William Blair

AI INSIGHTS

Equity Research
Technology, Media, and Communications | Internet and Digital Media

January 22, 2026
Industry Report

Ralph Schackart +1 312 364 8753
rschackart@williamblair.com

Jason Ader +1 617 235 7519
jader@williamblair.com

Arjun Bhatia +1 312 364 5696
abhatia@williamblair.com

Jack Brenczewski +1 312 364 5321
jbrenczewski@williamblair.com

Justin Whitney +1 312 364 5464
jwhitney@williamblair.com

Beyond the Build

The Next Wave of Generative AI Applications and Investment

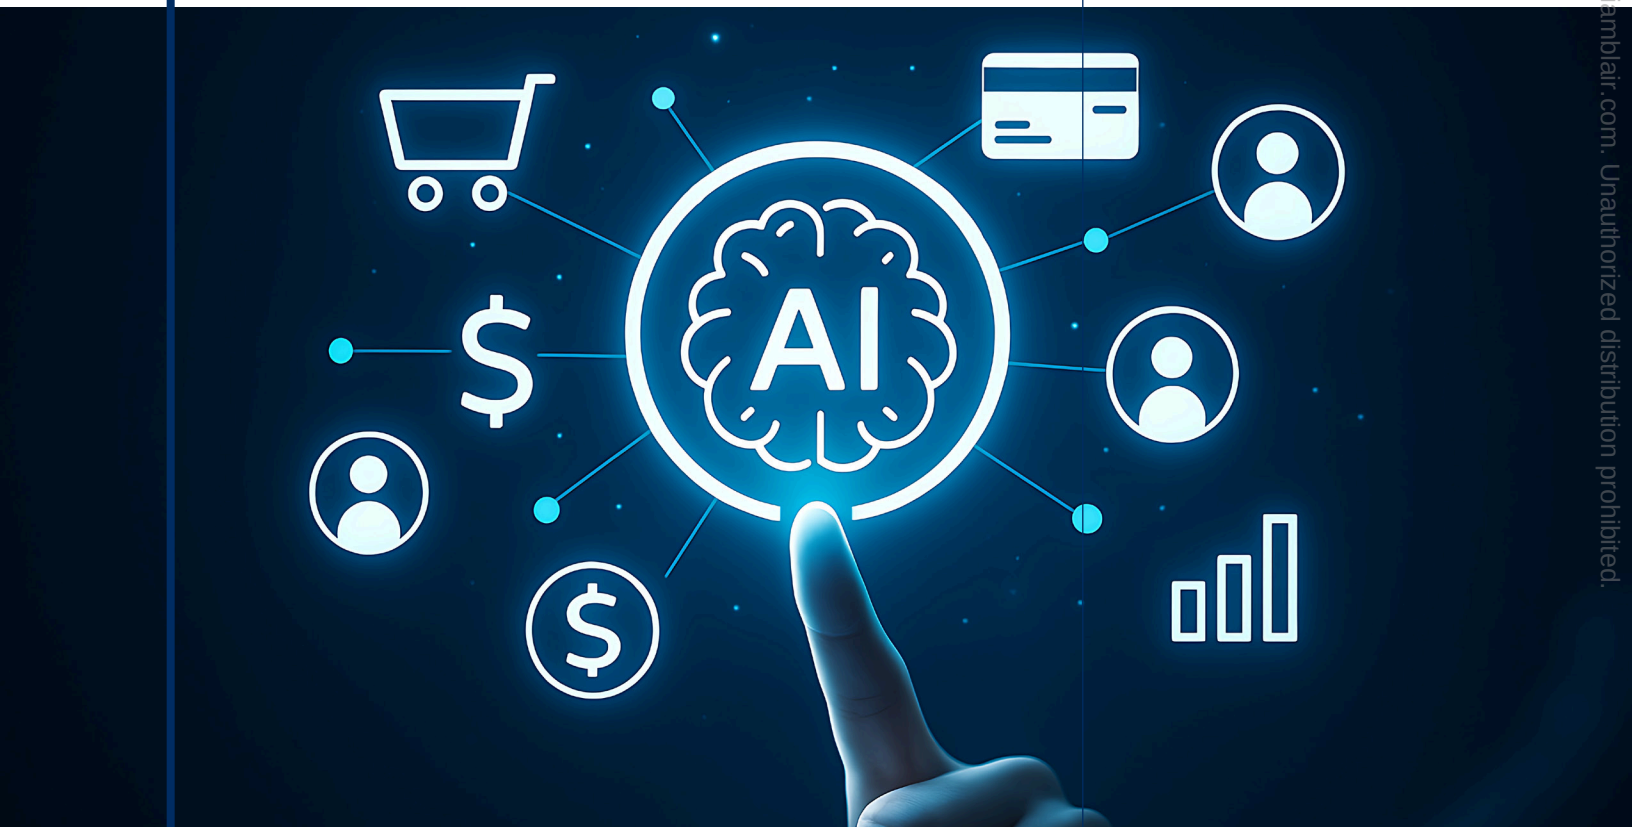

Contents

Overview

Artificial intelligence has dominated the investing zeitgeist over the past few years. While model developers and infrastructure providers have captured most of the headlines and initial investments, a growing number of GenAI use-cases are emerging and will grow in importance. While the infrastructure layer needs to be built first, as the pace of the infrastructure buildout slows, investor attention is likely to shift toward tangible applications and innovative new companies that drive the technology forward, in our view.

In this report, we dive into five key topics encompassing the current state of AI, what the next three to five years may entail, and industries facing major disruption. We preface this report by noting that the non-infrastructure GenAI landscape is evolving rapidly, often daily. This report is our effort to capture that momentum, reflect on recent developments, and highlight areas for investors to focus on in the future.

Below we highlight the key takeaways from the five primary sections in our report.

Key Takeaways

1. **AI integration is rapidly advancing as a result of greater accessibility, fueling industry growth.** Cost reductions, open-source models, and innovations like Mixture-of-Experts have boosted AI use in both existing applications and new start-ups. At the end of this report, we spotlight roughly 200 companies that are now active in this space.

2. **AI-native companies are rewriting the revenue playbook, scaling faster, and some more profitably, than in prior tech cycles.** Consumer-focused companies are leveraging usage-based pricing models and upselling into businesses with specific business-to-business (B2B) pricing plans. Vertical-specific consumer companies are employing creative revenue models to seek differentiation in a competitive landscape.

3. **There will soon be more AI agents than the U.S. population, and they are moving beyond simple tasks and will transform consumer experiences.** Agentic AI is evolving beyond simple task execution toward autonomous, personalized workflows that could transform both enterprise and consumer experiences over the next three to five years. We highlight the potential for AI agents to act independently, collaborate, and integrate into edge devices like wearables and home products, enabling continuous, context-aware support.

4. **Voice will become the main way consumers interact with AI and will unlock more AI use-cases.** In our view, voice is emerging as the primary interface for consumer AI, driven by technical breakthroughs in speech recognition, latency reduction, and expressive text-to-speech. We also highlight current use-cases across therapy, language learning, and recruiting.

5. **AI has permanently disrupted the traditional search and digital advertising landscape, although Google is likely to continue to be a winner in the space.** GenAI is reshaping the search landscape, from the rise of AI Overviews and chat-based queries to the emergence of AI browsers and zero-click behavior. We explore the implications for publishers, advertisers, and brands, including the shift from search engine optimization to generative engine optimization, and highlight new tools and strategies that companies are adopting to maintain visibility and performance in an AI-driven search environment.

Exhibit 1
Generative AI Report
Key Predictions and Trends to Watch for Investors

Existing and New Applications				
Incumbents will drive widespread AI adoption by embedding models into existing platforms, creating sticky user experiences and operational efficiencies.	Purpose-built AI start-ups will accelerate and may attract outsized venture capital funding as they expand into dedicated consumer use-cases, potentially leading to increased M&A activity.	MoE architecture will become the industry standard, reshaping model economics and enabling smaller players to potentially compete with hyperscalers.	Open-source models will accelerate innovation and competition, especially in cost-sensitive geographies and verticals.	Companies that successfully combine performance with cost efficiency will gain rapid market traction and developer loyalty.
Revenue Models				
Consumer AI companies will achieve net revenue retention figures previously only achieved by B2B software companies.	Subscription-based consumer AI companies will scale faster and more profitably than Web 2.0 companies.	Vertical-specific monetization strategies (e.g., affiliate, marketplace, licensing) will gain traction as companies tailor models to niche use-cases.	Hybrid subscription models with usage-based pricing could become the norm, driving strong net revenue retention and margin expansion for AI companies and increasing scale.	Personalization will become a monetization tool, increasing switching costs and customer lifetime value.
Agentic AI				
The importance of apps and brand affinity may wane with AI emerging as the primary interface.	Every consumer will likely have a personal assistant/agent.	Agents may begin to disrupt the existing online marketplace ecosystem.	Meta glasses adoption can be a major catalyst for consumer AI adoption if the price drops below $199/unit.	Edge devices and in-home AI hardware will be the first step in achieving broad adoption of consumer AI.
AI-Powered Voice				
Voice may become the dominant modality for consumer AI once general purpose voice models achieve sub-200ms latency, which could happen within two to three years.	Verticals like therapy, education, and recruiting will see early adoption of voice-native AI agents with measurable ROI.	Edge computing will enable real-time voice interactions, pushing demand for low-latency, privacy-first hardware solutions.	Companies that combine expressive voice with agentic capabilities will differentiate in a crowded consumer marketplace.	Monetization of AI-powered voice will lag initially but accelerate as trust, usability, and vertical specialization improve.
GenAI Search				
GEO will replace SEO as the dominant strategy for digital visibility, creating demand for new marketing platforms and analytics tools tailored to AI-driven search.	Zero-click behavior will erode publisher traffic, accelerating the need for licensing deals and alternative monetization models.	Companies will increase digital advertising spend to replace degraded organic traffic, representing a tailwind to Google, Meta, and AppLovin, among others.	Google will shift toward a marketplace-based revenue model, monetizing transactions rather than clicks, which could reshape digital ad bidding strategies and margins.	Gemini will replace traditional Google search within three to five years.

