok 4.3	Grok integrated into X, which has \$2.26B ad business (2025). Ads appear in X platform surrounding Grok interactions. No confirmed direct ad insertion inside Grok chat responses in production, though Musk announced paid placements inside Grok responses (Aug 2025) and, on Apr 30, 2026, X launched an xAI-powered ad-platform rebuild with Grok handling contextual and semantic ad targeting, announced and being built out rather than independently confirmed in production. Jan 2026: Grok restricted image generation to paying subscribers after CSAM crisis, but restrictions found "porous" by Common Sense Media.	LIKELY: X platform data broadly used for ad targeting. Grok integrated into X's data ecosystem since Mar 2025 merger. Certainty: MEDIUM.	MINIMAL / NONE. xAI responses to press inquiries: auto-replies reading "Legacy Media Lies." No published information about Grok-specific ad- targeting data practices. Certainty: LOW.	CRITICAL SAFETY FAILURES. Common Sense Media rated Grok "among the worst" for child safety (Jan 2026). CDDH: Grok generated 3M+ sexualized images in under 2 weeks, including 23,000+ appearing to depict children. "Kids Mode" exists in mobile app but not on web or X; ineffective even when enabled. No age verification. Regulatory escalation (ongoing through Apr 2026): California AG investigation (Jan 14, 2026); bipartisan 35-AG demand letter (Jan 26, 2026); EU Commission DSA investigation and document preservation order through end of 2026; UK Ofcom investigation under Online Safety Act (Jan 12, 2026); France CSAM investigation; India 72-hour compliance demand; Malaysia temporary block; Brazil 30-day compliance order; Indonesia access contingent on strict requirements. CA Senator Padilla introduced SB 300 citing Grok. FTC 6(b) inquiry issued (Sep 2025). EU Commission separately ordered X to preserve all documents related to Grok and its image tools until at least end of 2026. Youth AI safety document: None published. xAI has published no documentation about Kids Mode or its guardrails.	[15][28] [29][30] [49][50] [51][52] [53][56] [57][65]

Company	Chat Product	LLM	Ads Inside Chat?	Chat Data for Ad Targeting?	Key Disclosures & Certainty	Youth-Specific Notes & Protections	Sources
Anthropic	Claude (claude.ai / Claude app)	Claude Opus 4.8 / Sonnet 4.6 / Haiku 4.5	NO. Anthropic pledged (Feb 4, 2026) that Claude will remain ad-free. No sponsored links, no advertiser influence on responses, no third-party product placements. Caveat noted: “Should we need to revisit this approach, we’ll be transparent.”	NO. Conversations not used for advertising. Revenue from enterprise contracts and paid subscriptions. Certainty: HIGH for current commitment.	Public blog post and Super Bowl campaign (Feb 4, 2026) committed to ad-free model. Sam Altman (OpenAI) called the campaign “dishonest” and Anthropic an “authoritarian company.” Certainty: HIGH.	Anthropic notes “an appreciable portion” of Claude conversations involve sensitive or deeply personal topics. Claude products marketed for 18+. No specific youth-targeted features or parental controls publicly documented. Free tier expanded Feb 2026, potentially increasing youth access. Youth AI safety document: None teen-focused; Claude.ai restricted to 18+. CSAM-focused commitments via April 2024 Safety by Design alignment with Thorn; B2B “Guidelines for Organizations Serving Minors” for API customers building youth-facing products.	[31][32] [33][66]
TikTok (ByteDance)	No standalone AI chatbot; AI features embedded in platform	Proprietary (internal models)	N/A for standalone chat. TikTok’s broader ad ecosystem is AI-driven. AI-generated ad content requires disclosure via AIGC label since Jul 2024.	YES (platform-level). TikTok’s recommendation algorithm and ad system extensively AI-driven. Off-TikTok activity data used for personalization. Certainty: HIGH for platform; N/A for chat-specific.	Mandatory commercial content disclosure for all promotional posts (2025). AI content labeling required. “Disconnect Advertisers” feature lets users stop specific off-TikTok data sharing. Certainty: HIGH.	Personalized ads restricted for teens (13-17). TikTok does not deliver personalized ads to 13-17s in the EEA, UK, Switzerland, Brazil, Mexico, and Colombia (on- or off-TikTok activity); in the U.S., Canada, and the rest of Latin America, advertiser targeting of under-18s is limited to broad parameters (location, language, device). The EEA/UK restrictions date to 2023; U.S. targeting limits to 2024. Minimum age 16 for livestreaming; 18 for receiving gifts. TikTok faces ongoing regulatory scrutiny and state lawsuits regarding youth data practices. Youth AI safety document: None standalone, since TikTok has no chatbot. Detailed “Teen Safety and Wellbeing” advertising policy.	[34][35] [36][67] [85]

ENDNOTES AND SOURCES

- [1] Meta, “Improving Your Recommendations on Our Apps With AI at Meta,” October 1, 2025, <https://about.fb.com/news/2025/10/improving-your-recommendations-apps-ai-meta/>.
- [2] TechCrunch, “Meta Plans to Sell Targeted Ads Based on Data in Your AI Chats,” October 1, 2025, <https://techcrunch.com/2025/10/01/meta-plans-to-sell-targeted-ads-based-on-data-in-your-ai-chats/>.
- [3] *Fortune*, “Meta just Tied your Private AI Chats to its Ad Business. The Next Step? Designing Bots That Keep You Talking, Expert Says,” October 2, 2025, <https://fortune.com/2025/10/02/meta-ai-chatbot-update-exploits-privacy-monetize-chat-data-facebook-instagram-messenger-ray-ban-display-glasses/>.
- [4] EPIC/CDD et al., “Letter Calling for FTC Oversight and Suspension of Meta’s AI Chatbot Advertising Practice,” October 30, 2025, <https://epic.org/documents/letter-calling-for-ftc-oversight-and-suspension-of-metas-ai-chatbot-advertising-practice/>.
- [5] EPIC, “Press Release: Advocates Urge FTC to Halt Meta’s Plan to Use AI Chatbot Data for Ads,” October 30, 2025, <https://epic.org/press-release-advocates-urge-ftc-to-halt-metas-plan-to-use-ai-chatbot-data-for-ads/>.
- [6] Adweek, “Exclusive: Google Tells Advertisers It’ll Bring Ads to Gemini in 2026,” December 8, 2025, <https://www.adweek.com/media/google-gemini-ads-2026/>.
- [7] Search Engine Land, “Google Denies Ads Are Coming to Gemini in 2026,” December 8, 2025, <https://searchengineland.com/google-corrects-report-claiming-ads-are-coming-to-gemini-in-2026-465856>.
- [8] Google Ads, “Google Ads Highlights of 2025,” <https://support.google.com/google-ads/answer/16756291>.
- [9] Camphouse, “Google Says Gemini Will Improve Its Ad Business,” February 2026, <https://camphouse.io/news/google-gemini-ad-business>.
- [10] Euronews, “ChatGPT Will Now Show You Adverts,” February 10, 2026, <https://www.euronews.com/next/2026/02/10/chatgpt-will-now-show-you-adverts-heres-everything-you-need-to-know>.
- [11] OpenAI, “Our Approach to Advertising and Expanding Access to ChatGPT,” January 16, 2026, <https://openai.com/index/our-approach-to-advertising-and-expanding-access/>.
- [12] TechCrunch, “ChatGPT Rolls Out Ads,” February 9, 2026, <https://techcrunch.com/2026/02/09/chatgpt-rolls-out-ads/>.
- [13] The Register, “OpenAI Introduces Ads...for the People!” February 10, 2026, <https://www.theregister.com/on-prem/2026/02/10/openai-introduces-ads-for-the-people/>.
- [14] StubGroup, “How to Advertise on ChatGPT: Complete Guide 2026,” February 2026, <https://stubgroup.com/blog/how-to-advertise-on-chatgpt-the-complete-guide-for-2026/>.
- [15] TechCrunch, “‘Among the Worst We’ve Seen’: Report Slams xAI’s Grok over Child Safety,” January 27, 2026, <https://techcrunch.com/2026/01/27/among-the-worst-weve-seen-report-slams-xais-grok-over-child-safety-failures/>.
- [16] Microsoft Advertising, “Conversational AI Redefines Audience Engagement,” August 6, 2025, <https://about.ads.microsoft.com/en/blog/post/august-2025/73-higher-ctrs-why-advertisers-need-to-pay-attention-to-conversational-ai>.
- [17] PPC Hero, “How Microsoft Is Transforming Search and Advertising with AI,” June 2025, <https://ppchero.com/how-microsoft-is-transforming-search-and-advertising-with-ai/>.
- [18] *Fortune*, “Amazon Says Its AI Shopping Assistant Rufus Is on Pace for \$10B in Sales,” November 2, 2025, <https://fortune.com/2025/11/02/amazon-rufus-ai-shopping-assistant-chatbot-10-billion-sales-monetization/>.
- [19] My Amazon Guy, “Amazon Rufus AI Updates Drive \$10B Sales Lift,” November 21, 2025, <https://myamazonguy.com/news/amazon-rufus-ai-updates/>.
- [20] Amazon Ads, “What Agentic Shopping Means for Advertising,” 2026, <https://advertising.amazon.com/library/news/agentic-shopping-advertising>.
- [21] Retail Technology Innovation Hub, “Rufus and the AI Shopping War,” January 6, 2026, <https://retailtechinnovationhub.com/home/2026/1/5/rufus-and-the-ai-shopping-war-why-amazons-assistant-reveals-the-battle-for-customer-intent>.
- [22] Kinzoo, “What to Know about Snapchat AI in 2025,” <https://www.kinzoo.com/blog/what-to-know-about-the-new-snapchat-ai>.
- [23] Utah Dept. of Commerce, “Utah Sues Snapchat for Unleashing Experimental AI Technology on Young Users,” June 30, 2025, <https://commerce.utah.gov/2025/06/30/utah-sues-snapchat-for-unleashing-experimental-ai-technology-on-young-users-while-misrepresenting-the-safety-of-the-platform/>.
- [24] TACD, “Snapchat’s AI Data Grab: Why Teens Are at Risk,” August 22, 2025, <https://tacd.org/snapchats-ai-data-grab-why-teens-are-at-risk-and-regulators-are-silent/>.
- [25] UK ICO, “We Warn Organisations Must Not Ignore Data Protection Risks - Snap My AI,” May 2024, <https://ico.org.uk/about-the-ico/media-centre/news-and-blogs/2024/05/ico-warns-organisations-must-not-ignore-data-protection-risks-as-it-concludes-snap-my-ai-chatbot-investigation/>.
- [26] FTC, “Statement of the Commission Regarding the Snap Complaint Referral to the DOJ,” January 2025, <https://www.ftc.gov/news-events/news/press-releases/2025/01/statement-commission-regarding-snap-complaint-referral-doj>; Reuters, “US FTC Refers Snapchat Complaint over AI Chatbot to Justice Dept,” January 16, 2025, <https://www.reuters.com/technology/us-ftc-refers-snapchat-complaint-over-ai-chatbot-justice-dept-2025-01-16/>.
- [27] CNN, “Snapchat to Let Parents Decide Whether Teens Can Use AI Chatbot,” January 11, 2024, <https://www.cnn.com/2024/01/11/tech/snapchat-parents-myai-family-center/index.html>.
- [28] The Register, “Anthropic Keeps Claude Ad-Free,” February 4, 2026, <https://www.theregister.com/software/2026/02/04/anthropic-keeps-claude-ad-free/4658051>.
- [29] CNN, “Elon Musk’s AI Chatbot Grok under Fire for Failing to Rein In Digital Undressing,” January 8, 2026, <https://www.cnn.com/2026/01/08/tech/elon-musk-xai-digital-undressing>.
- [30] NBC News, “Dark Web Users Cite Grok as Tool for Making Criminal Imagery of Kids,” January 9, 2026, <https://www.nbcnews.com/tech/internet/grok-x-bikini-make-imagine-ai-elon-musk-rcna252864>.
- [31] CNBC, “Anthropic Takes Aim at OpenAI’s Ad Push in Super Bowl Commercial,” February 4, 2026, <https://www.cnbc.com/2026/02/04/anthropic-no-ads-claude-chatbot-openai-chatgpt.html>.
- [32] Axios, “Anthropic Pledges Ad-Free Claude, with Super Bowl Shot at ChatGPT,” February 4, 2026, <https://www.axios.com/2026/02/04/anthropic-ads-openai-super-bowl>.
- [33] MacRumors, “Anthropic Promises Claude Will Remain Ad-Free,” February 4, 2026, <https://www.macrumors.com/2026/02/04/anthropic-claude-no-ads/>.
- [34] Social Media Today, “TikTok Implements Restrictions on Ads Targeted at Teens,” July 7, 2024, <https://www.socialmediatoday.com/news/tiktok-implements-restrictions-ads-targeted-teens-ai-disclosure/720667/>.
- [35] TikTok Ads, “Protecting Minors on TikTok: Advertising Initiatives,” <https://ads.tiktok.com/help/article/protecting-minors-on-tiktok-advertising-initiatives>.

- [36] Darkroom Agency, "TikTok 2025 Policy Update." <https://www.darkroomagency.com/observatory/what-brands-need-to-know-about-tiktok-new-rules-2026>
- [37] TechCrunch, "Amazon CEO Wants to Put Ads in Your Alexa+ Conversations," August 1, 2025. <https://techcrunch.com/2025/07/31/amazon-ceo-wants-to-put-ads-in-your-alexa-conversations/>
- [38] Thurrott, "Amazon May Introduce Ads in Conversations with Alexa+," August 1, 2025. <https://www.thurrott.com/amazon/324086/amazon-may-introduce-ads-in-conversations-with-alexa>
- [39] The Register, "Study: How Amazon Uses Echo Interactions to Target Ads," April 28, 2022; original study: Iqbal et al., "Your Echos Are Heard." https://www.theregister.com/2022/04/27/amazon_audio_data/
- [40] Laurent Giret, "How Alexa Echo Is Changing the Game in Ad Targeting," Anstrex, August 1, 2025. <https://www.anstrex.com/blog/how-alexa-echo-is-changing-the-game-in-ad-targeting>
- [41] FTC, "Out of the Mouths of Babies?" May 2023; U.S. v. Amazon.com (Alexa), July 2023, \$25M settlement. <https://www.ftc.gov/business-guidance/blog/2023/05/out-mouths-babies-ftc-says-amazon-kept-kids-alexa-voice-data-forever-even-after-parents-ordered>
- [42] DLA Piper, "A New Era for Children's Privacy: FTC Finalizes Changes to COPPA," January 2025. <https://www.dlapiper.com/en/insights/publications/2025/01/ftc-finalizes-changes-to-coppa>
- [43] TechCrunch, "Instagram and Facebook Introduce More Limits on Targeting Teens with Ads," January 10, 2023 (under-18 ads restricted to age and location; gender and in-app activity removed, effective February 2023), <https://techcrunch.com/2023/01/10/instagram-and-facebook-introduce-more-limits-on-targeting-teens-with-ads/>; Meta Business Help Center, "About Advertising to Teens"; WhatsApp Blog, "Introducing Parent-Managed Accounts on WhatsApp," March 11, 2026, <https://blog.whatsapp.com/introducing-parent-managed-accounts-on-whatsapp>
- [44] Storyboard18, "Google Says Irrelevant Ads Reduced by 40% via Gemini," March 12, 2026; HookedOnMedia, "Now Ads Could be Coming to Google Gemini," March 2026 (Nick Fox statement). <https://www.storyboard18.com/amp/advertising/google-gemini-ai-cuts-irrelevant-ads-by-40-percent-boosts-ad-quality-ws-l-91966.htm>
- [45] OpenAI, "Introducing Parental Controls," 2025; OpenAI, "Updating Our Model Spec with Teen Protections," December 2025. <https://openai.com/index/introducing-parental-controls/>
- [46] Search Engine Land, "OpenAI Updates Privacy Policy as Ads Expand in ChatGPT," March 9, 2026. <https://searchengineland.com/openai-updates-privacy-policy-as-ads-expand-in-chatgpt-471150>
- [47] Amazon Ads, "Alexa+ Tests Product Discovery through Conversational AI," August 2025; GeekWire, "Pick Alexa's Personality," February 25, 2026. <https://advertising.amazon.com/resources/whats-new/alexa-test-product-discovery-through-conversational-ai>
- [48] Engadget, "Snap and Perplexity Sign \$400 Million Deal," November 5, 2025; Snap Form 8-K FY2026 (age verification, Australia). <https://www.engadget.com/social-media/snap-and-perplexity-sign-400-million-deal-to-put-ai-search-directly-in-snapchat-221101734.html>
- [49] California AG, "Attorney General Bonta Launches Investigation into xAI, Grok," January 14, 2026. <https://oag.ca.gov/news/press-releases/attorney-general-bonta-launches-investigation-xai-grok-over-undressed-sexual-ai>
- [50] NY AG James, "Attorney General James Demands More Action from xAI," January 26, 2026 (bipartisan 35-AG coalition). <https://ag.ny.gov/press-release/2026/attorney-general-james-demands-more-action-xai-stop-grok-chatbot-producing>
- [51] Al Jazeera, "EU Probes Musk's Grok AI Feature over Deepfakes of Women, Minors," January 26, 2026. <https://www.aljazeera.com/news/2026/1/26/eu-launches-probe-into-grok-ai-feature-creating-deepfakes-of-women-minors>
- [52] CNBC, "India, EU Investigate Musk's X after Grok Created Deepfake Child Porn," January 6, 2026. <https://www.cnbc.com/2026/01/05/india-eu-investigate-musks-x-after-grok-created-deepfake-child-porn.html>
- [53] FTC, Section 6(b) Resolution re Generative AI Companion Products, September 2025. <https://www.ftc.gov/reports/6b-orders-file-special-report-regarding-advertising-safety-data-handling-practices-companies>
- [54] Search Engine Land, "ChatGPT Hits \$100 Million in Ad Revenue," March 26, 2026, <https://searchengineland.com/chatgpt-hits-100-million-in-ad-revenue-and-is-opening-self-serve-access-in-april-472797>
- [55] Digiday / ppc.land, "OpenAI's Ads Manager is Live," April 10-13, 2026, <https://ppc.land/openais-ads-manager-is-live-and-the-barrier-to-entry-just-dropped/>
- [56] TechPolicy.Press, "Regulators are Going after Grok and X," January 26, 2026, <https://www.techpolicy.press/regulators-are-going-after-grok-and-x-just-not-together/>
- [57] ComplexDiscovery, "The Grok Stress Test: Global Regulators Confront AI Sexual Deepfakes," January 2026; EU Commission document preservation order through end of 2026. <https://complexdiscovery.com/the-grok-stress-test-global-regulators-confront-ai-sexual-deepfakes/>
- [58] Digital Applied, "Meta Advantage+ AI Chat Signals Strategy Guide 2026," April 2026, <https://www.digitalapplied.com/blog/meta-advantage-plus-ai-chat-signals-strategy-2026>
- [59] OpenAI, "Our Approach to Advertising and Expanding Access to ChatGPT," January 16, 2026, <https://openai.com/index/our-approach-to-advertising-and-expanding-access/> ("During our test, we will not show ads in accounts where the user tells us or we predict that they are under 18").
- [60] Meta, "Our Approach to Teen AI Safety: Empowering Parents, Protecting Teens," October 17, 2025, <https://about.fb.com/news/2025/10/teen-ai-safety-approach/>; Maxwell Zeff, "Meta Updates Chatbot Rules to Avoid Inappropriate Topics with Teen Users," TechCrunch, August 29, 2025.
- [61] Microsoft Support, "Microsoft Copilot Age Limits and Parental Controls," <https://support.microsoft.com/en-us/topic/microsoft-copilot-age-limits-and-parental-controls-f79b47a6-288a-4513-8c01-afe4d16db900> (minor accounts "do not get shown personalized ads. Ads may still appear in Copilot, but they will not be personalized").
- [62] Common Sense Media, "AI Risk Assessment: Gemini with Teen Protections," September 5, 2025, <https://www.common Sense Media.org/sites/default/files/featured-content/files/csm-ai-risk-assessment-gemini-with-teen-protections-09052025.pdf>
- [63] Clare Duffy, "Microsoft AI CEO: We're Making an AI that you Can Trust your Kids to Use," CNN Business, October 23, 2025. <https://edition.cnn.com/2025/10/23/tech/microsoft-ai-copilot-updates-teen-safety>
- [64] Amazon, "Amazon's Efforts to Combat Child Sexual Abuse Material," 2024 transparency report, <https://www.aboutamazon.com/news/policy-news-views/amazon-csam-transparency-report-2024> (participation in Generative AI Principles to Prevent Child Abuse and Thorn Safety by Design).
- [65] TechCrunch, "'Among the Worst We've Seen': Report Slams xAI's Grok over Child Safety," January 27, 2026 (Common Sense Media risk assessment finding Kids Mode ineffective; xAI has not published documentation about Kids Mode or its guardrails). <https://techcrunch.com/2026/01/27/among-the-worst-weve-seen-report-slams-xais-grok-over-child-safety-failures/>

- [66] Anthropic, “Aligning on Child Safety Principles,” April 23, 2024, <https://www.anthropic.com/news/child-safety-principles>; Anthropic, “Responsible Use of Anthropic’s Models: Guidelines for Organizations Serving Minors,” <https://support.claude.com/en/articles/9307344>.
- [67] TikTok Ads, “Teen Safety and Wellbeing,” <https://ads.tiktok.com/help/article/tiktok-ads-policy-youth-safety> (“We do not handle underage user data to feature personalized ads to them”).
- [68] OpenAI, “Introducing the Child Safety Blueprint,” April 8, 2026, <https://openai.com/index/introducing-child-safety-blueprint/>.
- [69] CNBC, “Amazon Ditches Rufus AI Chatbot in Favor of Alexa Shopping Agent,” May 13, 2026, <https://www.cnbc.com/2026/05/13/amazon-ditches-rufus-ai-chatbot-in-favor-of-alexa-shopping-agent.html> (Amazon retired the standalone Rufus chatbot and merged Rufus and Alexa+ into “Alexa for Shopping,” embedded in the Amazon search bar and on the app, website, and Echo Show; Rufus helped over 300 million customers in 2025; ads to appear where relevant).
- [70] PPC Land, “Amazon Sponsored Prompts Exit Beta and Become Billable under CPC,” March 25, 2026; see also Perpetua, “Alexa for Shopping (Amazon Rufus): The Complete Guide for Brands and Sellers,” 2026, <https://perpetua.io/blog-alexa-for-shopping-amazon-rufus-the-complete-guide-for-brands-and-sellers/> (Sponsored Products and Sponsored Brands prompts inside the assistant became billable under cost-per-click bidding on March 25, 2026).
- [71] PPC Land, “ChatGPT Ads Go Live in the UK as OpenAI Expands Pilot Beyond US,” June 6, 2026, <https://ppc.land/chatgpt-ads-go-live-in-the-uk-as-openai-expands-pilot-beyond-us/> (UK launch June 6, 2026, OpenAI’s first European market; free and Go tiers only; cost-per-click bidding added alongside CPM on May 5, 2026; cost-per-action bidding activated June 5, 2026; conversion tracking and Conversions API shipped).
- [72] Digital Applied, “Ads in ChatGPT: OpenAI’s New Advertiser Toolkit 2026,” June 2026, <https://www.digitalapplied.com/blog/openai-chatgpt-ads-audience-creative-tools-2026-guide> (on or around June 17, 2026, OpenAI published Ad Tools Terms defining Audience Tools for advertisers).
- [73] Digiday, “ChatGPT Ads Land in U.K. as OpenAI Outlines EU Privacy Rules,” June 8, 2026, <https://digiday.com/marketing/chatgpt-ads-land-in-u-k-as-openai-outlines-eu-privacy-rules/> (OpenAI’s EU ads policy, added June 2, 2026, serves ads only to users who explicitly opt in and relies on consent, rather than legitimate interest, as the legal basis for ad personalization).
- [74] gHacks, “Google Considers Ads in Gemini App After Testing Ad Format in AI Mode First,” May 4, 2026, <https://www.ghacks.net/2026/05/04/google-considers-ads-in-gemini-app-after-testing-ad-format-in-ai-mode-first/> (Alphabet Chief Business Officer Philipp Schindler, on the Q1 2026 earnings call, confirmed ads in the Gemini app are under active consideration, with AI Mode as the testing ground first; no format or timeline announced; the Gemini app remains ad-free).
- [75] Peter Adams, “ChatGPT Ads Get Measurement Boost from New LiveRamp Partnership,” Marketing Dive, June 10, 2026, <https://www.marketingdive.com/news/chatgpt-ads-get-measurement-boost-from-new-liveramp-partnership/822416/> (OpenAI advertisers gain access to LiveRamp’s Conversions API Hub for measurement and signal-loss mitigation via privacy-safe, server-to-server connections; LiveRamp is among the earliest third-party ad-tech firms working with OpenAI; Criteo was the first ad-tech partner in March 2026; LiveRamp agreed in May 2026 to be acquired by Publicis Groupe for approximately \$2.2 billion, a deal expected to close by end of 2026).
- [76] Meta, “2026: AI Drives Performance,” Meta Newsroom, January 29, 2026, <https://about.fb.com/news/2026/01/2026-ai-drives-performance/> (in Q4 2025 the combined revenue run-rate of Meta’s video generation tools reached \$10 billion; US click-to-message ads revenue grew more than 50 percent year over year).
- [77] Reuters, “Meta’s AI Rules have let Bots Hold ‘Sensual’ Chats with Kids,” August 14, 2025, <https://www.reuters.com/investigates/special-report/meta-ai-chatbot-guidelines/> (internal “GenAI: Content Risk Standards” document, authenticity confirmed by Meta, permitted chatbots to engage a child in conversations that are romantic or sensual; Meta removed the relevant portions after Reuters’ inquiry).
- [78] Alphabet Inc., “Q4 2025 Earnings Call,” February 4, 2026, <https://abc.xyz/investor/events/event-details/2026/2025-Q4-Earnings-Call-2026-Dr-C033hS6/default.aspx> (CEO Sundar Pichai: the Gemini app surpassed 750 million monthly active users).
- [79] Google, “A New Generation of Ads for the AI Era of Search,” The Keyword, 2026, <https://blog.google/products/ads-commerce/google-marketing-live-search-ads/>, (“Since launching the Direct Offers pilot in January 2026”); Direct Offers debuted at the National Retail Federation, January 11, 2026.
- [80] Bloomberg, “Google Pushes AI Shopping Features in Search and Gemini Chatbot,” February 11, 2026, <https://www.bloomberg.com/news/articles/2026-02-11/google-pushes-ai-shopping-features-in-search-and-gemini-chatbot> (sponsored shopping ad format for AI Mode unveiled in Vidhya Srinivasan’s third annual advertising letter, February 11, 2026).
- [81] Google Ads Help, “About AI Max for Search Campaigns,” <https://support.google.com/google-ads/answer/15910366>, and “How AI Max for Search campaigns works,” <https://support.google.com/google-ads/answer/15910187>; Mike Ryan, “The End of Keywords? Ginny Marvin on the Shift to AI Max for Search,” February 2026, <https://smarter-ecommerce.com/blog/en/google-ads/the-end-of-keywords-ginny-marvin-on-the-shift-to-ai-max-for-search/> (keywords function as a signal of intent; AI Max expands beyond keywords with broad match and keywordless matching rather than eliminating keyword targeting).
- [82] TechCrunch, “Snap Says its \$400M Deal with Perplexity ‘Amicably Ended,’” May 6, 2026, <https://techcrunch.com/2026/05/06/snap-says-its-400m-deal-with-perplexity-amicably-ended/> (Snap Q1 2026 earnings: the relationship ended in Q1 2026 and guidance assumes no Perplexity contribution; the integration was tested with a subset of users and never fully rolled out). Snap subsequently promoted AI Sponsored Snaps as its native in-chat ad format.
- [83] Reuters, “OpenAI’s US Ad Pilot exceeds \$100 Million in Annualized Revenue in Six Weeks,” March 26, 2026, <https://money.usnews.com/investing/news/articles/2026-03-26/openai-us-ad-pilot-exceeds-100-million-in-annualized-revenue-in-six-weeks>; Digital Applied, “ChatGPT Ads Hit \$100M Revenue,” March 28, 2026, <https://www.digitalapplied.com/blog/chatgpt-ads-100m-revenue-self-serve-april-guide> (approximately 900 million weekly active users as of February 2026; self-serve lowered the managed-access minimum to \$50K in April 2026, then removed it entirely on May 5, 2026 when the Ads Manager opened to U.S. advertisers of all sizes).
- [84] GeekWire, “Amazon Unifies Alexa+ and Rufus as AI Rivals Move into Online Shopping,” May 13, 2026, <https://www.geekwire.com/2026/amazon-unifies-alexa-and-rufus-as-ai-rivals-move-into-online-shopping/> (Amazon Q4 2025 earnings, Feb. 7, 2026: Rufus drove nearly \$12 billion in incremental annualized sales in 2025 and was used by more than 300 million customers).
- [85] TikTok, “Protecting Minors on TikTok: Advertising Initiatives,” last updated April 2025, <https://ads.tiktok.com/help/article/protecting-minors-on-tiktok-advertising-initiatives> (no personalized ads to users 13-17 in the EEA, UK, Switzerland, Brazil, Mexico, and Colombia; advertiser targeting limits for under-18s in the U.S., Canada, EEA, UK, Switzerland, and Latin America).