Source: William Blair Equity Research

AI in Existing and Purpose-Built Applications

The accessibility of AI has increased dramatically since the release of ChatGPT, lowering barriers to entry for both incumbents and start-ups. As costs fall and open-source models gain traction, there are two emerging trends: existing companies are integrating AI into legacy platforms, and AI-native start-ups are creating purpose-built applications.

Accessibility to AI Has Increased Over the Past Few Years and Is Likely to Increase Further
Since the release of ChatGPT, the rapid pace of innovation in the AI space has increased accessibility to the technology, allowing more companies to compete in the evolving landscape. This is a result of increasing model efficiencies, competition, the emergence of open-source models, and broad enterprise and consumer adoption reducing token prices. This "democratization" of the technology drives costs down and allows existing companies to invest in AI using existing revenue streams while lowering barriers to entry for start-ups wanting to build. In our view, existing companies adopting and reinventing with AI will drive continued adoption by both enterprises and consumers, while AI-native companies building purpose-built applications will unlock new possibilities and use-cases.

The emergence of ***DeepSeek*** in late 2024 into 2025 created a shift in the economics of the industry. At the time, DeepSeek claimed that it built its V3 model, released in December 2024, for about $6 million (a widely disputed, but not disproven, figure) using less advanced H800 chips developed by Nvidia. This called into question the massive capital expenditures by hyperscalers and marked a pivotal moment for AI, in our view, as it sharpened focus on cost structures of foundational models. DeepSeek was able to achieve this cost efficiency primarily through employing a Mixture-of-Experts (MoE) approach. MoE is an AI model architecture in which only subnetworks of parameters (the "experts") are activated and trained for each query, chosen by a trained gate network.

Although not confirmed by ***OpenAI,*** it is widely speculated that GPT-5 employs this MoE approach (OpenAI mentions a "routing-first" design that resembles an expert gate network). This would explain the drastic decrease in inference and token cost relative to model intelligence. While the exact parameter count of GPT-5 is unconfirmed by OpenAI, experts estimate that the model employs more parameters than GPT-4, with some sources indicating magnitudes in the trillions. However, through use of the MoE approach, OpenAI can leverage the cost structure of running its model to experience some of the cost efficiencies enjoyed by DeepSeek. This has provided a tailwind for companies leveraging GPT-5 in products. For example, Duolingo calls on OpenAI models to build its GenAI products, according to management. Since the release of GPT-5, it has seen costs decrease, giving management flexibility in product pricing and capability without materially affecting margins.

While the cost efficiencies displayed in the release of DeepSeek's models have invited a sharpened focus on developing more-efficient models, another feature of the models is the open-source nature. Open-source models are those whose parameters, architecture, and code are released publicly and available for download and modification, typically with guidelines and rules for redistribution and commercialization. Open-source models increase accessibility to AI, in our view, since they lower barriers to entry for developers wishing to build on existing models, accelerating the pace of competition and innovation. However, DeepSeek was not the first company to open source its models. Another notable company that touts open-source AI is Meta, which positions itself as the domestic leader in open-source models. Meta's Llama models are widely used in research and experimentation, representing an indirect monetization lever as Meta creates a competitive advantage by allowing developers and researchers to build with its models, reducing the likelihood that those same users choose to build their own model from scratch instead. Further, Meta can provide managed services and enterprise support for developers and researchers building with its models. In any respect, it is not surprising to see Meta choose to open source its models given its traditional monetization strategy of prioritizing scale. By making open-source models, Meta can grow its user base and attract developers and researchers away from its competitors, deepening its strategic moat in the industry, in our view.

The increasing efficiencies in AI models, sparked by the emergence of DeepSeek and the embrace of open-source models by leading AI companies, create pressure to efficiently accelerate AI research, which increases the pace of innovation for the industry overall.

Cost Efficiencies Lead to Increased Adoption Within Existing Applications

The increasing accessibility to AI has led to more integration into existing applications as incumbents are able to inexpensively integrate AI into previously installed bases. According to information from FactSet, 306 companies in the S&P 500 index mentioned AI during second quarter 2025 earnings calls, above the 5-year average of 136 and the 10-year average of 86.

Exhibit 2
Generative AI Report
Number of S&P 500 Companies Citing "AI" on Earnings Call

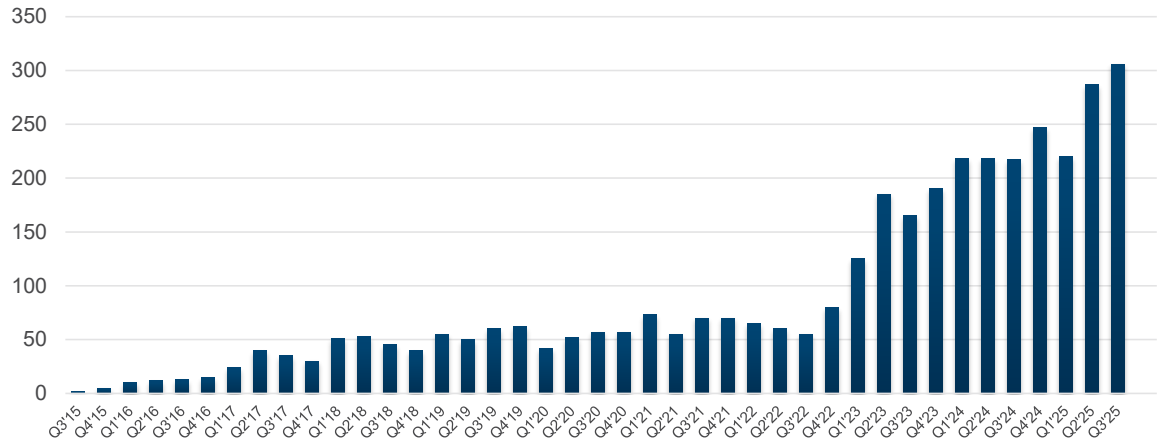

Sources: FactSet and William Blair Equity Research

The sharp increase in mentions of AI began in the fourth quarter of 2022, as seen in exhibit 2, which coincides with the release of ChatGPT. Since then there has been a gradual increase in the number of companies citing AI on earnings calls. The range of use-cases is far-reaching. Some companies are building out new AI tools to complement legacy services, some are integrating AI into existing services to increase efficiency, and some are using the technology internally within research-and-development or customer service departments to lower costs. The investments made by these companies are largely funded out of existing revenue, which mitigates risks, although further proof points regarding the efficacy of this technology are still needed. Exhibit 3 includes a broad overview of AI integration across our team's coverage list for public equity investors.

Exhibit 3
Generative AI Report
Level of AI Integration Across Internet and Digital Media
Coverage List

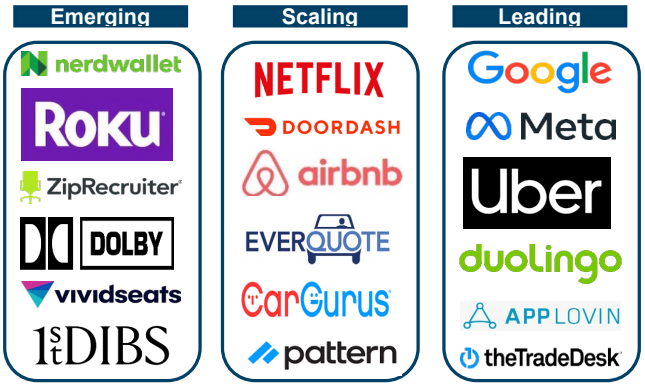

Source: William Blair Equity Research

Despite the different stages of integration, it is notable that every company in our coverage has implemented AI into its business in some form. Overall, investments in AI systems by established companies underscore the importance of this technology, in our view. Positive results in terms of revenue growth, new product capabilities, and cost efficiencies further incentivize investment and innovation, creating a flywheel effect. However, while the increased adoption among existing companies is encouraging for the industry overall, the emerging and transformative technologies that will push the industry forward are likely to be born out of start-ups creating purpose-built applications.

Increased Accessibility Has Led to a Boom in Purpose-Built AI Applications

Less expensive and more accessible AI technology also lowers barriers to entry for start-ups, enabling further experimentation and research. As a result, start-ups using AI can create technological advancements, more funding opportunities, and the possibility of outsized returns. Data sourced from PitchBook suggests that AI/ML has become the most popular category among private investors and has grown roughly 75% year-over-year, as evidenced in exhibit 4 on the following page.

New use-cases for AI spring up almost every day across a wide range of industries. Below, we spotlight four AI-native, consumer-focused start-ups taking advantage of the increasing accessibility to AI and building applications that have generated significant interest and funding.