Note: This landscape is evolving rapidly. Policies, features, and ad implementations may change between the date of this analysis and the reader’s review. Always verify current platform documentation.

APPENDIX A1-KEY FINDINGS:

CONVERSATIONAL AI AS AN ADVERTISING ENVIRONMENT

Summary and interpretation: key findings of Appendix A1-Table: AI Chat Products: Advertising Practices, Data Use for Ad Targeting, and Youth Protections. Current as of June 2026.

Children and teenagers are growing up inside a new advertising environment, one built into the conversational AI tools they use every day. The table accompanying this narrative surveys nine major platforms – **Meta**, **Google**, **OpenAI**, **Microsoft**, **Amazon** (Alexa for Shopping), **Snap**, **xAI**, **Anthropic**, and **TikTok** – and documents what each is doing with AI chat interactions for commercial purposes, and what protections, if any, exist for young users.

A note on how to read this analysis. This appendix documents what each platform does with AI chat data. That is a necessary starting point, but not a sufficient one. Read on its own, a data-focused account understates the problem, so this table should be read together with the companion capability analysis (Appendix A2-Key Findings), which examines what these systems are able to do to a young user regardless of whether any data is linked to a name.

The overall picture is this: seven of the nine platforms surveyed are either confirmed to serve ads inside AI chat interfaces, confirmed to use chat interaction data to target ads elsewhere on their platforms, or likely

doing one or both. Only one platform – **Anthropic** – has made a public commitment to remain ad-free.¹ Three platforms have published explicit, AI-specific protections against ad targeting of users identified as minors, though the protections vary substantially in scope: **OpenAI** has committed not to show ads to predicted-U18 accounts during its current testing phase;² **TikTok** prohibits personalized advertising to users ages 13 to 17 in the EU, the UK, and several other markets, and limits targeting of under-18s to broad parameters in the United States and Canada;³ and **Microsoft**, through its Family Safety policy, prevents personalized ads to minor accounts in Copilot, although non-personalized ads can still appear.⁴ **Amazon**, **Google**, **Snap**, and **xAI** have no comparable AI-specific protections disclosed. **Meta** has a narrower, platform-wide safeguard rather than none. Since 2023 it has limited ads shown to users it identifies as under 18 to age and location, excluding gender, interest, and behavioral targeting, and it keeps recognized minors out of the December 2025 chat-to-ads pipeline. That protection reaches only the minors Meta recognizes, so a young person on an adult or age-misrepresented account stays inside the ordinary pipeline. Two patterns

1 Anthropic blog post and Super Bowl campaign, February 4, 2026 (“There are many good places for advertising. A conversation with Claude is not one of them”); see Axios, “Anthropic Pledges Ad-Free Claude, with a Super Bowl Shot at ChatGPT,” February 4, 2026, <https://www.axios.com/2026/02/04/anthropic-ads-openai-super-bowl>.

2 OpenAI, “Our Approach to Advertising and Expanding Access to ChatGPT,” January 16, 2026, <https://openai.com/index/our-approach-to-advertising-and-expanding-access/>; Search Engine Land, “OpenAI Updates Privacy Policy as Ads Expand in ChatGPT,” March 9, 2026.

3 TikTok Ads, “Teen Safety and Wellbeing,” <https://ads.tiktok.com/help/article/tiktok-ads-policy-youth-safety> (“We do not handle underage user data to feature personalized ads to them”); Mediapost, “TikTok Limits Teen Ad-Targeting Methods,” July 9, 2024.

4 Microsoft Support, “Microsoft Copilot Age Limits and Parental Controls,” <https://support.microsoft.com/en-us/topic/microsoft-copilot-age-limits-and-parental-controls-f79b47a6-288a-4513-8c01-afe4d16db900> (minor accounts “do not get shown personalized ads. Ads may still appear in Copilot, but they will not be personalized”).

run through the youth protections detailed below. The protections that exist cluster on one question, whether a platform will target ads at a user it recognizes as a minor, and they tend to appear where law or regulatory pressure has already required them rather than as voluntary choices. They also depend on the platform correctly identifying a user as under 18, so a young person on an adult or age-misrepresented account falls outside them. Documentation is just as uneven. Only OpenAI has published a comprehensive, standalone youth AI safety document pitched as a model for legislation, while the other platforms offer platform-level updates, narrow CSAM-focused commitments, or nothing AI-specific at all.

I. Advertising and Data Collection Practices

Direct in-chat advertising confirmed and operational at three platforms. **OpenAI** launched sponsored ads in ChatGPT on February 9, 2026, appearing at the bottom of responses for users on the free and Go tiers. Ad selection draws on the user's current conversation, past chat history, and stored memories, with personalization on by default.⁵ As of June 2026, OpenAI has expanded the pilot beyond the United States, including the United Kingdom, its first European market, and has opened a self-serve ads manager that lets advertisers buy placements directly. Ad personalization is on by default in the United States; in the United Kingdom and the EU, OpenAI serves personalized ads only to users who give explicit opt-in consent, and users who decline

see contextual ads based on the current conversation rather than a personalized profile.⁶ **Microsoft Copilot** has served sponsored ads since 2024, with Microsoft Advertising reporting 73 percent higher click-through rates for conversational placements versus traditional search.⁷ **Amazon Rufus**, the AI shopping assistant embedded in Amazon's app and website, integrates sponsored products directly into responses, generating AI-written ad copy from product descriptions and campaign keywords. **Amazon Alexa+**, fully rolled out in February 2026, embeds Sponsored Products in shopping conversations on Echo Show screen devices, where ads are labeled as sponsored visually. On audio-only Echo devices, Alexa's older "By the Way" feature delivers contextual product suggestions after conversations without a verbal sponsorship disclosure, a gap that has drawn FTC and academic attention.⁸ On May 13, 2026, Amazon retired the standalone Rufus chatbot and merged Rufus and Alexa+ into a single product, Alexa for Shopping, so these conversational-ad practices now run through one assistant. Among the platforms surveyed, it is the only voice-based assistant, and so the only one delivering these ads in audio-first, shared household settings.⁹

Indirect monetization through chat data is the more consequential model for scale. **Meta** implemented a policy change on December 16, 2025, under which all conversations with Meta AI, across Facebook, Instagram, WhatsApp, and Messenger, in text and voice, feed its ad targeting system. Meta did not disclose this practice until October 2025, more than two years

5 OpenAI, "Our Approach to Advertising." <https://openai.com/index/our-approach-to-advertising-and-expanding-access/>

6 Digiday, "ChatGPT Ads Land in the UK as OpenAI Outlines EU Privacy Rules," June 2026 (personalized ads served only to users who explicitly opt in) <https://digiday.com/marketing/chatgpt-ads-land-in-u-k-as-openai-outlines-eu-privacy-rules/>; PPC Land, "ChatGPT Ads Go Live in the UK as OpenAI Expands Pilot beyond US," June 6, 2026 <https://ppc.land/chatgpt-ads-go-live-in-the-uk-as-openai-expands-pilot-beyond-us/>; OpenAI, "Testing Ads in ChatGPT" (self-serve Ads Manager).

7 Microsoft Advertising, "Conversational AI Redefines Audience Engagement," August 6, 2025. <https://about.ads.microsoft.com/en/blog/post/august-2025/73-higher-ctr-why-advertisers-need-to-pay-attention-to-conversational-ai>

8 Amazon Ads, "Alexa+ Tests Product Discovery Through Conversational AI," August 2025; Amazon Ads, "Conversational Sponsored Tiles Ads on Alexa+," October 2025. On visual "sponsored" labels: gHacks, "Amazon's Echo Show Devices Are Displaying Full-Screen Ads," October 2025; Amazon statement confirming labels. On audio-only disclosure gap: Iqbal et al., "Your Echos Are Heard: Tracking, Profiling, and Ad Targeting in the Amazon Smart Speaker Ecosystem" (University of Washington / UC Davis / UC Irvine / Northeastern, 2022); 10 of the skills surveyed had full privacy disclosure. The Verge, "Amazon's Echo is Now Displaying Full-Screen Ads," October 2025.

9 CNBC, "Amazon Ditches Rufus AI Chatbot in Favor of Alexa Shopping Agent," May 13, 2026 <https://www.cnbc.com/2026/05/13/amazon-ditches-rufus-ai-chatbot-in-favor-of-alexa-shopping-agent.html>; GeekWire, "Amazon Unifies Alexa+ and Rufus as AI Rivals Move into Online Shopping," May 13, 2026. <https://www.geekwire.com/2026/amazon-unifies-alexa-and-rufus-as-ai-rivals-move-into-online-shopping/>

after launching AI features in September 2023.¹⁰ As of April 2026, Meta's Advantage+ advertising system integrates anonymized chat intent signals from more than one billion monthly Meta AI users. Advertisers report 18 percent improvements in return on ad spend from these signals.¹¹ Meta did not make the change opt-in. Every user was included by default, with no setting to turn it off and no way to keep their conversations out of the targeting system. **Google** officially denies using Gemini chat data for advertising but has embedded Gemini into its ad optimization infrastructure across Search and its AI Mode product, which serves shopping ads with Direct Offers inside conversational search results. Google's SVP for Knowledge and Information has signaled that the company is "not ruling out" in-chat advertising in Gemini.¹² On Alphabet's first-quarter 2026 earnings call, Chief Business Officer Philipp Schindler said the company was actively considering ads in the Gemini app, with the format it is testing in AI Mode serving as the proving ground first and likely to carry over. The Gemini app itself remained ad-free as of June 2026.¹³ **Snap** collects geolocation data from My AI even when Ghost Mode is activated, and its own help documentation states that user data "may be used to personalize your experience, including ads."