Exhibit 4
Generative AI Report
Top 10 Verticals - Highest Total Funding in 2025

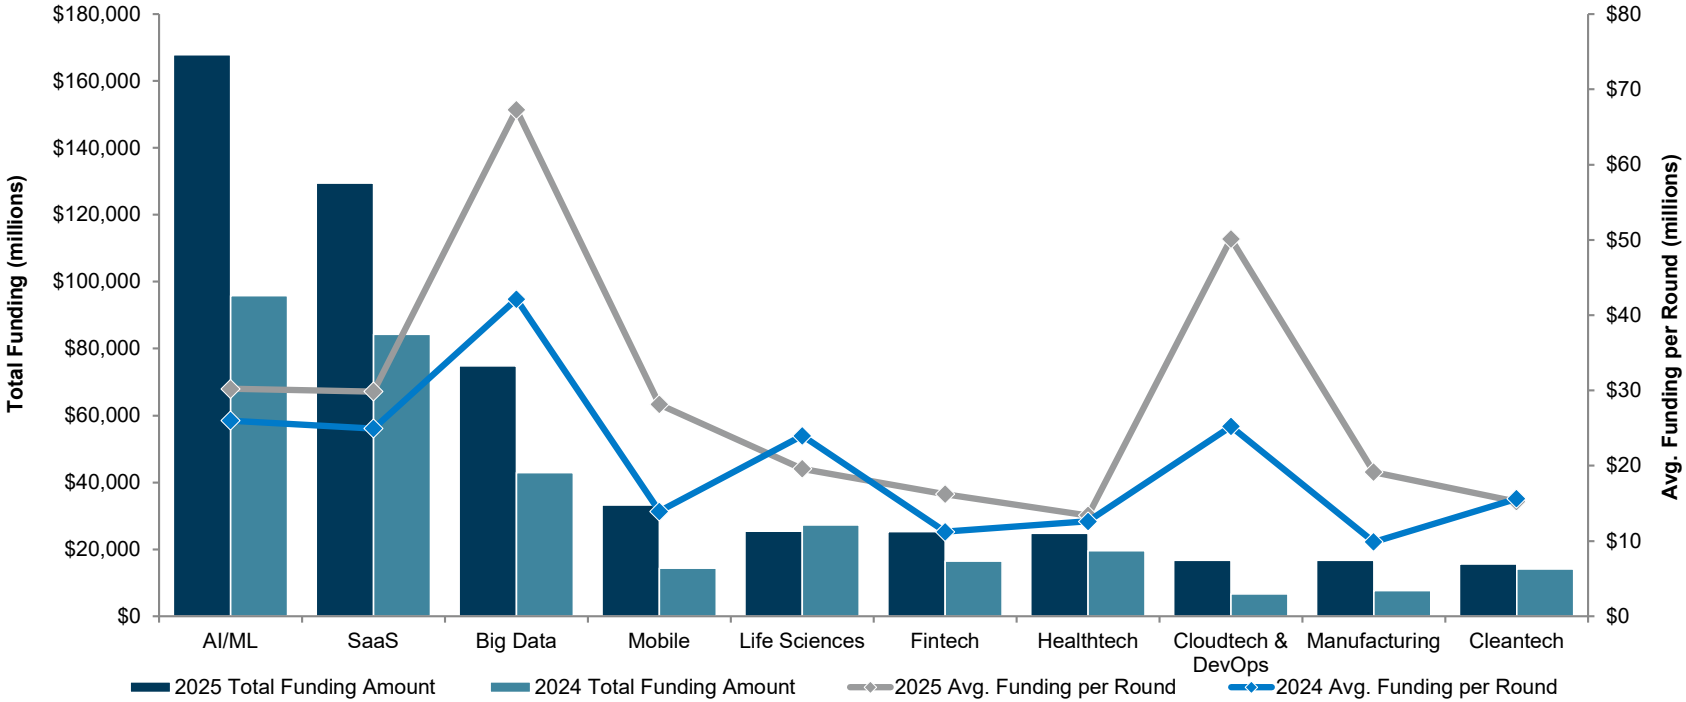

Source: PitchBook Data, Inc., and William Blair Equity Research

- ***Perplexity.*** An AI-powered search and answer engine that provides conversational answers to queries through a chatbot or its browser product, Comet. Perplexity differentiates itself by answering queries with real-time web results as opposed to training data. Comet acts as a more holistic personal assistant, integrating into every workflow a user might perform on a browser and adding capabilities beyond a simple answer engine. Comet can book meetings, follow up on emails, shop on a user's behalf, etc. The company is based in San Francisco and raised $200 million in September 2025 at a $20 billion valuation. Its products are offered through a freemium model with additional premium tiers.

- ***Runway.*** Develops AI-powered image and video tools that allow users to edit and transform existing videos and create characters, locations, and objects using natural language prompts. Creators and producers of the recent Amazon Prime Video show *House of David* used Runway's AI tools in production for previsualization and increasing resolution quality, creating both time and cost efficiencies. Runway is based in New York City and raised over $300 million in April 2025 at a $3 billion valuation.

- ***HeyGen.*** Specializes in AI-driven video generation focused on realistic spokesperson videos, avatars, and multilingual translations for corporate trainings, product marketing, and language learning. Trivago uses HeyGen to produce localized advertisements all over the world, a process that historically would have taken months of production and delivery time. Trivago was able to create 30 localized ads in roughly half the production time it would have previously taken. The company, based in Los Angeles, raised a $60 million series A in March 2024 at a $500 million valuation.

- ***Character.ai.*** A companion app that allows users to create customizable AI companions and characters with varying personalities, values, and conversation styles. Consumers mainly use the app for entertainment purposes, but it can also be used for life coaching, therapy, etc., by creating hyper-personalized companions. Character.ai is an end-to-end application, meaning it owns its models and the end-customer relationship. The company is based in Menlo Park, California, and raised $150 million in March 2023, valuing it at about $1 billion.

Despite the interesting use-cases and early adoption beyond enterprise, venture capital funding for consumer-focused AI companies has lagged enterprise-focused AI companies since the broader VC downturn in 2022. In general, high interest-rate environments increase scrutiny on venture capital investments as purse strings tighten. In this environment, enterprise AI is typically viewed as a safer investment than consumer AI. Enterprise AI start-ups are geared to solve specific pain points faced by businesses that have a clear pathway to generating revenue growth or cost savings, whereas consumer AI start-ups tend to focus on more emerging technologies. Further, enterprise models tend to have long-term contracts, creating stable and sticky recurring revenue, a highly desirable characteristic for a VC investor. Meanwhile, consumer start-ups can lag in monetization compared to enterprise. However, as we detail in the next section, we are seeing signs that this is changing in the current cycle.

Overall, greater accessibility to AI has accelerated its adoption by existing companies and led to the creation of more emerging technologies. We are still in the early stages of this emerging cycle, in our view. As AI continues to create cost efficiencies and value-added technologies, revenue generation will follow, in our view.

Key Predictions and Trends to Watch for Investors

1. Incumbents will drive widespread AI adoption by embedding models into existing platforms, creating sticky user experiences and operational efficiencies.

2. Purpose-built AI start-ups will accelerate and may attract outsized venture capital funding as they expand into dedicated consumer use-cases, potentially leading to increased M&A activity.

3. MoE architecture will become the industry standard, reshaping model economics and enabling smaller players to potentially compete with hyperscalers.

4. Open-source models will accelerate innovation and competition, especially in cost-sensitive geographies and verticals.

5. Model developers that successfully combine performance with cost efficiency will gain rapid market traction and developer loyalty.

Revenue Models

Introduction
The rapid pace of adoption and innovation is exciting, but financial returns are essential for companies to sustain innovation and continue pushing the frontier of possibilities with AI. There is evidence that AI-native companies monetize faster than previous technological cycles, driven by consumer familiarity with subscription models and the need to charge for higher operating costs.

Many AI-Native Companies Are Scaling Revenues Faster Than Previous Technology Cycles
In previous non-AI technology cycles, most consumer start-up companies followed a similar playbook: build a user base, generate scale, and monetize via advertising, (followed by transactional or subscription revenue), often waiting well over one year to generate revenue. Now, AI start-ups are generating revenue almost instantly, with standout companies generating millions of annual recurring revenue (ARR) within one year. According to data published in June 2025 from companies that venture capital firm a16z has invested in or interacted with, the median ARR of consumer GenAI start-ups in year one is over $4 million. Top-quartile performers average $8.7 million in ARR, while bottom-quartile companies average $2.9 million. What was previously considered best-in-class for ARR generation now reflects bottom-quartile performance. Top performers in the GenAI era are reporting **tens or even hundreds of millions of ARR within the first year.**

Some examples of these companies are ***Suno, Lovable,*** and ***Midjourney.***

- Suno, an AI music generator, reportedly quadrupled ARR in one year to roughly $150 million and raised a $250 million series C round of funding at a $2.45 billion valuation in November 2025 (despite legal headwinds).

- Lovable, a vibe coding platform that users can use to build functional apps and websites with natural language prompts, achieved $50 million of ARR within six months of operation and $200 million within one year of launch.

- Midjourney, a GenAI platform that turns text prompts into images, reached $500 million in revenue in 2025, up from $300 million in 2024, despite being bootstrapped since its founding in 2022.

Exhibit 5
Generative AI Report
Start-ups Scaling Fast

Company	Brief Description	ARR (Most Recent Reported)
SUNO	AI Music Generator	$150 million
Lovable	Vibe Coding Platform	$200 million
Midjourney	AI Image Generator	$500 million

Sources: Company websites and William Blair Equity Research

In our view, these figures underscore the investment opportunity and, assuming these companies can grow profitably, pave the way for potential IPOs for public equity investors in the future. In contrast to the dotcom bubble, an era often compared to this current AI cycle, companies are generating scaled financial returns rather than simply demonstrating user growth. This is likely to make it easier for investors to underwrite investments as they see a clear path to achieving returns, assuming this scaled growth continues.

This shift to early revenue generation is caused by two factors, in our view.

- First, AI-native start-ups no longer have the luxury of delaying revenue for multiple years to build up a user base since the costs of operating an AI-native company are higher than traditional software companies. Every time a user interacts with an AI-native company's product, the company pays incremental costs to run the models behind the scenes. This differs from software-as-a-service (SaaS) companies, which can deliver services to new users at almost no incremental cost once the software is built. As a result, AI-native companies have higher cost of goods sold than traditional consumer-oriented or software companies, underscoring the importance of monetizing quickly.

- Second, consumers and small and midsize businesses (SMBs) today are "trained" to pay subscriptions, which differs from previous technology cycles. Since the advent of SaaS, consumers and enterprises have become more comfortable paying a recurring subscription for products. Consumers today might subscribe to several different services for video/music streaming, ridesharing, gaming, etc., so the idea that they might subscribe to their favorite AI chat agent for access to better models or increased functionality is not unreasonable, in our view.

Many consumer AI companies are now charging roughly $20 per month for base-level subscriptions, which is about double pre-AI subscription models. Still, consumers are paying; a16z estimates that OpenAI and ***Anthropic*** alone are adding almost half as much new revenue in 2025 as the entirety of the public SaaS universe, excluding the Mag 7. In our proprietary consumer AI survey, about 85% of survey respondents indicated they are paying for an AI subscription. We attribute this to consumers being more comfortable with subscriptions while recognizing the value that these companies offer.

Exhibit 6
Generative AI Report
Question: Do you currently pay a subscription for any AI tools?

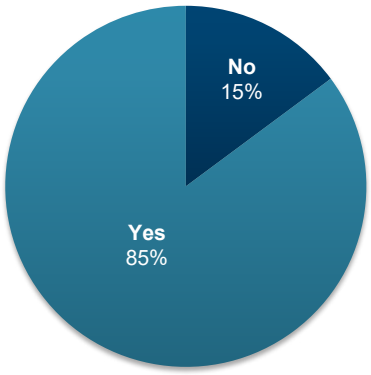

n = 155
Source: William Blair Equity Research

This shift toward early revenue generation in an increasingly competitive industry underscores the importance of employing the right revenue model, in our view. As competition for revenue and funding grows, companies are increasingly turning to creative revenue models as a means of differentiation. Next, we dive into different revenue models employed by AI-native companies.

Subscription-Based Revenue Models Still Dominate, but With a Twist

The most common revenue model employed by AI-native companies is the subscription model, popularized by SaaS companies in the early 2000s, along with other companies such as Netflix and Spotify, among others. For example, OpenAI offers plus, team, and enterprise plans that generally have more features and access to newer, better-performing models than the free plan. However, another growth driver being employed by consumer AI start-ups is usage-based pricing, where consumers can pay for additional capacity if they exceed the allotment of their subscription plan. This is a practical model for AI-native companies because of the incremental usage-based inference costs.

Midjourney, the image and video generator spotlighted previously, is one of many AI-native companies that employ this model. Midjourney offers four plans—basic, standard, pro, and mega—varying in monthly price from $10 to $120. Each plan offers increasing amounts of fast GPU time (defined as approximately one minute per image and eight minutes per video) per month. Beyond this, each plan also offers the ability to purchase extra GPU time at $4 per hour, representing the usage-based component of the subscription plan. This means that even the highest-paying subscribers can still spend on extra compute if they run out. Therefore, Midjourney possesses the ability to increase revenues of even its highest-paying customer cohort, an option that was not available for traditional subscription models (i.e., the highest Netflix subscription tier does not limit the amount of time a user can spend watching movies). Thus, according to a16z, consumer AI companies using hybrid subscription models with a usage-based component are **achieving net revenue retention over 100%,** a high-water mark that was previously achieved only by B2B software companies.

Blurring the Lines Between Business-to-Consumer (B2C) and B2B

For Web 1.0/2.0 companies, there was a clear dividing line between businesses focused on the consumer and businesses focused on SMBs and enterprises. Although it is early in this AI cycle, there is evidence of companies adopting hybrid models with some success, focusing first on consumer adoption before expanding into SMBs. Typically, consumers tend to be more price sensitive than businesses, and employees are being rewarded for being AI-native and bringing applications to their employers. Therefore, companies are finding success with capturing consumers, with either inexpensive or free plans, and then upselling into businesses, for significantly higher prices. Three companies that are taking this approach are ***Anysphere, Gamma,*** and ***ElevenLabs.***

- Anysphere, better known by its product name, Cursor, is a consumer-focused AI coding assistant that also has a B2B business. Cursor initially built a user base by focusing on individual developers with its hobby and pro pricing plans. Its hobby plan provides individuals with a one-week free trial, and its pro plan is priced at $20 per month. However, Cursor also offers two business plans, a teams plan and a more comprehensive enterprise plan. Its business plans incorporate added features such as security and analytics. The company's strategy of focusing on the product and making it accessible to consumers encouraged adoption and created customer loyalty that it was able to parlay into B2B adoption. Cursor reportedly achieved $1 billion in ARR as of November 2025, three years after its founding.

- Gamma, a popular AI design tool for presentations and websites, has both consumer- and B2B-focused plans. The company offers a free plan that does not require a credit card, limiting friction and encouraging consumers to try the platform. The free plan comes with limited AI creations and 400 credits. Gamma's next plan starts at $8 per seat per month, comes with unlimited AI creations, and has 1,000 monthly credits. This plan removes the Gamma watermark from presentations, which is an important requirement for corporate use, allowing for custom branding. The company also offers pro and ultra plans, which increase the monthly credit allotment and unlock other features, such as analytics and API access. According to the company's website, Gamma has $100 million of ARR and 70 million users, with a reported valuation of $2.1 billion.

- Lastly, ElevenLabs, a start-up that specializes in AI-powered voice models, took a similar approach, building a strong consumer user base with free and starter tiers, before quickly building B2B capabilities. For example, it added HIPAA compliance to its voice and conversational agents to position itself to serve healthcare companies. The company offers three different business plans tailored to businesses of different sizes. In October 2025, ElevenLabs officially announced a $100 million employee tender offer at a $6.6 billion valuation. In the blog post announcing the offer, cofounder Mati Staniszewski said the company had recently surpassed $200 million in ARR, was expecting to surpass $300 million by the end of the year, and was approaching a 50/50 revenue split between its enterprise and self-service customers, with enterprise having grown more than 200% in the previous year.

Exhibit 7
Generative AI Report
Usage-Based Subscription Pricing for B2C and B2B Plans

Company	Brief Description	B2C Pricing	B2B Pricing
CURSOR	AI Coding Assistant	-Hobby: Free -Pro: $20/month -Pro+: $60/month -Ultra: $200/month	-Teams: $40/month -Enterprise: Custom
Gamma	AI Design for Presentations and Websites	-Free: Free -Plus: $10/month -Pro: $25/month -Ultra: $100/month	-Plus: $10/month -Pro: $25/month -Ultra: $100/month
IIElevenLabs	AI-powered Voice Models	-Free: Free -Starter: $5/month -Creator: $11/month -Pro: $99/month	-Scale: $330/month -Business: $1,320/month -Enterprise: Custom

Sources: Company Websites and William Blair Equity Research

Overall, this combination of B2C and B2B subscription plans from companies that started with a consumer focus is an interesting shift from previous cycles where there was more separation. The result is faster monetization, stickier revenue, and increased market sizes. Further, these companies are raising money fast at growing valuations, meaning the investment community recognizes the benefits of this multidimensional revenue. We believe that businesses that generate loyal consumer adoption before offering B2B-focused plans will likely grow faster and have more stable revenue models.

Alternatives to Subscription Models Provide Flexibility for Niche Consumer Companies
Niche consumer AI-native companies that do not charge subscriptions are deploying varying and creative revenue models to differentiate from an increasingly competitive industry. These include:

- ***Advertising revenue*.** This revenue model, which has been a popular choice among consumer internet platforms for years, is proving increasingly successful in an AI world. ***Koah,*** which closed a $5 million seed round in September 2025, is focused on helping AI companies monetize with ads. Many AI companies cannot charge the $20 per month subscription that ChatGPT can charge, so they are searching for alternative monetization strategies. Koah serves ads into apps like ***Ollie*,** an AI assistant focused on planning meals and grocery lists, that appear at relevant moments in chats and are marked as sponsored content, not unlike a Google search. For example, a user on Ollie might be served an ad for an Amazon Fresh item on the user's grocery list. As more specialized use-cases of AI emerge, advertising will become an increasingly popular monetization model, in our view.

- ***Affiliate revenue*.** Affiliate revenue models are essentially referral programs, where brands pay companies a portion of the purchase price of a product for directing consumers to the purchase. This model has become popular among AI stylist companies. ***OneOff*** is an AI stylist that users can interact with in natural language. For example, a user might prompt the AI by saying, "I want to dress like Justin Bieber," and the platform directs the user to various brands and pieces that Justin Bieber has worn or currently wears. If the user was to purchase one of the products, OneOff would receive a referral fee. Another AI stylist, ***Alta,*** uses a similar revenue model. Upon signing up for Alta, the user uploads pictures of his or her current closet and details personal preferences and style. Alta then acts as a daily, personalized stylist that can recommend outfits for any event. Alta also offers try-on features for clothing the user

is considering buying. Like OneOff, Alta receives a referral fee if the user purchases an item through the app. This model makes sense for specialized use-cases where users have high purchasing intent, in our view.

- ***Marketplace models.*** Marketplace models connect buyers and sellers on a shared platform, typically taking a cut of each transaction, and are prevalent with existing internet companies. ***Replicate*** is a platform where developers publish AI models and users pay to use them on the cloud, with Replicate taking a percentage of the revenue. Potential purchasers can easily find specialized models for their use-case through the platform. For example, a user can search for models that create videos from images. Replicate then returns all the models that developers have posted on the marketplace and the number of times the model has been run. Once a model is selected, the user pays per use, with Replicate taking a portion of the revenue for facilitating the transaction.

- ***Personalization.*** Another contributor to revenue for AI companies, while not an explicit revenue model, is personalization, in our view. The higher the degree of personalization that agents have, the higher the switching costs for users. For example, the more users engage with ChatGPT and reveal information about themselves, the more tailored and personalized the responses become. Eventually, the agent starts to resemble a trusted resource who understands users' taste, preferences, goals, etc. Once this relationship is established, it can feel daunting to switch to another agent, even if that agent's model is "better" performing. Therefore, increased personalization becomes a way to generate sticky revenue. While it is not an explicit revenue-generating model, it can be a growing factor, in our view.

Exhibit 8
Generative AI Report
Nonsubscription Revenue Models

Model	Description	Company
Advertising Revenue	Placing advertisements in chatbot answers or other outputs targeted toward the user	Koah
Affiliate Revenue	Referral programs where brands pay AI companies a portion of the purchase price of a product for directing consumers to the purchase	Alta, OneOff
Marketplace Models	Connect buyers and sellers on a shared platform, taking a cut of each transaction	Replicate
Personalization	The higher the degree of personalization, the higher the switching costs for users, representing an indirect revenue stream	ChatGPT

Source: William Blair Equity Research

Overall, while many AI-native companies employ subscription-based pricing models that users have become trained to pay over the years, differentiation of revenue models has increased with industry competition. The rapid pace of change in the AI industry will make it interesting to uncover which revenue models will accompany the industry's future leaders.

Key Predictions and Trends to Watch for Investors

1. Consumer AI companies will achieve net revenue retention figures previously only achieved by B2B software companies.

2. Subscription-based consumer AI companies will scale faster and more profitably than Web 2.0 companies.

3. Vertical-specific monetization strategies (e.g., affiliate, marketplace, licensing) will gain traction as companies tailor models to niche use-cases.

4. Hybrid subscription models with usage-based pricing could become the norm, driving strong net revenue retention and margin expansion for AI companies and increasing scale.