Who is advertising. The brands entering these spaces are mainstream consumer companies with large youth audiences. Confirmed early advertisers in the ChatGPT

pilot include **Target, Audible, HelloFresh, Albertsons, Adobe, Ford, Mazda, and Williams-Sonoma**. A Sensor Tower analysis of the first two weeks of ChatGPT ad activity found retail brands account for approximately 44 percent of advertisers, with food delivery, entertainment, software, and travel rounding out the top categories.¹⁴ The three largest advertising agency holding companies, **WPP, Omnicom** (with more than 30 client brands in the pilot), and **Dentsu**, are official OpenAI pilot partners, representing hundreds of consumer brands. Ad technology firm **Criteo**, which connects to more than \$1 trillion in annual commerce transactions, joined the ChatGPT pilot in March 2026, reporting that users referred from ChatGPT convert at approximately 1.5 times the rate of other referral channels. In June 2026 OpenAI added LiveRamp as a measurement partner, giving ChatGPT advertisers server-to-server conversion tracking that ties ad exposure to offline purchases. LiveRamp agreed in May 2026 to be acquired by the advertising holding company Publicis, a deal expected to close by the end of 2026.¹⁵ OpenAI's self-serve ads manager launched on April 10, 2026, lowering the entry threshold from the \$200,000 to \$250,000 enterprise-only floor at launch to \$50,000, and on May 5, 2026 OpenAI removed the minimum entirely and opened the manager to U.S. advertisers of all sizes, widening the pool of brands able to reach ChatGPT's roughly 900 million weekly users. The ad pilot crossed \$100 million in annualized revenue within its first six weeks.¹⁶

10 Meta, "Improving Your Recommendations on Our Apps with AI at Meta," October 1, 2025. <https://about.fb.com/news/2025/10/improving-your-recommendations-apps-ai-meta/> Practice effective December 16, 2025. EPIC/CDD et al., "Letter Calling for FTC Oversight," October 30, 2025. <https://epic.org/documents/letter-calling-for-ftc-oversight-and-suspension-of-metas-ai-chatbot-advertising-practice/>

11 Digital Applied, "Meta Advantage+ AI Chat Signals Strategy Guide 2026," April 2026. <https://www.digitalapplied.com/blog/meta-advantage-plus-ai-chat-signals-strategy-2026>

12 HookedOnMedia, "Now Ads Could Be Coming to Google Gemini," March-April 2026 (Nick Fox, SVP Knowledge and Information). *Adweek*, "Google Tells Advertisers it'll Bring Ads to Gemini in 2026," December 8, 2025.

13 gHacks, "Google Considers Ads in Gemini App after Testing Ad Format in AI Mode First," May 4, 2026. Meg Flippin, "Ads Everywhere - Google Says it may Put Advertisements in its Gemini AI App," Yahoo Finance / Benzinga, May 6, 2026, <https://finance.yahoo.com/sectors/technology/articles/ads-everywhere-google-says-may-193015841.html> (reporting CBO Philipp Schindler's remarks on the Q1 2026 earnings call).

14 Sensor Tower analysis, February 9-23, 2026, cited in almcop.com, "Retail & Grocery Brands Claim 44% of ChatGPT Ads," March 9, 2026. Confirmed launch partners: Marketing4eCommerce, February 12, 2026; StubGroup, "How to Advertise on ChatGPT," 2026.

15 Criteo press release, "Criteo Joins OpenAI Advertising Pilot in ChatGPT," March 2, 2026. CNBC, "ChatGPT's Ad Pilot Has the Industry Excited," March 20, 2026 (WPP, Omnicom 30+ clients, Dentsu). Peter Adams, "ChatGPT ads get measurement boost from new LiveRamp partnership," Marketing Dive, June 10, 2026, <https://www.marketingdive.com/news/chatgpt-ads-get-measurement-boost-from-new-liveramp-partnership/822416/> (LiveRamp, among the earliest ad-tech firms working with OpenAI, agreed in May 2026 to be acquired by Publicis Groupe, a deal expected to close by end of 2026).

16 Search Engine Land, "ChatGPT Hits \$100 Million in Ad Revenue," March 26, 2026; Digiday / ppc.land, "OpenAI's Ads Manager Is Live," April 10-13, 2026 (threshold lowered to \$50K). OpenAI blog, March 26, 2026. Launch floor reported at \$200,000 (OpenAI to *Adweek*) to \$250,000 (some agencies); minimum removed entirely on May 5, 2026 when the manager opened to all U.S. advertisers: Digiday, "OpenAI Opens Up ChatGPT Ads Manager to the U.S.," May 6, 2026, <https://digiday.com/marketing/openai-opens-up-chatgpt-ads-manager-to-the-u-s-while-promising-third-party-measurement-cpa-bidding/>.

II. Youth Protections: What the Table Finds

The table documents youth protections, or their absence, across all nine platforms. Three of them, OpenAI, TikTok, and Microsoft, have AI-relevant protections for minors in place; the rest disclose none specific to AI. Where protections exist they share two limits: they apply only to users a platform recognizes as a minor, and several are tied to a testing phase or to a setting a parent must switch on. The company-by-company findings follow.

OpenAI has the most developed protections of any ad-serving platform in the survey, although key elements are scoped to its current testing phase. In its January 2026 advertising announcement, OpenAI committed that “during our test, we will not show ads in accounts where the user tells us or we predict that they are under 18.”¹⁷ The company uses a behavioral age-prediction model to make this determination. OpenAI has also invested in parental controls, allowing parents to link accounts, set content restrictions, establish blackout hours, and receive safety notifications, and has published a Teen Safety Blueprint (November 2025) and a Child Safety Blueprint focused on AI-enabled CSAM (April 2026).^{18, 19} Three caveats apply to the U18 ad exclusion: it is scoped to OpenAI’s testing phase and has not been committed to as a durable policy beyond it; the age-prediction model is acknowledged to be imperfect, with false negatives meaning that actual minors classified as adults will receive ads; and these are self-declared policies that have not been externally audited.

TikTok does not deliver personalized advertising to users ages 13 to 17 in the EEA, the UK, Switzerland, Brazil, Mexico, and Colombia, and in the United States, Canada, and the rest of Latin America it limits advertiser targeting of under-18s to broad parameters such as location, language, and device type.²⁰ TikTok has no standalone AI chatbot; these protections apply to its AI-driven platform ad system.

Microsoft, through Microsoft Family Safety, prevents personalized ads to Copilot accounts classified as minor, although non-personalized ads can still appear and the company discloses no separate prohibition on ad targeting of minors using non-personalized signals.²¹ Microsoft minor accounts also do not receive personalized experiences and conversations are not used for model training. Microsoft has not published a comprehensive AI-specific youth safety blueprint. Microsoft AI CEO Mustafa Suleyman publicly articulated a differentiated stance against romantic, flirtatious, or erotic AI content, even for adults, in an October 23, 2025 CNN interview, framing this as Microsoft’s alternative to competitors building “young user” modes.²²

Meta states an overall under-18 advertising policy under which ads shown to users it identifies as under 18 are restricted to age and location, with no gender, interest, or behavioral targeting.²³ That policy, and the protections that follow from it, reach only the users Meta actually recognizes as minors. On WhatsApp, the parent-managed accounts Meta introduced for children under 13 in March 2026 disable Meta AI and serve no ads on that product, but they are opt-in and set up by a parent. A child on an account that misstates their age

17 OpenAI, “Our Approach to Advertising.” <https://openai.com/index/our-approach-to-advertising-and-expanding-access/>

18 OpenAI, “Protecting Teen ChatGPT Users: OpenAI’s Teen Safety Blueprint,” November 2025, <https://openai.com/index/teen-safety-freedom-and-privacy/>; Ashley Gold, “Exclusive: OpenAI Unveils Blueprint for Teen AI Safety Standards,” Axios, November 6, 2025.

19 OpenAI, “Introducing the Child Safety Blueprint,” April 8, 2026, <https://openai.com/index/introducing-child-safety-blueprint/> (developed with NCMEC, the Attorney General Alliance, and Thorn; focused on AI-enabled CSAM).

20 TikTok Ads, “Teen Safety and Wellbeing.” <https://ads.tiktok.com/help/article/tiktok-ads-policy-youth-safety>

21 Microsoft Support, “Microsoft Copilot Age Limits and Parental Controls.” <https://support.microsoft.com/en-us/topic/microsoft-copilot-age-limits-and-parental-controls-f79b47a6-288a-4513-8c01-afe4d16db900>

22 Clare Duffy, “Microsoft AI CEO: We’re Making an AI that you Can Trust your Kids to Use,” CNN Business, October 23, 2025, <https://us.cnn.com/2025/10/23/tech/microsoft-ai-copilot-updates-teen-safety/>.

23 TechCrunch, “Instagram and Facebook Introduce More Limits on Targeting Teens with Ads,” January 10, 2023 (under-18 ads restricted to age and location; gender and in-app activity removed, effective February 2023), <https://techcrunch.com/2023/01/10/instagram-and-facebook-introduce-more-limits-on-targeting-teens-with-ads/>; Meta Business Help Center, “About Advertising to Teens.” WhatsApp Blog, “Introducing Parent-Managed Accounts on WhatsApp,” March 11, 2026, <https://blog.whatsapp.com/introducing-parent-managed-accounts-on-whatsapp>.

as adult sits outside this regime and inside the ordinary Meta AI and advertising pipeline, because Meta does not recognize them as a minor. Notice, consent, and age-gating all fail at the population the system cannot or will not identify, which is why what matters is what the system is permitted to do with any chat data, regardless of a user's known age. That capability question is developed in the companion analysis, Appendix A2-Key Findings. In October 2025, Meta announced "Our Approach to Teen AI Safety," adding parental controls over teen AI character interactions on Instagram, with rollout beginning in early 2026.²⁴ The announcement followed an August 2025 Reuters investigation that documented internal Meta policies permitting sexual conversations between chatbots and minors and prompted interim training-rule changes.²⁵

Amazon has not publicly disclosed an Alexa+-specific age-gating mechanism for conversational ads. It paid a \$25 million COPPA settlement in 2023 after the FTC and DOJ alleged it retained children's Alexa voice recordings and failed to honor parental deletion requests.²⁶ The amended COPPA Rule (effective June 2025, compliance date April 2026) requires separate verifiable parental consent before disclosing children's personal information to third parties for targeted advertising. Amazon states it will not serve interest-based ads to a child using an Amazon child profile, but it has not disclosed how that protection applies to

Alexa+ conversational ad formats on shared household devices, where a child profile may not be active. Its own ad materials describe Sponsored Products in Alexa+ shopping conversations and Conversational Entertainment Ads on eligible Echo Show devices. Amazon has published no Rufus- or Alexa+-specific youth safety framework addressing AI-driven advertising; its public youth-related generative-AI commitments are concentrated on CSAM prevention.²⁷

Snap's My AI is pinned to the top of the chat interface above real friends and cannot be fully removed. Snap is the subject of a Utah state lawsuit alleging it profits from design features created to addict children, a concluded UK ICO investigation that required improved data protections, and an FTC Section 6(b) inquiry. Snap's \$400 million deal to add Perplexity as a conversational search agent, announced in November 2025, reached only a subset of users and was amicably ended in the first quarter of 2026 without a full rollout. Snap's current in-chat advertising play is AI Sponsored Snaps, advertiser AI agents that appear inside users' Chat threads. In late 2025 Snap also began enforcing platform-level age verification, removing roughly 400,000 accounts in Australia under a new minimum-age law and testing Apple's Declared Age Range API. Snap added parental controls over My AI access in January 2024,²⁸ but has not produced a standalone AI-specific youth safety document; its teen protection

24 Meta, "Our Approach to Teen AI Safety: Empowering Parents, Protecting Teens," October 17, 2025, <https://about.fb.com/news/2025/10/teen-ai-safety-approach/>.

25 Maxwell Zeff, "Meta Updates Chatbot Rules to Avoid Inappropriate Topics with Teen Users," TechCrunch, August 29, 2025, <https://techcrunch.com/2025/08/29/meta-updates-chatbot-rules-to-avoid-inappropriate-topics-with-teen-users/> (interim training-rule changes following Reuters investigation documenting internal Meta policies permitting sexual conversations between chatbots and minors).

26 FTC, U.S. v. Amazon.com (Alexa), July 2023, \$25M settlement for COPPA violations including retaining children's voice recordings and failing to honor parental deletion requests. Amended COPPA Rule, 16 C.F.R. Part 312 (effective June 23, 2025; compliance date April 22, 2026), requiring separate verifiable parental consent before disclosing children's personal information to third parties for targeted advertising, <https://www.federalregister.gov/documents/2025/04/22/2025-05904/childrens-online-privacy-protection-rule>. Amazon, Children's Privacy Disclosure (no interest-based ads to a child using an Amazon child profile). On Alexa+ ad formats, Amazon Ads materials state that Sponsored Products are "automatically integrated into shopping conversations" on Alexa+, where a request for a product recommendation can return both organic and sponsored options, <https://advertising.amazon.com/resources/whats-new/alexa-test-product-discovery-through-conversational-ai>; Conversational Entertainment Ads place sponsored content within entertainment-related voice queries on eligible Echo Show devices, where users can engage by voice or touch to subscribe to, buy, or rent content, <https://advertising.amazon.com/resources/whats-new/conversational-entertainment-ads-on-alexa-plus>.

27 Amazon, "Amazon's Efforts to Combat Child Sexual Abuse Material," 2024 transparency report, <https://www.aboutamazon.com/news/policy-news-views/amazon-csam-transparency-report-2024> (describing participation in Generative AI Principles to Prevent Child Abuse and Thorn Safety by Design).