5. Personalization will become a monetization tool, increasing switching costs and customer lifetime value.

So far, we have covered the technology driving the industry and how companies are monetizing. Next, we look at where the technology is headed: agentic AI that can plan and act independently, and voice interfaces that make these systems more natural and accessible. This shift moves GenAI from a case-by-case tool to an integrated personal assistant.

Agentic AI Moves Beyond Simple Tasks

Introduction

At the current rate of development, we believe the next three to five years will bring about a form of agentic AI that has the potential to revolutionize the way humans interact with digital marketplaces and online ordering. Agentic AI encompasses autonomous systems, embedded deep within operational workflows, that are equipped with contextual awareness of the prompt. In this section, we seek to explore what a world with unique digital agents can look like in the next three to five years, compared with where development is currently, and how edge devices like wearables and home products will look to familiarize consumers with AI technology.

What Will This Look Like?

AI agents have the potential to be the most monetizable product of artificial intelligence in the near future, in our view. Consumers (and businesses) have found several use-cases for AI to enhance productivity, planning, and decision-making with the existing tools available. However, these systems are still heavily reliant on user inputs and often struggle to interpret directions and adapt in unfamiliar circumstances. The next steps for agentic AI will revolve around acting autonomously. We believe these systems will be integrated across several enterprise domains, hallucinate less, and eventually collaborate with others just like a human coworker would. Traditional versions of AI (e.g., Siri and Alexa products) rely on predefined workflows programmed by humans, such as accessing Spotify and playing the song a user verbally requests or mapping a route based on a given address. These workflows become agentic when the human decision-maker is replaced by an LLM that can reason independently. Therefore, instead of a human defining where an LLM should source information, which tools to use, and in what order, an AI agent would automatically iterate on its workflow to produce the best possible response for a given prompt.

Exhibit 9
Generative AI Report
Agentic Workflow Diagram

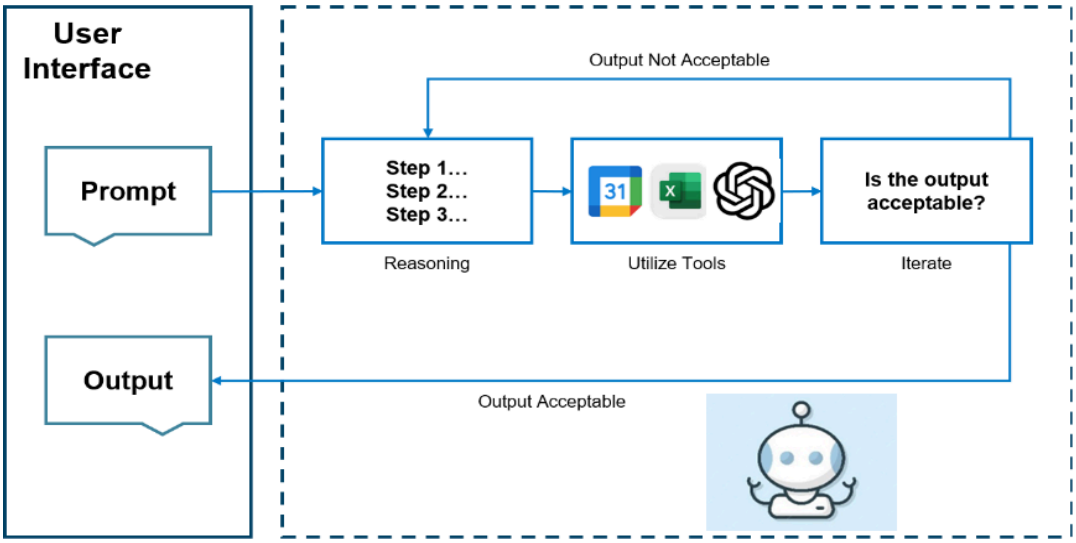

Source: William Blair Equity Research

Agentic AI is meant to be self-reinforcing and generalize based off the prompts of a user rather than reciting from a list of outputs. For example, if a user had an agentic workflow that compiles relevant news based off a provided list of sources each morning, the agent might iterate on its workflow by suggesting other news sources the user frequents or reorganizing the user's morning tear sheet based on what flows best. Today, investors are more concerned with how agentic AI can transform business or enterprise responsibilities (managing ad campaigns, customer support, creating reports, etc.), but we believe that in the hands of the consumer, the opportunity for expansion is substantial. Nvidia estimates that the AI agent market will represent about a trillion dollars, roughly the same as the enterprise software market the company predicts the agents will replace. We believe this market could be significantly larger when combined with consumer use-cases. So, while efficiencies and monetization opportunities from agentic AI are becoming evident in an enterprise setting, it will be exciting to see how they can perform end-to-end processes for consumers as well.

For an AI assistant to have the necessary capabilities to be fully agentic, users need to become comfortable sharing more with their AI to understand context (time of day, normal preferences/routine, and even emotional state) and act proactively. Fortunately, AI technology is particularly skilled at understanding the intentions of the user; it can take vague and disorganized references, compare it with the corpus of information at its disposal, and produce specific outcomes using the user's context. In our opinion, the most impactful feature of LLMs for consumers is the ability to learn and retain knowledge from past conversations. However, since every consumer is different (and likely to use LLMs differently), conversational AI systems need to adapt to specific use-cases. This presents another hurdle as consumers' concerns are likely exacerbated by AI agents' need to integrate with tools like calendars, emails, personal files, etc. Therefore, it will likely take time for consumers to become familiar with this level of personalization, but the process can be sped up if companies like OpenAI stress the importance of protecting user data and put guardrails in place to combat a leak or its use in training data.

Another important consideration in the dissemination of consumer AI is how apps evolve given that AI technology will likely become users' primary interface for acquiring or accessing information, services, and goods. Said otherwise, rather than using a phone to navigate through different apps, AI likely becomes the bridge that enables us to access the content we desire. This is not to say that AI will not be integrated into the hardware and software of mobile phones, but we would imagine another, more natural interface spawns from the widespread use of AI in the future. Apps are currently essential to everyday life because of how computing has evolved until now, but AI can orchestrate this navigational process since it knows what services we already use and is better at inferring what the consumer wants. For example, a consumer may express to its AI assistant that it wants to listen to a certain song; in an AI-first world, the consumer likely does not care which music platform the song is streamed from as long as the audio quality is good. Therefore, the AI assistant would source the song from the highest-quality service at the lowest price, assuming internet app subscriptions become a thing of the past. As a result, navigating across multiple different services and brands' apps may fall out of favor since elements like the user interface, brand trust, and specific features will not matter as much to consumers as they are not operating the interface.

While we do not believe that the use of apps will disappear anytime soon, it is likely that LLMs will enable access to most of our desired content and/or information, whether that is being sourced from an app or not. Meta CTO Andrew Bosworth projects that many of the most impactful consumer AI apps will take shape once consumer AI adoption is widespread, consumers realize the dead ends of AI capabilities, and developers learn to bridge that consumer demand to a marketplace that prioritizes performance over brand affinity. As this occurs, we assume that most of the heavy lifting in developing these apps will be performed by AI and the last mile of development will serve as the difference-maker between an impactful consumer AI app and not. However, the world is currently built on phones and apps, and if this changes, we believe it will evolve over time versus a dramatic near-term change. In the short to medium term, we believe three-sided marketplaces with a physical component (such as Airbnb) will remain largely intact with added GenAI modifications. However, in the longer term, apps based on asset-light business models may need to transition as brands become less important due to LLMs serving as the gateway to information and consumer products and services. The result is perhaps new marketplaces that prioritize price per performance, which we view as a net positive for consumers.

What Does This Currently Look Like?

The latest generation of AI tools is pushing the boundaries when it comes to performing tasks on behalf of users. The newest models from developers like OpenAI, Google, and Anthropic are beginning to act on behalf of users when making travel plans, making purchases, or communicating with others. We believe near-constant access to AI will open even more possibilities and use-cases for consumers leveraging this technology. The current phase of AI app development is focused mostly on seeing what works and what provides the best value for consumers today. Once further advancements in AI emerge, there may be a consolidation in the number of apps that dominate each aspect of consumer lives, in our view.

Walmart released a blog post in May 2025 outlining some of the changes the company expects in the next three to five years due to agentic AI and what it is doing now to prepare. The company is working to optimize its e-commerce site to adhere to several digital shopping agents navigating the website at once on behalf of Walmart customers. The biggest finding regarding the company's testing, in our opinion, is that agents work best when implemented for very specific tasks as the output from several task-specific agents can be stitched together for Walmart to solve complex workflows. Today, the company is using AI agents informed by trend signals (e.g., what teenagers have frequently purchased in the past month) to help source its products. We also expect that soon AI agents will purchase products on consumers' behalf as well, creating a natural opportunity for AI agents to interact with each other. There are several companies that have found task-specific agents to be the best approach for implementation in an enterprise; for example, financial services

platform BNY has about 100 digital employees overseen by a human manager that are equipped with unique logic credentials and can communicate via email or Microsoft Teams. These agents often tackle low-complexity problems and relieve the brainpower of their human peers to focus on mission-critical tasks.

Gaining access to agentic capabilities on the consumer side has become fairly easy in recent months, especially given the release of AI browsers with built-in agentic capabilities. In addition, there are several no-code or low-code platforms that consumers can use to create agentic workflows, but most commonly, prebuilt agents are being offered as part of a subscription to LLMs. A good example of a low-code platform for orchestrating workflows is ***n8n*,** which allows users to develop multistep AI agents for a variety of use-cases. This work is performed within a drag-and-drop canvas by connecting nodes (e.g., apps, APIs, LLMs) and allows the user to add custom logic using JavaScript or Python, if desired. The company raised a $180 million series C round in October 2025 at a $2.5 billion valuation. As consumers' familiarity with AI technology grows, the proliferation of platforms that can combine all the necessary tools and LLMs into automated intelligence will become more widespread. In the early days, however, the development of AI browsers appears to be the primary catalyst in bringing agentic capabilities to the hands of the consumer. One example of this is OpenAI offering "agent mode" exclusively as a feature of its Atlas browser, which was released in October 2025.