28 Clare Duffy, "Snapchat to Let Parents Decide Whether Their Teens Can Use the App's AI Chatbot," CNN Business, January 11, 2024. On the Perplexity arrangement (announced November 2025; ended in Q1 2026 without full rollout): TechCrunch, "Snap Says Its \$400M Deal with Perplexity 'Amicably Ended,'" May 6, 2026, <https://techcrunch.com/2026/05/06/snap-says-its-400m-deal-with-perplexity-amicably-ended/>; Engadget, "Snap and Perplexity Sign \$400 Million Deal to Put AI Search Directly in Snapchat," November 5, 2025, <https://www.engadget.com/social-media/snap-and-perplexity-sign-400-million-deal-to-put-ai-search-directly-in-snapchat-221101734.html>. On age verification: Snap, Form 8-K (Australia minimum-age enforcement, approximately 400,000 accounts removed; Apple Declared Age Range API testing).

infrastructure (Family Center, Snap Council for Digital Well-Being, Safety Advisory Board) is platform-wide rather than AI-specific.

Note on xAI/Grok: Grok’s ad-data practices are not disclosed. The platform’s documented failure to prevent generation of sexualized images of minors, which triggered formal investigations in at least eight jurisdictions, is relevant context for any regulatory analysis of AI chat products accessible to children. xAI has published no documentation about its “Kids Mode” or its guardrails, and Common Sense Media’s January 2026 risk assessment described Grok as “among the worst” of major chatbots for child safety, finding Kids Mode ineffective.²⁹

Youth-Specific AI Safety Documents Across the Industry

Among the nine platforms surveyed, OpenAI is alone in having produced a comprehensive, branded, standalone youth AI safety document framed as a model for legislation. OpenAI’s Teen Safety Blueprint (November 2025) covers age estimation, U18 content policies, default-to-U18 treatment when age is uncertain, parental controls, and well-being design, and is pitched explicitly to lawmakers as a starting point for setting AI youth-safety standards.³⁰ OpenAI separately published a Child Safety Blueprint focused on AI-enabled CSAM in April 2026³¹ and an open-source teen safety policy pack for AI developers in March 2026.³² The other platforms have produced either platform-level youth safety

updates that incidentally apply to AI features, narrowly scoped CSAM-focused commitments, or no AI-specific youth safety documentation at all:

Meta announced “Our Approach to Teen AI Safety” on October 17, 2025, a platform-specific framework for parental supervision of teen AI character interactions, with Instagram rollout in early 2026. It is not framed as an industry standard, and followed the August 2025 Reuters investigation noted above.^{33, 34}

Google has no comprehensive youth AI safety document. Gemini relies on Google’s Family Link supervision for under-13 accounts and content-policy filters layered on the adult model. Common Sense Media rated Gemini “high risk” for kids and teens in September 2025.³⁵

Microsoft has not published a youth AI safety blueprint. Its youth-related Copilot policy lives in a support article codifying age limits, parental controls, and a no-personalized-ads policy for minor accounts (with non-personalized ads still permitted).³⁶ The October 2025 CNN stance from Microsoft AI CEO Mustafa Suleyman against romantic or erotic AI content, noted above, is positioning rather than a youth-safety framework.³⁷

Amazon has published no AI-specific youth safety framework for Rufus or Alexa+; its youth-related AI commitments are CSAM-focused (the 2024 Generative AI Principles to Prevent Child Abuse and Thorn Safety by Design).³⁸

29 TechCrunch, “‘Among the Worst We’ve Seen’: Report Slams xAI’s Grok over Child Safety Failures,” January 27, 2026 (citing Common Sense Media risk assessment finding Kids Mode ineffective and noting xAI has not published documentation about Kids Mode or its guardrails). <https://techcrunch.com/2026/01/27/among-the-worst-weve-seen-report-slams-xais-grok-over-child-safety-failures/>

30 OpenAI, “Teen Safety Blueprint,” (November 2025). <https://openai.com/index/introducing-the-teen-safety-blueprint/>

31 OpenAI, “Introducing the Child Safety Blueprint,” April 8, 2026. <https://openai.com/index/introducing-child-safety-blueprint/>

32 OpenAI, “Teen Safety Policies for AI Developers,” March 23, 2026, released on GitHub for use with gpt-oss-safeguard; developed with Common Sense Media and everyone.ai.

33 Meta, “Our Approach to Teen AI Safety.” <https://about.fb.com/news/2025/10/teen-ai-safety-approach/>

34 Zeff, “Meta Updates Chatbot Rules,” TechCrunch, August 29, 2025. <https://techcrunch.com/2025/08/29/meta-updates-chatbot-rules-to-avoid-inappropriate-topics-with-teen-users/>

35 Common Sense Media, “AI Risk Assessment: Gemini with Teen Protections,” September 5, 2025, <https://www.common Sense Media.org/sites/default/files/featured-content/files/csm-ai-risk-assessment-gemini-with-teen-protections-09052025.pdf>; rated Gemini “high risk” for kids and teens.

36 Microsoft Support, “Microsoft Copilot Age Limits and Parental Controls.” <https://support.microsoft.com/en-us/topic/microsoft-copilot-age-limits-and-parental-controls-f79b47a6-288a-4513-8c01-afe4d16db900>

37 Duffy, “Microsoft AI CEO,” CNN Business, October 23, 2025.

38 Amazon, “Efforts to Combat Child Sexual Abuse Material,” 2024 transparency report. <https://www.aboutamazon.com/news/policy-news-views/amazon-csam-transparency-report-2024>

Snap has platform-wide teen safety infrastructure and added parental controls over My AI access in January 2024,³⁹ but has not produced a standalone AI-specific youth safety document.

xAI has published no documentation about Grok’s “Kids Mode” or its guardrails; its practices are opaque. Common Sense Media’s January 2026 risk assessment described Grok as “among the worst” of major chatbots for child safety and found Kids Mode ineffective.⁴⁰

Anthropic restricts Claude.ai to users 18 and over and has not produced a teen-focused AI safety document. Its youth-related commitments are CSAM-focused (the April 2024 Safety by Design alignment with Thorn) and a B2B-oriented “Guidelines for Organizations Serving Minors” for API customers building youth-facing products.⁴¹ In February 2026, Anthropic publicly pledged that Claude will remain ad-free, in a Super Bowl marketing campaign explicitly contrasted with OpenAI’s ad rollout.⁴²

TikTok has detailed teen-specific advertising policies (“Teen Safety and Wellbeing” advertising policy) prohibiting personalized advertising to users ages 13 to 17 in covered jurisdictions⁴³ but has not produced an AI-specific youth safety blueprint, since TikTok does not have a standalone AI chatbot.

The result is a landscape in which OpenAI’s documents are functioning, by default, as the de facto industry template, even where their specific commitments are scoped to a testing phase or depend on behavioral age-prediction tools the company itself acknowledges are imperfect. No competitor has produced a comparable comprehensive document, and no independent body has evaluated whether the OpenAI commitments are being implemented as described.

III. What the Findings Mean

What this means for regulators. The protections catalogued above share a shape. They are organized around data – around whether a platform holds identifying information about a user it knows to be a minor, and whether it chooses to use that information – and they appear mainly where the law has already required them. Data protection is therefore a reasonable place to begin and an incomplete place to stop. The question is not only what a platform knows about a young user, but what it is permitted to do with whatever it knows. Engaging these platforms only on data-protection terms, asking what personal information they collect and whether a child consented, leaves untouched what actually shapes a young user’s experience, which is what the system is able to do with any chat data, identified or not. The companion capability analysis, Appendix A2-Key Findings, documents that mechanism. A regulatory conversation anchored to data possession alone will keep arriving after the fact and keep missing it.

Bottom line. Most of the youth protections in this table rest on two fragile conditions: that a young person has opted in, or that the platform has both recognized their age and chosen to limit how it uses their data. Where neither holds, and for the many children on adult or age-misrepresented accounts neither does, the protections simply do not apply, and the ordinary advertising and data-use machinery runs as it would for any adult. A data-focused reading of this table will therefore understate the exposure, because the deeper question is not what a platform holds about a young user but what it is free to do with it.

See the accompanying table for full platform-by-platform detail, sourcing, and certainty ratings.

39 Duffy, “Snapchat to Let Parents Decide,” CNN Business, January 11, 2024. <https://www.cnn.com/2024/01/11/tech/snapchat-parents-myai-family-center/index.html>

40 TechCrunch, “‘Among the Worst We’ve Seen’: Report Slams xAI’s Grok over Child Safety,” January 27, 2026. <https://techcrunch.com/2026/01/27/among-the-worst-weve-seen-report-slams-xais-grok-over-child-safety-failures/>

41 Anthropic, “Aligning on Child Safety Principles,” April 23, 2024, <https://www.anthropic.com/news/child-safety-principles>; Anthropic, “Responsible Use of Anthropic’s Models: Guidelines for Organizations Serving Minors,” <https://support.claude.com/en/articles/9307344>.

42 Anthropic, ad-free announcement and Super Bowl campaign, February 4, 2026. <https://www.axios.com/2026/02/04/anthropic-ads-openai-super-bowl>

43 TikTok Ads, “Teen Safety and Wellbeing.” <https://ads.tiktok.com/help/article/tiktok-ads-policy-youth-safety>

APPENDIX A2-TABLE:

CAPABILITIES: WHAT THESE SYSTEMS CAN DO TO YOUNG USERS – AND WHAT USERS CAN ACTUALLY FIND OUT

Companion to the Appendix A1-Table (the data-practices table). Current as of June 2026, developed over a three-month period from April to June 2026. This table maps four capabilities – inference, cohort classification, real-time optimization, and differential treatment – across the same nine platforms (Amazon’s Alexa for Shopping appears across its shopping-assistant and voice surfaces), and records separately (a) whether the capability is documented and with what certainty, and (b) whether an ordinary parent or teen could find this out from what the company discloses. “Not disclosed” is a finding, not a missing cell. Read together, the two appendices show both what companies do with the data and what their systems are able to do. For analysis of what these findings mean, see the companion narrative, Appendix A2-Key Findings.

How to read the ratings

Capability: Yes (documented) · Likely (strong indirect evidence) · Unclear (insufficient public evidence) · No (does not occur / no commercial use). For the four capabilities as generic LLM behaviors, the rating reflects documented commercial or operational deployment, not the bare technical possibility.

Consumer accessibility: Plainly disclosed · Disclosed but buried (in policy, settings, or support article) · Only in third-party reporting (press, academic, advocacy, litigation) · Not disclosed (no public account from the company). Shaded cells flag stronger capability or weaker disclosure for quick scanning.

Meta (Facebook, Instagram, Messenger, WhatsApp) – Meta AI

Llama variants. Ads inside chat: not yet, but chat data feeds ad targeting since Dec 16, 2025; no opt-out.

Capability dimension	Capability	Evidence (what / source)	Consumer accessibility	Where disclosed (or not)
1. Inference of user characteristics	Yes	Meta builds detailed user profiles and now folds Meta AI conversation content (text and voice) into them to refine ad targeting. Its Andromeda/ GEM ad systems infer preferences from creative engagement and behavioral signals. Sensitive categories (religion, politics, health, sexual orientation) are stated to be excluded from chat-derived targeting. [Endnote 1, 2, 3]	Disclosed but buried	A consumer can see assigned interest categories via Settings > Ads > Ad Preferences and use “Why am I seeing this ad?” The use of AI-chat content for targeting is stated in a policy update (Oct 2025; effective Dec 16, 2025), but the inference machinery (Andromeda, cohort modeling) appears only in ad-trade press. There is no opt-out for chat-data use. [Endnote 4, 5]
2. Cohort / segment classification	Yes	Advantage+ replaces manual interest audiences with algorithmic audience discovery; the system sorts users into conversion-likelihood segments from first-party behavioral and chat-intent signals across 1B+ Meta AI users. [Endnote 3, 6]	Only in third-party reporting	The shift to algorithmic cohorting and the integration of chat-intent signals into Advantage+ are documented in marketing/ad-tech trade press, not in any consumer-facing disclosure. A parent or teen cannot determine from Meta directly which commercial segments they have been placed in. [Endnote 6]
3. Real-time optimization	Yes	Advantage+ adjusts creative, targeting, placement, and budget continuously during campaign delivery; Andromeda evaluates creative signals in real time to match users. [Endnote 6, 7]	Not disclosed	No consumer-facing surface describes live per-session optimization. Knowable only from advertiser documentation and trade analysis. [Endnote 6]
4. Differential treatment	Yes	Different users receive different ad creative, frequency, and placement based on inferred profile. Meta operates a Variance Reduction System (VRS) for housing/employment/credit ads following a DOJ settlement, an explicit acknowledgment that delivery differs by group. [Endnote 8]	Disclosed but buried	The VRS exists and is described publicly, but only for the three regulated ad categories; differential commercial treatment in ordinary advertising is not disclosed at the individual level. Teens can set some Ad Topic Controls (“See Less”) but cannot see how they are differentially treated. [Endnote 8, 9]

OpenAI – ChatGPT

GPT-5.5 (current flagship, April 2026). Ads launched Feb 9, 2026 (Free/Go tiers); expanded to the UK in June 2026, with personalized ads in the UK and EU served only on explicit opt-in; personalization on by default in the US; under-18 accounts excluded during the test.