For the agent to perform such a broad range of tasks for consumers, it was given a text browser to parse through large amounts of text content for simpler, reasoning-based queries. It was given a visual browser to interact with websites and navigate graphical interfaces like a human would. On top of that, the agent was given access to its own terminal for coding along with generating and analyzing files. The model moves between these capabilities through a virtual computer that exists within a browser-native virtual computer and mimics a traditional operating system. The model was trained using reinforcement learning to effectively navigate between its capabilities and accomplish multistep workflows autonomously. The agent can call on public APIs and private APIs (e.g., GitHub, Google Drive, SharePoint) that have been connected to ChatGPT to access various datasets. It was also given access to image generation APIs to create visuals.

Exhibit 10
Generative AI Report
OpenAI's Atlas Browser

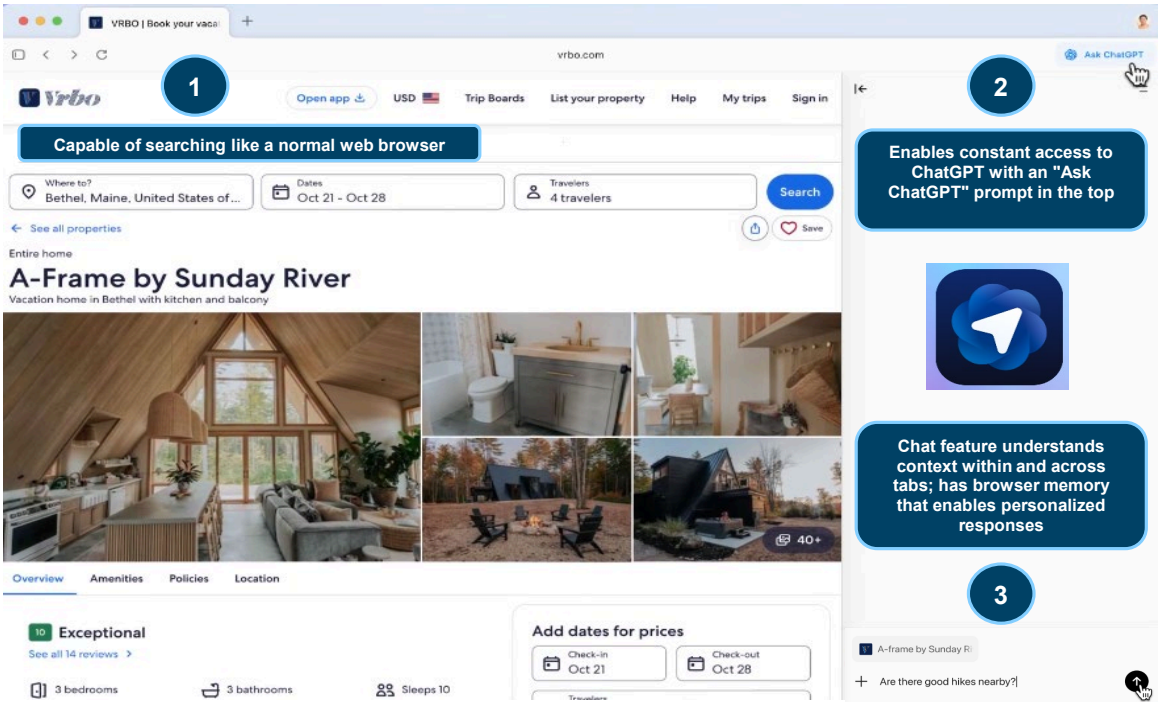

Sources: ChatGPT and William Blair Equity Research

Other consumer-facing agentic products include Perplexity's Comet, ***The Browser Company's*** Dia, Microsoft's Copilot, and Google's Gemini in Chrome, among others. Comet and Dia are AI-native browsers that offer an always-available AI assistant, provide cross-tab interconnectivity, and perform multistep tasks. In contrast, Copilot and Gemini in Chrome have existing AI models with capabilities that have been integrated into the respective companies' existing browsers. The features that are currently available across several of the AI search engines released in 2025 are largely the same. Many of these AI-driven search products emerged as Google's Chrome browser faced legal challenges. In response, Google intensified its efforts by releasing several of its own AI solutions that are performing better than expected. While it is likely that the hype behind many of the competing AI browsers and consumer tools from start-ups will fade over time, a select few may endure, potentially carving a path to profitability for the most resilient developers.

AI Integration in Hardware and Edge Devices Could Be the Needed Catalyst for Adoption

One of the biggest hurdles in enabling personalized AI assistants is training them to be unique to the user. A large value-add for consumers using AI agents is saving time and energy on daily activities and tasks. For AI agents to do so, they must be equipped with knowledge and context. As referenced earlier, context can refer to a user's daily routine, preferences, attitude, or anything that makes the user the person he or she is. Primed with this information, AI agents can plan and optimize day-to-day activities and intuitively assist the user. We believe that edge devices and in-home hardware can be the first step in familiarizing consumers with this technology and will allow AI to gather the necessary knowledge and context. This can include glasses, watches, home devices, pins, thermostats, etc. Devices like these can gather the necessary context through visual and auditory inputs while simultaneously understanding users' intentions based on the physical world around them. Further, progress in robotics could be a more advanced use-case, especially as companies such as Tesla scale development of robots, in our view.

Our team generally subscribes to Andrew Bosworth's viewpoint that AI will usher in an evolution of the technological interface we interact with every day (i.e., the mobile phone). Bosworth noted that AI technology is not domain-specific and can be applied to solve a multitude of problems across different interfaces. As referenced earlier, mobile phones are the result of the evolution of compute thus far, and it is likely that humans would prefer an interface that prioritizes our eyes and ears over a keyboard and touchscreen. We believe that AI offers this possibility as it is a flexible technology that can be applied across different interfaces like edge devices, and that is validated by the multitude of wearable products released over the past year. One of the latest releases is a smart ring developed by ***Sandbar*** that is designed to be worn around the index finger and can record any thoughts or ideas that are spoken into the ring. Any words recorded by Sandbar's ring are organized and made accessible through the company's app. This product displays how consumer AI products are still heavily reliant on the phone for accessibility but are beginning to familiarize consumers with everyday use of AI technology outside the mobile phone.

The LLMs that most consumers are familiar with currently require inference to be performed on the cloud, which means they need to be connected to the internet. This is not always possible or feasible, especially with edge devices. The solution is likely a much smaller, more optimized model that we will refer to as a small language model (SLM). These systems are designed to handle specific functions and commands locally by using high-quality data and the right set of tools and APIs. SLMs can allow edge devices to run without a network connection at lower latencies, while better protecting the user's privacy. This shift aligns with the broader evolution of edge computing, which aims to move inference closer to the device to reduce reliance on centralized servers and internet connection. In some cases, edge-cloud collaborative computing is used to split tasks in real time between the device (for immediate tasks) and an off-premises data center (for heavier workloads). For example, Apple recently announced that GPT-5 will be incorporated into future generations of its iOS. Within this partnership, training will remain cloud-based, but inference is being pushed locally. This hybrid approach mitigates hardware constraints while improving responsiveness and security.

Overall, the trade-off in using edge devices is that they cannot perform tasks that require heavy processing and are therefore siloed to simple function-calling tasks like speech and gesture recognition, summarization, object detection, etc. An example of this approach is Meta's lightweight Llama 3.2 model, which has only 1 billion parameters, enabling it to run efficiently on smartphones, wearables, or home devices. Limitations to these models include narrow scope, risk of bias from training on smaller datasets, and reduced complexity and ambiguity. Despite these hurdles, AMD CTO Mark Papermaster predicts that most inference will be on-device by 2030. We believe that as technology advances, high-quality data coupled with model compression/inference techniques can enhance the capabilities of localized edge devices and potentially pave the way for widespread AI adoption, as the ease-of-use becomes impossible to ignore.

Can Meta's AI Glasses Be a Consumer AI Adoption Catalyst?

We believe Meta's AI glasses are the front-runner in hardware devices with the potential to kickstart the AI consumer boom. While investors still debate if the form factor will be successful, Meta remains optimistic about this opportunity. For reference, the Meta Ray-Ban Display glasses, released at the end of September 2025, sold out within 48 hours of launch. However, the challenge with creating products in any technology cycle is gaining consumer adoption. Requiring consumers to wear AI glasses, at least initially, is an obvious concern to adoption. However, Meta continues to highlight an increasing rate of adoption, and we believe the form factor in terms of style, battery life, etc. will only improve. Meta CEO Mark Zuckerberg has gone as far as to say that people who do not wear AI glasses in the future will be at a cognitive disadvantage, similar to visually impaired people who do not wear prescription glasses or contacts. In addition, we believe that features and capabilities such as calling, texting, listening to music, capturing image and video, and chatting with Meta AI will multiply with time and better hardware design.

Exhibit 11
Generative AI Report
Meta Ray-Ban Display Glasses

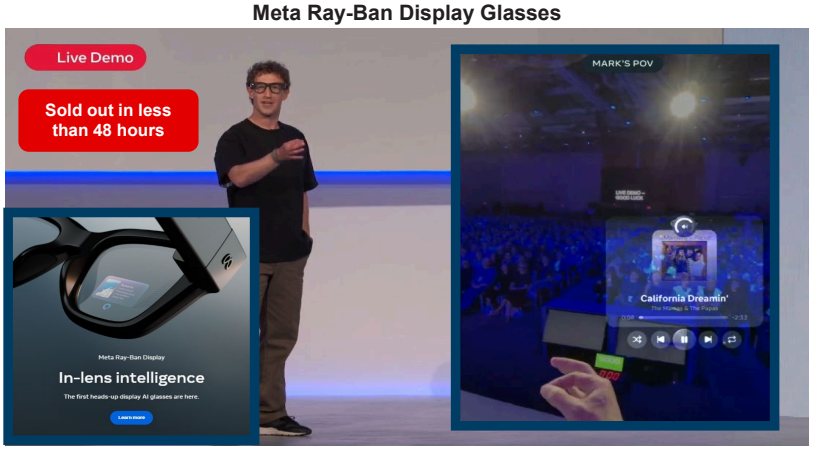

Sources: Meta Connect and William Blair Equity Research

As referenced, Bosworth has expressed confidence that devices outside our phone and TV will increasingly be used to interact with content. He was positive about the Meta Ray-Ban glasses before the proliferation of AI, when the team at Meta insisted that AI would make interactions richer with an assistant that understands not just device inputs, but also the world surrounding the user. He believes that glasses are the ideal form factor and that they enable better experiences and interactions with the content users seek. Despite this, he acknowledges that completely replacing mobile phones within five years does not seem feasible. During an interview with a16z, he offered a unique perspective on consuming content where users can directly access the information they seek rather than navigating through certain apps and potentially getting derailed. For example, while scrolling Instagram, users might come across a video highlight of their favorite basketball team, prompting them to seek the box score for the game. Rather than looking up the box score on the ESPN app and then potentially navigating to YouTube's app to watch additional highlights, the AI would likely populate all the relevant information in their display almost instantaneously. This would fundamentally alter how people interact with devices; as opposed to users navigating from app to app seeking the information desired, AI would retrieve it with an unparalleled understanding of the format in which users like to consume content and the sources users seek it from.