Capability dimension	Capability	Evidence (what / source)	Consumer accessibility	Where disclosed (or not)
1. Inference of user characteristics	Yes	OpenAI runs a behavioral age-prediction model that estimates whether an account is under 18 from usage patterns, account age, time-of-day activity, and query types, producing a probability estimate, an inferred sensitive characteristic. [Endnote 10, 11]	Plainly disclosed	Unusually for the industry, OpenAI emailed the policy to users and published a dedicated age-prediction explainer. Users flagged as adults-misclassified can verify via Persona. This is among the most transparent inference disclosures in the survey, though the model is acknowledged to be imperfect and is not externally audited. [Endnote 11, 12]
2. Cohort / segment classification	Likely	Ad selection runs a relevance auction over advertiser-submitted ads matched to conversation context, past chats, and stored memories (when personalization is on). Whether users are placed in reusable commercial segments is not stated; OpenAI says advertisers receive only aggregated metrics. [Endnote 13, 14]	Disclosed but buried	The privacy policy describes context- and memory-based personalization and an “anonymized engagement signals” approach, but does not disclose any audience-segment taxonomy. A user cannot see what cohort, if any, they are in. [Endnote 13]
3. Real-time optimization	Yes	Ads are selected per response by matching current conversation topics to eligible ads, with the most contextually relevant ad surfaced first, live, within the session. [Endnote 14]	Disclosed but buried	The mechanism is described in OpenAI/ secondary explainers (context matching at response time) and personalization can be toggled off in settings; temporary chats generate no ads. The live auction logic itself is not surfaced to ordinary users. [Endnote 14, 15]
4. Differential treatment	Yes	Accounts predicted under 18 are treated differently by default: no ads, restricted sensitive content, and parental-control hooks. Adult Free/ Go users receive personalized ads; paid tiers receive none. [Endnote 11, 13]	Plainly disclosed	The under-18 differential treatment and tier-based ad exposure are stated directly in the announcement and policy. Caveat: the U18 ad exclusion is scoped to the testing phase and not committed as durable. [Endnote 11, 13]

Google – Gemini app

Gemini 3.1 Pro. No ads in the Gemini app currently; ads served in AI Mode / AI Overviews; company is actively considering Gemini-app ads (CBO Philipp Schindler, Q1 2026 earnings call), with AI Mode as the testing ground.

Capability dimension	Capability	Evidence (what / source)	Consumer accessibility	Where disclosed (or not)
1. Inference of user characteristics	Likely	Gemini sits inside Google's broader inference and signals infrastructure. AI Mode uses Gemini to match user intent to advertiser landing pages without keyword targeting (AI Max). Inference capability is evident; commercial use of Gemini-app chat content for ads is officially denied. [Endnote 16, 17]	Not disclosed	Google states Gemini chats are "not being used to show ads," but provides no consumer-facing account of what is inferred from Gemini interactions. The capability is held under a revocable corporate position rather than disclosed or constrained. [Endnote 16]
2. Cohort / segment classification	Likely	Google operates audience-segment systems across its ad stack; AI Max and Direct Offers in AI Mode infer intent and match commercially. Gemini-specific segmentation is undisclosed. [Endnote 17]	Not disclosed	No consumer surface ties Gemini interactions to commercial segments. Ad Settings cover the broader Google account, not the Gemini app specifically. [Endnote 16, 17]
3. Real-time optimization	Yes	AI Mode places shopping ads with Direct Offers inside conversational search results, optimized live to the query; AI Max matches intent dynamically. [Endnote 17]	Only in third-party reporting	Live ad matching in AI Mode is documented in ad-trade and press reporting (e.g., Adweek), not in plain consumer terms. [Endnote 16, 17]
4. Differential treatment	Unclear	Differential commercial treatment within the Gemini app is not established because the company denies ad use there. Family Link applies content-level differences for supervised minors. [Endnote 18]	Not disclosed	No Gemini-specific disclosure of differential commercial treatment. Common Sense Media rated Gemini "high risk" for kids and teens (Sept 2025), an external assessment, not a company disclosure. [Endnote 18]

Microsoft – Copilot

OpenAI + proprietary models. Sponsored ads inside Copilot conversations since 2024; whole-thread context analyzed.

Capability dimension	Capability	Evidence (what / source)	Consumer accessibility	Where disclosed (or not)
1. Inference of user characteristics	Likely	Microsoft Advertising analyzes the entire Copilot conversation thread, not just the last query, to determine ad relevance, implying inference from conversational context. [Endnote 19, 20]	Disclosed but buried	Microsoft Advertising markets Copilot as an ad surface and offers an "Ad Voice" explanation of relevance, but the consumer-facing account of what is inferred is thin and lives in advertiser materials and support articles. [Endnote 19]
2. Cohort / segment classification	Likely	Copilot ads run through Microsoft Advertising's audience and Performance Max infrastructure, which uses audience segmentation. Copilot-specific segment disclosure is absent. [Endnote 19, 20]	Not disclosed	No consumer-facing surface discloses Copilot-derived commercial segments. [Endnote 19]
3. Real-time optimization	Yes	Sponsored placements appear within conversations and beneath AI responses, selected against live conversational context; Microsoft reports markedly higher CTR/conversion for conversational placements. [Endnote 19, 20]	Only in third-party reporting	Performance numbers and live placement logic appear in Microsoft Advertising blogs and trade press, not plain consumer disclosures. [Endnote 19, 20]
4. Differential treatment	Yes	Microsoft Family Safety prevents personalized ads to minor accounts in Copilot (non-personalized ads can still appear) and withholds personalized experiences and model-training use, an explicit differential treatment by account age. [Endnote 21]	Disclosed but buried	The minor-account ad policy is stated in a support article rather than a prominent youth-safety document. A parent must locate Family Safety settings to find it; there is no comprehensive AI-specific youth blueprint. [Endnote 21]

Amazon – Alexa for Shopping: shopping-assistant surface (formerly Rufus)

Amazon Nova + Claude Sonnet via RAG. Sponsored ads in conversation; real-time personalization from shopping history; agentic auto-buy. On May 13, 2026 Amazon merged Rufus and Alexa+ into one product, Alexa for Shopping, sharing a single persistent shopper profile across surfaces; this row covers the shopping-assistant (chat/app/web) surface, analyzed separately because its capability and disclosure profile differ from the voice surface below.

Capability dimension	Capability	Evidence (what / source)	Consumer accessibility	Where disclosed (or not)
1. Inference of user characteristics	Yes	Rufus infers price sensitivity, brand preference, and use-case context from each shopper's individual Amazon activity, browsing, cart, and conversation, then tailors answers. Practitioner/patent analysis indicates multimodal inference (e.g., OCR of listing images). [Endnote 22, 23, 24]	Disclosed but buried	Amazon's own page states Rufus uses "insights from each customer's individual Amazon shopping activity" for "predictive" personalization, general language that does not specify inferred traits like price sensitivity. The detail comes from seller/ad-tech analysis. [Endnote 24, 22]
2. Cohort / segment classification	Likely	Rufus surfaces buyer-segment insight to sellers (e.g., revealing a product's "dominant buyer segment"), indicating users are classified into commercially meaningful groups. [Endnote 25]	Not disclosed	Segment classification is visible to sellers as an optimization output, never to the shopper. A teen cannot learn which buyer segment they have been assigned to. [Endnote 25]
3. Real-time optimization	Yes	Same query yields different product results by shopper history; Rufus runs agentic actions, monitoring prices roughly every 30 minutes and auto-buying at target prices, live optimization at the point of decision. [Endnote 22, 26]	Only in third-party reporting	The per-shopper variability and agentic price-monitoring cadence are documented in seller guides and enterprise analyses, not in consumer disclosures. [Endnote 22, 26]
4. Differential treatment	Yes	Recommendations and the curated product set differ per user based on inferred profile and history. The assistant has historically favored Amazon's own catalog, though as of May 2026 it can also route to external retailers via Shop Direct and complete external purchases via Buy for Me. [Endnote 22, 27]	Not disclosed	No youth protections or differential-treatment disclosures specific to Rufus/Alexa for Shopping. Auto-buy raises autonomous-purchasing concerns for minors with no Rufus-specific age-gating disclosed. [Endnote 27]

Amazon – Alexa for Shopping: voice surface (formerly Alexa+; Echo, smart home, mobile)

Amazon Nova + Claude Sonnet. Fully rolled out Feb 2026; conversational ads on Echo Show; voice ad auctions. Now part of the unified Alexa for Shopping product (merged May 13, 2026); this row covers the voice surface, kept separate because audio-first, screen-optional delivery carries a distinct disclosure profile (no screen, no labels for many placements).

Capability dimension	Capability	Evidence (what / source)	Consumer accessibility	Where disclosed (or not)
1. Inference of user characteristics	Yes	Alexa interaction records (voice commands, usage patterns, device and smart-home activity, media consumption) inform ad relevance; research finds Alexa data raises ad-auction bids dramatically versus no Echo data. [Endnote 28, 29]	Disclosed but buried	Amazon states interaction records "may inform relevant ads," and offers an interest-based-ads opt-out, but academic research finds practice often inconsistent with the stated policy. Voice setting is uniquely opaque: no screen, no labels for many placements. [Endnote 29]
2. Cohort / segment classification	Likely	Echo-derived data is shared into ad-auction systems with 40+ advertisers, implying segmentation/audience use. [Endnote 28]	Only in third-party reporting	The advertiser-sharing and auction-bid findings come from research, not consumer disclosure. [Endnote 28]
3. Real-time optimization	Yes	Sponsored Products auto-integrate into shopping conversations; CEO has said ads will play a broader role in multi-turn conversations; Conversational Entertainment Ads launched Oct 2025. [Endnote 30]	Not disclosed	No consumer disclosure of live in-conversation ad optimization; audio placements may carry no verbal sponsorship disclosure. [Endnote 30]
4. Differential treatment	Yes	Outputs and product suggestions vary by inferred household/user profile. Voice assistants are accessible to children of all ages with no age verification or login. [Endnote 28, 31]	Not disclosed	No age-gating or youth protections disclosed for Alexa+ conversational ads, despite a \$25M 2023 COPPA settlement over children's voice data and the updated COPPA Rule (eff. June 2025) requiring opt-in parental consent for targeted ads. [Endnote 31]

Snap Inc. – My AI (+ AI Sponsored Snaps)

GPT + Gemini. The Perplexity in-Snapchat integration (announced Nov 2025) was amicably ended in Q1 2026 without full rollout; the live in-chat advertising play is now AI Sponsored Snaps, advertiser-operated AI agents inside the Chat tab (announced Apr 28, 2026; Experian alpha partner). Young user base; My AI pinned and not fully removable.

Capability dimension	Capability	Evidence (what / source)	Consumer accessibility	Where disclosed (or not)
1. Inference of user characteristics	Likely	My AI retains access to location data in Ghost Mode where the user has enabled location sharing with Snapchat (Ghost Mode hides location from friends, not from Snap or My AI), and it processes conversational content; Snap states data “may be used to personalize your experience, including ads.” Academic and advocacy work documents emotional engagement of younger teens with My AI. [Endnote 32, 33, 34]	Disclosed but buried	The personalization-including-ads language sits in the help center; Snap’s disclosures were found inadequate by the UK ICO (My AI investigation concluded May 2024). What is inferred is not specified to users. [Endnote 33, 35]
2. Cohort / segment classification	Unclear	Snap runs platform ad segmentation; whether My AI conversation content specifically feeds segment classification is not established in public sources. [Endnote 33]	Not disclosed	No consumer disclosure ties My AI to commercial segments. [Endnote 33]
3. Real-time optimization	Likely	Snap has tested sponsored links and product recommendations within My AI conversations and, as of April 2026, places AI Sponsored Snaps, advertiser-operated AI agents, directly in the Chat tab, with brand agents extended into DMs and an agentic ad suite announced June 18, 2026. [Endnote 32]	Only in third-party reporting	In-conversation sponsored-content testing is documented in press/advocacy, not plain consumer disclosure. [Endnote 32, 36]
4. Differential treatment	Likely	Utah’s June 2025 suit alleges design features built to addict children and that My AI gave harmful advice to minors and used dark patterns to extract information, treatment calibrated to vulnerable young users. Snap is also among the companies named in the FTC’s September 2025 6(b) inquiry into AI companion chatbots’ child- and teen-safety practices. [Endnote 36, 37, 50]	Only in third-party reporting	These findings come from litigation and advocacy reports (Utah AG; Voicebox “Coded Companions”), not Snap disclosures. Parents can limit My AI via Family Center, but cannot see differential treatment. [Endnote 36, 38]

xAI – Grok (in X / standalone)

Grok 4.3. Integrated into X’s ad ecosystem; in-response ad placements announced (Aug 2025) and an xAI-powered ad-platform rebuild (Apr 30, 2026); documented child-safety failures.