The glasses form factor appears to be picking up steam as Google recently announced that it expects to release its first pair of AI glasses in 2026. The company is doing so with help from Gentle Monster and Warby Parker and has committed $75 million, which will be met with another $75 million and an equity stake if Warby Parker achieves certain milestones. The only details offered about the products are that one model will allow the user to communicate with Gemini and take photos and another will offer an in-lens display capable of directions and closed captioning. This is a far cry from the current capabilities of Meta's AI glasses, but we believe that the Google development team will eventually create a rival product. Further, investments into AI glasses by two hyperscalers with a particular focus on the consumer is another bullish indicator that glasses can become the ideal form factor for AI in the future.

Wearing glasses/goggles has proved ideal for consumers seeking augmented reality (AR) and virtual reality (VR) experiences in past years. We believe that AI can enable even better experiences in these formats, like watching your favorite basketball team courtside without needing to travel or purchase a ticket. We believe there will be advancements in other AR form factors as well, such as enabling AR via contact lenses. For example, ***Mojo Vision*** is developing technology to enable AR capabilities in contact lenses; however, there is currently not a consumer product launched. In our view, one near- to medium-term hurdle to mass market adoption is when/if the Meta glasses

can sell for $199. Historically, this is when mass market adoption occurs for consumer devices, measured at 30% of the population. There will likely be additional AI form factors available to consumers, such as devices pinned to clothing, contact lenses, or hardware that sits on a desk for those who do not want to wear glasses.

Overall, we believe agentic AI is poised to fundamentally reshape consumer experiences by enabling autonomous, context-aware digital agents that can streamline daily tasks and decision-making. While most of the advancements in agentic AI are being applied to enterprise use-cases, we believe it is only a matter of time until the technology extends to the hands of the consumer. As hardware and edge devices become more integrated with AI, widespread adoption will depend on balancing personalization, privacy, and ease-of-use to unlock the full potential of these transformative technologies.

Key Predictions and Trends to Watch for Investors

1. The importance of apps and brand affinity may wane with AI emerging as the primary interface.

2. Every consumer will likely have a personal assistant/agent.

3. Agents will disrupt the existing online marketplace ecosystem.

4. Meta glasses adoption can be a major catalyst for consumer AI adoption if the price drops below $199 per unit.

5. Edge devices and in-home AI hardware will be the first step in achieving broad adoption of consumer AI.

Voice Modality

Introduction
As edge devices and other hardware become more prominent, we believe that voice will be the most natural and efficient interface for interacting with AI. Adoption will be driven by breakthroughs in speech recognition, latency reduction, and expressive text-to-speech (TTS). As voice capabilities improve, new consumer use-cases and hardware integrations will accelerate adoption, in our view.

Breakthroughs in AI-Powered Voice Technology Have Contributed to Rising Consumer Adoption
In our view, AI-powered voice technology has the potential to transform everyday consumer experiences by enabling hands-free interactions. One of the most important advancements thus far to enable this has been the improvement of speech-recognition accuracy. Experts indicate that error rates have been reduced to that of a human, and sometimes better, which has solved a limitation that substantially restricted the usability of voice modalities. Accuracy rates have improved to about 95% with LLMs, up from roughly 50% 10 years ago. Further, these improvements extend to various languages, accents, and intonations, making conversational voice more widely adoptable. As a result, speech-recognition accuracy is no longer a differentiating factor, but rather the bare minimum needed for voice models.

The focus in AI-powered voice has shifted recently toward enhancing the natural feel of conversations and minimizing latency. Researchers have been particularly focused on tailoring model behavior to the more subtle aspects of human interaction (e.g., pausing between breaths, speech being cut off, taking time to think) that make speaking with AI feel more seamless and intuitive.

On the AI side, user speech gets converted to text, a response is generated in the LLM using context management and reasoning where appropriate, and then the text is converted to speech and output to the user.

Exhibit 12
Generative AI Report
AI-Powered Voice Flow Diagram

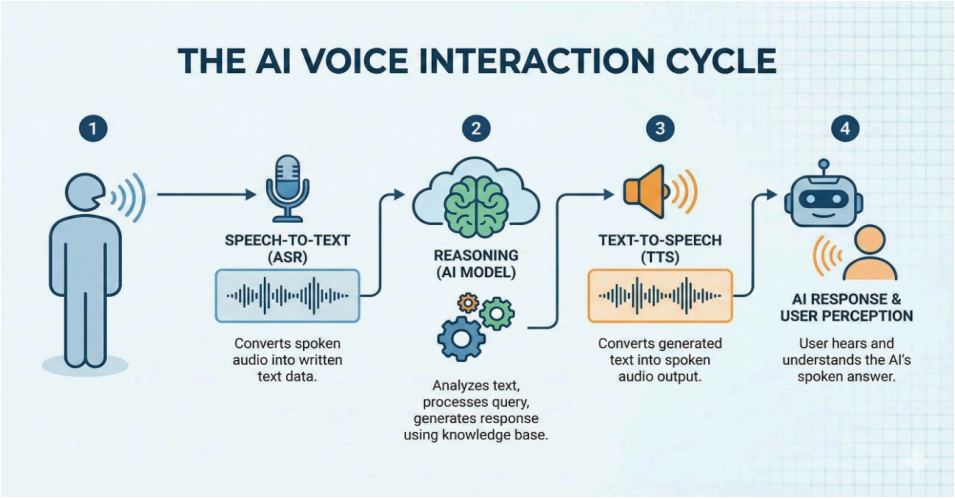

Sources: Gemini and William Blair Equity Research

Speech-to-text (STT), also known as automatic speech recognition, is an important part of reducing latency in voice models. Developments in STT have made it fast and accurate, with LLMs combining large datasets, multilingual training, and transformer architectures for near-human-level transcription accuracy. More recently, this has evolved in edge STT, which is particularly relevant for smaller hardware devices. For example, Google Cloud's STT On-Device models are kept to under 1 GB but provide high transcription accuracy and low latency for use on small devices. The on-premises nature of the solution is promising and will help in-home devices like Siri and Alexa evolve into more conversational, GenAI solutions, in our view. However, storage remains a key barrier. If the AI cannot respond to a query using the on-device model, it will still require a connection to the cloud, increasing latency.

Another breakthrough in voice centers on expressive TTS, which includes tone, intonation, prosody, and emotions, making interactions with AI more human-like. The ChatGPT voice modality has become fairly realistic through the evolution of the product, in our view. Its combination of relatively low latency and expressive STT provides a human-like experience, particularly for easy conversations where the AI does not have to "think" too hard. These progressions have been made possible by replacing rigid decision trees with open-ended processing capabilities. Open-ended processing capabilities allow LLMs to understand the intent and context behind a question or command and then generate an appropriate response in real time. This evolution makes any AI conversation feel less robotic and more familiar to the user.

These technical breakthroughs in accuracy, speed, expression, and open-ended capabilities have created AI agents that are able to hold dynamic, rather than predetermined, conversations. However, one of the major obstacles remaining is latency. Human conversations tend to experience latency of 200ms to 250ms. LLMs are making progress toward that goal, with OpenAI's GPT-4o achieving an average latency of roughly 300ms in its voice demo. Edge models have the potential to be even faster, although the data storage limitations remain. Continued development of AI models

will likely reduce latency to a point that is indistinguishable from human conversations. That kind of innovation will likely occur only as consumers increasingly trust in AI, leading to greater consumer adoption, and more on-device and edge inference possibilities, in our view.

Specific Consumer Use-Cases Are Emerging With Voice as the Primary Modality

As AI-powered voice has improved, numerous use-cases have emerged. In general, voice represents a more natural interface to interact with AI, in our view. Voice allows for faster inputs and more fluid interactions, as opposed to typing out a query. Further, hands-free and eyes-free interaction allows for multitasking and improves productivity, which is ultimately the goal of AI assistants. This technology can also empower users with visual, motor, or literacy issues. However, outside the generic use-case of voice as a primary mode of interaction, there are several specific consumer AI use-cases where voice, particularly a human-like voice, plays an intrinsic role, including in therapy/life coaching, language learning, and interviews, among others.

Exhibit 13
Generative AI Report
AI-Powered Voice Use-Cases

Company	Vertical	Use-Case
Ash	Therapy	AI therapist that users can speak with
duolingo	Language Learning	AI tutor that users can practice conversing with
Alex, Lightscreen, WAYFASTER	Interviewing	Résumé sorting and screening interviews

Source: William Blair Equity Research

- ***Therapy/life coaching.*** Ash, an emotionally intelligent AI therapist, is on the forefront of changing consumer norms, raising a $93 million series A in July 2025 led by Radical Ventures and Forerunner Ventures. Developed by ***Slingshot AI,*** Ash markets itself as an always-on therapist. The company's mission is to increase access to therapy, although it does not claim to replace traditional therapy. Therapy is a difficult space for AI to tackle, given regulatory and privacy concerns and the personal relationship patients have with their therapists. However, clinicians are keeping an open mind about the tool given that it broadens access to mental health help. Since speaking openly is essential to getting the help one needs, increasing consumers' trust of the tool by blurring the line between AI and human voice is likely quite important, in our view, and relies on continued innovations in AI-powered voice.