Capability dimension	Capability	Evidence (what / source)	Consumer accessibility	Where disclosed (or not)
1. Inference of user characteristics	Likely	Grok is integrated into X’s data ecosystem (since the March 2025 merger), which performs extensive ad-targeting inference; the April 2026 ad-platform rebuild has Grok itself performing contextual and semantic ad targeting. Grok-specific inference for commercial purposes remains undocumented by xAI. [Endnote 39]	Not disclosed	xAI has published no information about Grok-specific ad-targeting data practices; press inquiries reportedly met with auto-replies. Lowest transparency in the survey. [Endnote 40]
2. Cohort / segment classification	Unclear	X operates ad segmentation; Grok-specific contribution is unestablished publicly. [Endnote 39]	Not disclosed	No public information. [Endnote 40]
3. Real-time optimization	Likely	Musk announced paid placements inside Grok responses (Aug 2025); on Apr 30, 2026 X launched an xAI-powered ad-platform rebuild with Grok handling contextual and semantic targeting, auction, and ranking. Live in-response ad optimization is announced and being built out rather than independently confirmed in production. [Endnote 39]	Not disclosed	No public information. [Endnote 40]
4. Differential treatment	Unclear	Commercial differential treatment not established. A nominal “Kids Mode” exists in the mobile app (not on web/X) but was found ineffective and “porous” by Common Sense Media. [Endnote 40, 41]	Not disclosed	xAI has published no documentation about Kids Mode or its guardrails, despite regulatory escalation across multiple jurisdictions. Opacity here coincides with documented safety failures. [Endnote 40, 41]

Anthropic – Claude

Claude Opus 4.8 / Sonnet 4.6 / Haiku 4.5. Public pledge (Feb 4, 2026) to remain ad-free; marketed 18+.

Capability dimension	Capability	Evidence (what / source)	Consumer accessibility	Where disclosed (or not)
1. Inference of user characteristics	Likely (capability) / No commercial use	As an LLM, Claude processes conversational context and can infer characteristics in-session; Anthropic notes an “appreciable portion” of conversations involve sensitive or personal topics. No use of this for advertising. [Endnote 42, 43]	No commercial use disclosed	Anthropic states conversations are not used for advertising and Claude carries no ads, with the caveat “Should we need to revisit this approach, we’ll be transparent.” The capability exists but is governed by a revocable pledge, not an external constraint. [Endnote 42, 44]
2. Cohort / segment classification	No (commercial)	No advertising segmentation; revenue is from enterprise and subscriptions. [Endnote 42]	No commercial use disclosed	No ad segments to disclose under current policy. [Endnote 42]
3. Real-time optimization	No (commercial)	No advertiser influence on responses; no ad selection occurs. [Endnote 42]	No commercial use disclosed	Stated publicly in the ad-free commitment and Super Bowl campaign (Feb 4, 2026). [Endnote 43, 44]
4. Differential treatment	No (commercial)	No commercial differential treatment. Claude.ai is marketed 18+; no youth-specific features or parental controls publicly documented. [Endnote 42, 45]	No commercial use disclosed	Free tier expanded Feb 2026, which may increase youth access despite the 18+ posture; youth handling relies on B2B “Guidelines for Organizations Serving Minors” for API customers. [Endnote 45]

TikTok (ByteDance)

No standalone chatbot; AI features embedded platform-wide. Personalized ads prohibited for users 13–17 in covered regions.

Capability dimension	Capability	Evidence (what / source)	Consumer accessibility	Where disclosed (or not)
1. Inference of user characteristics	Yes (platform)	TikTok’s recommendation and ad systems are extensively AI-driven and infer interests from on- and off-platform activity. No standalone chatbot, so inference is platform-level rather than conversational. [Endnote 46, 47]	Disclosed but buried	TikTok offers a “Disconnect Advertisers” control and discloses use of off-TikTok activity, but the inference detail is general. [Endnote 47]
2. Cohort / segment classification	Yes (platform)	The platform ad system uses audience segmentation for adults; for teens, only broad parameters (location, language, device) are permitted. [Endnote 48]	Disclosed but buried	The teen restriction is stated in TikTok’s teen-safety ad policy; the adult segmentation taxonomy is not user-visible. [Endnote 48, 49]
3. Real-time optimization	Yes (platform)	Recommendation and ad delivery optimize continuously to engagement. AI-generated ad content requires an AIGC label (since July 2024). [Endnote 46]	Disclosed but buried	Mandatory commercial-content and AI-content labeling are disclosed; the live optimization logic is not surfaced to users. [Endnote 46]
4. Differential treatment	Yes (with youth limit)	Adults receive personalized targeting; users 13–17 are limited to broad targeting in the U.S., Canada, EEA, UK, Switzerland, and Latin America (EEA/UK since 2023; U.S. since 2024), an explicit youth differential. [Endnote 48]	Plainly disclosed	The prohibition on personalized ads for 13–17s is clearly stated in a dedicated teen-safety advertising policy, the clearest youth differential-treatment disclosure in the survey. [Endnote 48, 49]

ENDNOTES AND SOURCES

- [1] Meta, “Improving your Recommendations on Our Apps with AI at Meta,” October 1, 2025, <https://about.fb.com/news/2025/10/improving-your-recommendations-apps-ai-meta/>.
- [2] Devesh Beri, “Meta Finds Another Way to Show You Targeted Ads,” Yahoo/ExtremeTech, 2025, <https://tech.yahoo.com/ai/meta-ai/articles/meta-finds-another-way-show-120000570.html> (Meta AI conversations added to ad profiles effective December 16, 2025; no opt-out).
- [3] DigitalApplied, “Meta Advantage+ AI Chat Signals: Strategy Guide 2026,” March 31, 2026, <https://www.digitalapplied.com/blog/meta-advantage-plus-ai-chat-signals-strategy-2026> (Advantage+ integrates anonymized AI-chat intent signals from 1B+ Meta AI users; 18% ROAS improvement reported).
- [4] TechCrunch, “Meta Improves its Consumer-Facing Tool that Explains Why You’re Seeing that Ad,” <https://techcrunch.com/> (“Why am I seeing this ad?” summarizes on/off-platform activity feeding machine-learning ad models).
- [5] EPIC/CDD et al., “Letter Calling for FTC Oversight and Suspension of Meta’s AI Chatbot Advertising Practice,” October 30, 2025, <https://epic.org/documents/letter-calling-for-ftc-oversight-and-suspension-of-metas-ai-chatbot-advertising-practice/>.
- [6] ALM Corp, “Meta’s AI-Driven Advertising Infrastructure: Inside Andromeda and GEM Systems,” April 20, 2026, <https://almcorp.com/blog/meta-andromeda-gem-ai-advertising-system-guide/> (Andromeda evaluates creative signals to identify likely-engaging users; real-time optimization).
- [7] Adligator, “Meta Broad Targeting 2026: Why Advantage+ Audiences Replace Interest Targeting,” April 11, 2026, <https://adligator.com/blog/meta-broad-targeting-advantage-plus-audiences-2026> (algorithmic audience discovery replaces manual segments; continuous adaptation).
- [8] TechRadar, “Meta’s New Ad Policy Further Protects Teen Privacy and Tackles Discrimination,” <https://www.techradar.com/news/metas-new-ad-policy-further-protects-teen-privacy-and-tackles-discrimination> (Variance Reduction System for housing/employment/credit following DOJ settlement; teen Ad Topic Controls).
- [9] Meta Ad Preferences / “Why am I seeing this ad?” consumer tooling; interest categories at Settings & Privacy > Settings > Ads > Ad Settings <https://about.fb.com/news/2023/02/increasing-our-ads-transparency/>.
- [10] OpenAI, “OpenAI’s Privacy Update,” effective February 2026, summarized in Analyst’s Corner, February 17, 2026, <https://medium.com/analysts-corner/openais-privacy-update-reveals-the-endgame-fd096256c5d8> (active age-prediction model; new inferred-data category).
- [11] CNBC, “OpenAI is Rolling Out Age Prediction for ChatGPT Consumer Plans,” January 20, 2026, <https://www.cnbc.com/2026/01/20/open-ai-age-prediction-chatgpt.html>, (behavioral + account-level signals; under-18 protections; Persona verification).
- [12] OpenAI advertising announcement, January 16, 2026 (commitment not to show ads to predicted-under-18 accounts during the test; scoped to testing phase) <https://openai.com/index/our-approach-to-advertising-and-expanding-access/>.
- [13] SearchEngineLand, “OpenAI Updates Privacy Policy as Ads Expand in ChatGPT,” March 9, 2026, <https://searchengineland.com/openai-updates-privacy-policy-as-ads-expand-in-chatgpt-471150> (personalization via in-ChatGPT signals; advertisers receive only aggregated metrics).
- [14] ALM Corp, “ChatGPT Advertising: Complete Implementation Guide,” February 11, 2026, <https://almcorp.com/blog/chatgpt-advertising-implementation-guide-privacy-business-impact-2026/> (per-response relevance auction over current context, past chats, memories; under-18 accounts receive no ads).
- [15] ALM Corp, “OpenAI’s Privacy Policy Update and ChatGPT Ads Expansion,” March 10, 2026, <https://almcorp.com/blog/openai-privacy-policy-update-chatgpt-ads-2026/> (settings controls; temporary chats generate no ads).
- [16] Adweek, “Exclusive: Google Tells Advertisers It’ll Bring Ads to Gemini in 2026,” December 8, 2025 <https://www.adweek.com/media/google-gemini-ads-2026/>; Search Engine Land, “Google Denies Ads are Coming to Gemini in 2026,” December 8, 2025 (denial softened to “not ruling out” per SVP Nick Fox) <https://searchengineland.com/google-corrects-report-claiming-ads-are-coming-to-gemini-in-2026-465856>. On the company moving to active consideration: gHacks, “Google Considers Ads in Gemini App after Testing Ad Format in AI Mode First,” May 4, 2026 (CBO Philipp Schindler, Q1 2026 earnings call; Gemini app still ad-free, AI Mode as testing ground) <https://www.ghacks.net/2026/05/04/google-considers-ads-in-gemini-app-after-testing-ad-format-in-ai-mode-first/>.
- [17] Google Ads, AI Mode / AI Max documentation (shopping ads with Direct Offers inside conversational search; keyword-free intent matching), 2026 <https://blog.google/products/ads-commerce/google-ai-max-for-search-campaigns/>.
- [18] Common Sense Media, Gemini risk assessment rating Gemini “high risk” for kids and teens, September 2025 <https://www.common SenseMedia.org/sites/default/files/featured-content/files/csm-ai-risk-assessment-gemini-with-teen-protections-09052025.pdf>; Google Family Link supervision documentation <https://support.google.com/gemini/answer/16109150>.
- [19] Microsoft Advertising, “Conversational AI Redefines Audience Engagement,” August 6, 2025 (whole-thread analysis; Copilot as ad surface; Ad Voice; 73% higher CTR, 16% stronger conversion for conversational placements; Performance Max beneath AI responses) <https://about.ads.microsoft.com/en/blog/post/august-2025/73-higher-ctrs-why-advertisers-need-to-pay-attention-to-conversational-ai>.
- [20] PPC Hero, “How Microsoft is Transforming Search and Advertising with AI,” June 2025 <https://ppchero.com/how-microsoft-is-transforming-search-and-advertising-with-ai/>.
- [21] Microsoft Family Safety support documentation (no personalized ads to minor accounts in Copilot; non-personalized ads may still appear; no personalized experiences or model-training use for minors) <https://support.microsoft.com/en-us/topic/microsoft-copilot-age-limits-and-parental-controls-f79b47a6-288a-4513-8c01-afe4d16db900>.
- [22] Tinuiti, “Amazon Rufus AI: Optimization Strategies for 2026,” March 30, 2026, <https://tinuiti.com/blog/amazon/amazon-rufus-ai/> (Rufus infers price sensitivity, brand preference, use-case; same query yields different results by shopper).
- [23] EvolveAMZ, “Amazon Rufus Optimization 2026,” April 24, 2026, <https://evolveamz.com/amazon-rufus-optimization-guide/> (multimodal inference incl. OCR per patent/practitioner analysis).
- [24] Amazon, “Amazon’s AI Shopping Assistant Gets Smarter and More Personalized,” November 18, 2025, <https://www.aboutamazon.com/news/retail/amazon-rufus-ai-assistant-personalized-shopping-features> (renamed “Alexa for Shopping”) May 13, 2026; “predictive” personalization from individual activity).
- [25] Seller Labs, “How to Optimize Amazon Listings for Rufus AI,” December 5, 2025, <https://www.sellerlabs.com/blog/amazon-rufus-ai-listing-optimization-2026/> (Rufus surfaces “dominant buyer segment” to sellers).
- [26] Genrise, “Amazon Rufus in 2026: The Enterprise Brand Guide,” 2026, <https://genrise.ai/insights/amazon-rufus-ai-2025> (agentic auto-buy; price monitoring approximately every 30 minutes).
- [27] Retail Technology Innovation Hub, “Rufus and the AI Shopping War,” January 6, 2026 (recommendations skew to Amazon’s own products; auto-buy concerns) <https://retailtechinnovationhub.com/home/2026/1/5/rufus-and-the-ai-shopping-war-why-amazons-assistant-reveals-the-battle-for-customer-intent>.

[28] Research on Alexa advertising data sharing (voice interactions shared with 40+ advertisers; ad-auction bids up to ~30x higher with Echo data): Iqbal et al., “Your Echoes Are Heard: Tracking, Profiling, and Ad Targeting in the Amazon Smart Speaker Ecosystem” (2022) <https://arxiv.org/abs/2204.10920>; see also the companion Appendix A1-Table.

[29] Amazon interest-based-ads policy and opt-out <https://www.amazon.com/adprefs>; academic research finding practice inconsistent with stated privacy policy (Iqbal et al., 2022) <https://arxiv.org/abs/2204.10920>; see also the companion Appendix A1-Table.

[30] Amazon, Alexa+ rollout and Conversational Entertainment Ads <https://advertising.amazon.com/resources/whats-new/conversational-entertainment-ads> (Oct 2025) <https://content-na1.emarketer.com/amazon-eyes-ads-subscriptions-monetize-alexa>; CEO Andy Jassy on ads in multi-turn conversations <https://www.aboutamazon.com/news/company-news/amazon-ceo-andy-jassy-2025-letter-to-shareholders>; see also the companion Appendix A1-Table.

[31] FTC, 2023 COPPA action and \$25M settlement re: Alexa children’s voice recordings <https://www.ftc.gov/news-events/news/press-releases/2023/05/ftc-doj-charge-amazon-violating-childrens-privacy-law-keeping-kids-alexa-voice-recordings-forever>; amended COPPA Rule effective June 2025 (opt-in parental consent for targeted advertising using children’s data) <https://www.federalregister.gov/documents/2025/04/22/2025-05904/childrens-online-privacy-protection-rule>.

[32] Kinzoo, “What to Know about Snapchat AI in 2025” <https://www.kinzoo.com/blog/what-to-know-about-the-new-snapchat-ai>; Snap help center (“data may be used to personalize your experience, including ads”; geolocation collected in Ghost Mode). On AI Sponsored Snaps (advertiser-operated AI agents in the Chat tab, announced April 28, 2026, with Experian as alpha partner; agentic ad suite and brand agents in DMs announced June 18, 2026): Engadget, “Snapchat is Rolling Out Sponsored AI Agents,” April 28, 2026 <https://www.engadget.com/ai/snapchat-is-rolling-out-sponsored-ai-agents-162720124.html>; Snap Newsroom <https://newsroom.snap.com/ai-sponsored-snaps>.

[33] [My AI – Safety and Privacy Hub | Snapchat](#) / privacy disclosures re: My AI personalization.

[34] Vanhoutte et al., “Teens, Tech, and Talk: Adolescents’ Use of and Emotional Reactions to Snapchat’s My AI Chatbot,” PMC12382817, <https://pmc.ncbi.nlm.nih.gov/articles/PMC12382817/> (younger adolescents more likely to use My AI; emotional engagement).

[35] UK ICO, “We Warn Organisations Must Not Ignore Data Protection Risks – Snap My AI,” investigation concluded May 2024 (Snap required to improve risk assessment) <https://ico.org.uk/about-the-ico/media-centre/news-and-blogs/2024/05/ico-warns-organisations-must-not-ignore-data-protection-risks-as-it-concludes-snap-my-ai-chatbot-investigation/>.

[36] Utah Dept. of Commerce, “Utah Sues Snapchat for Unleashing Experimental AI Technology on Young Users,” June 30, 2025 <https://commerce.utah.gov/2025/06/30/utah-sues-snapchat-for-unleashing-experimental-ai-technology-on-young-users-while-misrepresenting-the-safety-of-the-platform/>.

[37] TACD, “Snapchat’s AI Data Grab: Why Teens Are at Risk,” August 22, 2025 <https://tacd.org/snapchats-ai-data-grab-why-teens-are-at-risk-and-regulators-are-silent/>.

[38] Snap Family Center documentation (parental controls over My AI access) <https://www.cnn.com/2024/01/11/tech/snapchat-parents-myai-family-center/index.html>. The Perplexity in-Snapchat integration was amicably ended in Q1 2026 without full rollout; see TechCrunch, “Snap Says Its \$400M Deal with Perplexity ‘Amicably Ended,’” May 6, 2026 <https://techcrunch.com/2026/05/06/snap-says-its-400m-deal-with-perplexity-amicably-ended/>.

[39] xAI/Grok integration into X data ecosystem (March 2025 merger); X ad-targeting; see also the companion Appendix A1-Table. On the April 30, 2026 xAI-powered ad-platform rebuild (Grok handling contextual and semantic targeting, auction, and ranking): Social Media Today, “X Introduces Rebuilt AI-Powered Ad Platform,” April 30, 2026 <https://www.socialmediatoday.com/news/x-introduces-rebuilt-ai-powered-ad-platform/819029/>. On paid placements inside Grok responses: Search Engine Land, “X Will Put Ads Inside Grok AI Answers,” August 2025 <https://searchengineland.com/x-ads-grok-ai-answers-460520>.

[40] TechCrunch, “‘Among the Worst We’ve Seen’: Report Slams xAI’s Grok over Child Safety,” January 27, 2026 (no published documentation about Kids Mode or guardrails; press-inquiry auto-replies) <https://techcrunch.com/2026/01/27/among-the-worst-weve-seen-report-slams-xais-grok-over-child-safety-failures/>.

[41] Common Sense Media risk assessment of Grok (Jan 2026), finding Kids Mode ineffective and “porous” <https://www.common SenseMedia.org/sites/default/files/ai-ratings/csm-ai-risk-assessment-grok-01222026.pdf>; CCDH findings on sexualized image generation <https://counterhate.com/research/grok-floods-x-with-sexualized-images/>; multi-jurisdiction regulatory actions.

[42] The Register, “Anthropic Keeps Claude Ad-Free,” February 4, 2026 (no sponsored links, advertiser influence, or product placements; “should we need to revisit this approach, we’ll be transparent”) <https://www.theregister.com/software/2026/02/04/anthropic-keeps-claude-ad-free/4658051>.

[43] Anthropic statement that an “appreciable portion” of Claude conversations involve sensitive or deeply personal topics <https://www.anthropic.com/news/claude-is-a-space-to-think>.

[44] CNBC, “Anthropic Takes Aim at OpenAI’s Ad Push in Super Bowl Commercial,” February 4, 2026 <https://www.cnbc.com/2026/02/04/anthropic-no-ads-claude-chatbot-openai-chatgpt.html>; Axios, “Anthropic Pledges Ad-Free Claude,” February 4, 2026 <https://www.axios.com/2026/02/04/anthropic-ads-openai-super-bowl>.

[45] Anthropic, “Responsible Use of Anthropic’s Models: Guidelines for Organizations Serving Minors,” <https://support.claude.com/en/articles/9307344>; Claude.ai marketed 18+; free tier expanded Feb 2026.

[46] TikTok recommendation and ad systems (AI-driven; off-TikTok activity used for personalization); AIGC labeling required since July 2024 <https://www.socialmediatoday.com/news/tiktok-implements-restrictions-ads-targeted-teens-ai-disclosure/720667/>.

[47] TikTok, “Disconnect Advertisers” control; commercial-content disclosure requirements (2025) <https://ads.tiktok.com/business/en-US/blog/enhancing-privacy-control-advertisers-users>.

[48] TikTok Ads, “Teen Safety and Wellbeing,” <https://ads.tiktok.com/help/article/tiktok-ads-policy-youth-safety> (personalized ads prohibited for users 13–17 in US, Canada, EEA, UK, Switzerland, Latin America; only broad targeting permitted).

[49] TikTok teen-safety advertising policy documentation <https://ads.tiktok.com/help/article/about-advertising-to-people-under-the-age-of-18>.

[50] FTC, 6(b) inquiry into AI companion chatbots (Sept 2025), issued to OpenAI, Meta/Alphabet, Character.AI, Snap, and xAI, among others, regarding child and teen safety practices <https://www.ftc.gov/news-events/news/press-releases/2025/09/ftc-launches-inquiry-into-ai-chatbots-acting-as-companions>; see Axios, “AI Chatbots Struggle with Subtle Mental Health Cues,” May 2026 <https://www.axios.com/2026/05/12/ai-chatbots-mental-health-cues>, and Stanford Medicine, “Why AI companions and young people can make for a dangerous mix,” by John Sanford, August 2025 <https://med.stanford.edu/news/insights/2025/08/ai-chatbots-kids-teens-artificial-intelligence.html> (cross-cutting source for the emotional-inference and differential-treatment dimensions).

Note: This landscape evolves rapidly. Capabilities and disclosures may change between the date of this analysis (June 2026) and review. Cells marked “Likely” or “Unclear” reflect deliberate caution where public evidence is indirect or absent; they should not be read as confirmed denials. Several findings rest on ad-trade, academic, advocacy, or litigation sources because no first-party disclosure exists.

APPENDIX A2-KEY FINDINGS:

POWERFUL SYSTEMS, INVISIBLE TO THE PEOPLE USING THEM

Summary and interpretation: key findings of Appendix A2-Table, the capability map. Current as of June 2026.

Children and teenagers are using conversational AI systems that can do far more than collect and store their data. The table accompanying this narrative maps nine major platforms – Meta, Google, OpenAI, Microsoft, Amazon (Alexa for Shopping), Snap, xAI, Anthropic, and TikTok – against four capabilities: whether each system can infer characteristics about a young user, sort that user into commercial groups, optimize against the user in real time, and treat the user differently from the next person. For each capability it records two things: what the system can do, and whether an ordinary parent or teenager could find that out from what the company discloses.

A note on how to read this analysis: Its companion, Appendix A1-Key Findings, documents what each platform does with AI chat data, whether it runs ads and whether it uses chat data to target them. This appendix asks the second question: regardless of what data a company holds or links to a name, what is each system able to do to a young user? The distinction matters because most of these capabilities do not require holding identifiable data about a particular child. They run instead on two things privacy law was not built to address. The first is aggregate analysis, inferences drawn not from a child's own record but from patterns learned across millions of other people and then applied to a user the system has never met. The second is memory, the context a system retains within a conversation and across sessions to adapt to a young user without ever assembling what the law would recognize as a profile. Some such inferences are kept as a durable profile, the kind of record privacy law was built to govern; others are made in the moment,

acted on, and discarded. A child can be inferred about, sorted, and steered by a system that holds no identifiable file on it at all. Read together, the two appendices show both what companies do with the data and what their systems are built to do. Every factual claim in this narrative is sourced in the companion Appendix A2 table; this narrative interprets that evidence rather than re-citing it.

The overall picture is this: across the nine platforms surveyed, the capabilities are near-universal and the disclosures are near-absent. Every ad-supported platform infers characteristics about its users, optimizes within the session, and applies some form of differential treatment; most also sort users into commercial cohorts. These are not edge cases or future risks; they are the normal operating mode of commercial conversational AI in 2026. Accessibility runs the opposite way. Almost nothing a young person would want to know is available from what the company tells them. The few capabilities that are plainly disclosed are almost all about one thing, the differential treatment of minors. For everything else, the inference engines, cohort taxonomies, and real-time optimization logic, the company provides no disclosure.

What is disclosed, and whether it can be understood.

Even where something is technically disclosed, disclosure and comprehension are not the same thing. Meta tells users it now uses their AI-chat conversations to target ads, but the statement sits in a policy update with no opt-out, and the actual machinery – the systems that decide what each person is shown – appears only in advertiser trade press. And the

optimization at work here does not require targeting in the conventional sense. Rather than singling out a named individual, these systems can model audiences and optimize delivery toward a desired outcome, shaping what a young person sees without ever matching an identifiable record to a name. Amazon says its shopping assistant, now part of the unified Alexa for Shopping product (Amazon merged Rufus and Alexa+ into one assistant in May 2026), uses “insights from each customer’s individual shopping activity” for “predictive” personalization; that language is accurate and almost meaningless to a shopper, who cannot tell from it that the system has inferred their price sensitivity or placed them in a buyer segment visible to sellers. OpenAI is the most genuinely transparent case, as it emailed users the policy and published a separate explainer on age prediction; yet even there the cohort logic and the live ad auction are described only in general terms. A disclosure a teenager cannot parse is, for practical purposes, not a disclosure.

What is hidden, and where the hiding is worst. The inference engines, the cohort taxonomies, and the real-time optimization logic, the parts that determine what a young person actually sees and how they are persuaded, are overwhelmingly hidden, surfacing only in press reports, academic studies, litigation, or documents written for advertisers rather than users. The contrast is consistent across platforms. Companies tend to describe the regulated and the already-documented (regulated-category ad delivery, minor-account rules, AI-content labels) while saying far less about the inferential machinery, the aggregate and in-session analysis that existing privacy and consumer-protection law was not designed to govern. The opacity is deepest exactly where the stakes are highest. Meta’s and Amazon’s cohort

and real-time systems are powerful and undisclosed. xAI’s Grok publishes nothing about its data practices or its “Kids Mode,” and that silence coincides with its documented child-safety failures. The emotional-state and vulnerability inference at the center of the report’s harms argument is real and documented by researchers, but appears in no platform’s own disclosure at all.

What this means for regulators. A regulator faces a softer version of the same wall. Much of what the map documents was assembled from the Federal Trade Commission’s own 6(b) inquiry, from state attorneys general lawsuits, and from academic and advocacy research, not from anything the companies disclose as a matter of course. That a regulator can eventually compel this information through investigation is not the same as it being transparent. It means the burden falls on enforcement after the fact, rather than on disclosure before harm. And it confirms why a regulatory approach anchored to whether a company has been transparent about holding identifying data misses the point: the operative fact is what the system can do to a young user, which is knowable to neither the user nor, without subpoena, the regulator.

Bottom line. Across the conversational-AI platforms children and teenagers actually use, the systems can infer who you are, sort you, optimize against you in real time, and treat you differently from the next user, and almost none of that is anything a young person, or their parent, can learn from what the company itself discloses. The capabilities are the norm; the transparency is the exception. That gap between what these systems do and what anyone outside the company can see is the strongest argument in this report for regulating capability rather than disclosure.