- ***Language learning.*** Duolingo's premium subscription tier, MAX, includes video calls with an AI chatbot named Lily, an animated character with a distinct personality. This feature allows users to apply what they have learned throughout Duolingo's courses in a lifelike video call with Lily. This feature builds on the Duolingo curriculum and allows the user an opportunity to practice speaking the language in real-world contexts.

- ***Interviews.*** Several companies have released AI agents that can handle screening interviews, including ***Alex, Wayfaster,*** and ***Lightscreen.*** These agents are not restricted by schedules, meaning they can interview a candidate almost immediately after résumé submission, while the candidate is still in search mode. These platforms can also handle other related tasks, including application intake, résumé scheduling, and résumé screening, and assist in hiring decisions, reducing the operational burden on recruiting teams. As voice continues to evolve and latency shrinks,

candidates will be less likely to distinguish between AI and human interviewers, particularly in early-stage screening calls where questions are often technical. This is critical to success, in our view, as trust between both parties in an interview is paramount. This use-case becomes even more pronounced for higher-volume roles, for which hiring teams sort through hundreds or thousands of résumés and may conduct many screening calls. An August 2025 study out of the Philippines indicated that for a sample set of 70,000 applicants, an AI-powered voice recruiter outperformed a human recruiter in hiring customer service representatives in terms of number of offers, job starts, and job retention. While it does not seem likely that voice agents will fully replace traditional recruiting teams, they certainly represent a tool to increase efficiency and performance without erasing trust between an interviewee and the hiring team.

Challenges Remaining for Voice to Achieve Widespread Adoption as the Primary Mode of Interaction With AI

Despite the progress made in AI-powered voice and the many potential use-cases, there are still several hurdles and challenges to overcome, including social acceptance, latency, privacy and regulatory concerns, monetization, and consumer trust. First, it might take time for users to become more comfortable speaking to AI assistants in public. The prevalence of people speaking on the phone through AirPods may alleviate the naivete of this concept slightly, but there is already a large population that has doubts about AI and are hesitant to fully adopt it, so making conversations with AI assistants via voice more socially acceptable will likely not happen overnight.

Second, latency remains a barrier. So far, general models have not achieved the sub-200ms bogey for latency. Narrow models with specific use-cases in ideal conditions have achieved it, proving the technological capability. However, general purpose models under live conditions still have not, as optimizing all layers of the pipeline is complex. Once consistent sub-200ms latency is achieved, we believe that voice may become the main modality for communicating with AI. Experts generally agree that the base-case scenario for this to be achieved, given the rate of technological progress, is two to three years away.

Third, there are privacy and regulatory concerns. This is less of a challenge for fully on-prem edge devices, but for hybrid inference models that still send information or tasks to the cloud, this is potentially a barrier.

Fourth is monetization. Since it is still early in the race to capture market share in the voice modality, companies are forced to keep prices low to gain a user base, which makes monetization a challenge, especially as they compete with hyperscalers and start-ups with substantial funding.

Last is consumer trust, since AI is still susceptible to hallucinations, and that is unchanged with AI-powered voice. One potential remedy here is vertical-specific AIs that have domain expertise, reducing the risks and prevalence of hallucinations. Further, challenging audio environments represent another roadblock in STT translation. To date, work around noise-augmented training, acoustic modeling, and voice activity detection has resulted in improvements. Being able to navigate challenging audio environments will also lead to increased trust among consumers, in our view. Despite these challenges, AI-powered voice has made significant progress and is quickly becoming the ideal form factor for consumer use-cases of GenAI technology.

Key Predictions and Trends to Watch for Investors

1. Voice may become the dominant modality for consumer AI once general-purpose voice models achieve sub-200ms latency, which could happen within two to three years.

2. Verticals like therapy, education, and recruiting will see early adoption of voice-native AI agents with measurable ROI.

3. Edge computing will enable real-time voice interactions, pushing demand for low-latency, privacy-first hardware solutions.

4. Companies that combine expressive voice with agentic capabilities will differentiate in a crowded consumer marketplace.

5. Monetization of AI-powered voice will lag initially but accelerate as trust, usability, and vertical specialization improve.

GenAI Is Changing Search and the Traditional Digital Advertising Ecosystem Landscape

Introduction

As AI technology continues to evolve, we believe the most disrupted vertical in our coverage universe will be search because of changes in how users discover and interact with information online. Traditional search engines are being disrupted by natural language chatbots and emerging AI browsers, creating ripple effects across advertising, publishing, and digital commerce. For investors, this shift presents both risks and opportunities as companies adapt to new user behaviors and monetization models.

State of the Search Landscape

Today, certain queries submitted to Google now trigger an AI Overviews response at the top of the results page, displaying a more conversational response than the traditional list of blue links.

Exhibit 14
Generative AI Report
AI Overviews Example

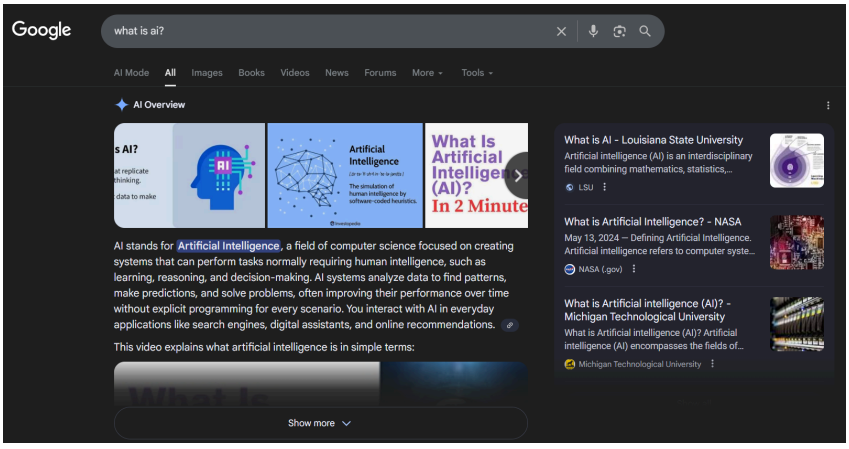

Sources: Google Search and William Blair Equity Research

This provides users with a more interactive and potentially streamlined experience, reducing the need to click on and review various links, and instead being presented with a summary answer and a handful of citations. The frequency with which queries trigger AI Overview has grown since the product's onset. A blog post published by Semrush noted that in November 2025, 15.69% of all queries triggered an AI Overview, up from 6.49% in January 2025. Many of these queries are informational, while just a fraction are navigational, although this subset is growing quickly. Informational queries (e.g., "what is AI?") represent low cost-per-click (CPC) and low-difficulty

queries. According to Semrush data pulled from January to March 2025, over 70% of the keywords that triggered AI Overviews had a CPC under $0.10, and almost 60% were in the most infrequent category of search volume. This indicates that Google is testing its AI Overviews on queries that are difficult to monetize and relatively infrequent, so as not to significantly disrupt its advertising revenue model over the short term. Industries experiencing the larger share growth of AI Overviews include science, health, people/society, and law/government. While navigational queries showing up in AI Overviews are lagging informational queries, the growth rate cannot be ignored. Semrush noted that the number of navigational queries that trigger AI Overviews increased more than tenfold between October 2024 and November 2025.

Alphabet's management has noted on recent earnings call that AI Overviews are now running fully on Gemini. Further, conversations we have held with a digital advertising industry participant, who manages over $1 billion of advertising spend across 250 companies, have indicated that Google is integrating its various properties (e.g., Gmail, Google Maps, Google Calendar), several with over a billion users each, into its training model, allowing Google to improve its ad targeting and the overall user experience. One example our participant mentioned centered on an integration with Gmail. Since about half of the world's personal email accounts are on Gmail, and most receipts are now delivered via e-mail, Google can integrate valuable data on what users are purchasing. It is not unreasonable then to assume this data could be used to improve AI Overviews, particularly when it comes to high-CPC commercial and transactional queries. Our participant estimates that in three to five years, Gemini will evolve to replace traditional search.

The rise of AI search, as evidenced by growth in overall queries in Google's search business, does not appear to be the result of its novelty, but because it is leading to better results. Several companies in our coverage have mentioned that leads generated through AI chatbots often represent higher-intent consumers, often with **3 to 5 times higher conversion**. This indicates that users are using chatbots to perform in-depth research on products and services and are trusting the responses enough to purchase the product. For example, NerdWallet observed materially higher monetization rates from users arriving via AI channels in the second quarter of 2025, while ZipRecruiter observed a 140% sequential increase in site visits from GenAI start-ups in the third quarter of 2025. These results only reinforce the case for AI search overtaking traditional search.

The Emergence of AI-Powered Browsers

AI browsers came into focus for many investors in the second half of 2025. Common features of current AI browsers include AI chat assistants, basic search functions that resemble traditional chatbots, personalization, cross-tab interconnectivity, app integrations, and task automation. As usage increases, the rate at which traditional search degrades may accelerate, as virtually all searches will eventually return a response akin to an AI Overview. In addition, if a user is shopping for a Titleist golf driver on Dicks Sporting Goods' website but wants to explore similar, less expensive options, the user could simply prompt the AI chat assistant to do so. This technology represents another inflection point in the shift from traditional search that companies and advertisers will be forced to adapt.

This shift may make the browser space more competitive for Chrome, as it relates to Google. However, we believe the company is still well positioned to adapt to growing competition. Historically, rival browsers like Microsoft Edge have not had material success in competing with Google Chrome. Recent data from Stat Counter shows that Google Chrome commands a roughly 70% worldwide market share, while Microsoft Edge remains below 5%. While browser competition will likely intensify, we believe Google will be successful in retaining a material market share following the integration of AI in Chrome, although it may cede some share if new browsers gain increased traction. In the fall of 2025, Google launched Gemini in Chrome to transition the browser into one that is AI-centric. Gemini in Chrome can integrate context among several tabs, answer

questions, and work with other popular Google services such as Google Docs and Google Calendar, with additional agentic capabilities to follow. Given Google's distribution capabilities and data advantage, we believe it is well positioned to deal with the increasing browser competition.

Exhibit 15
Generative AI Report
Gemini in Chrome

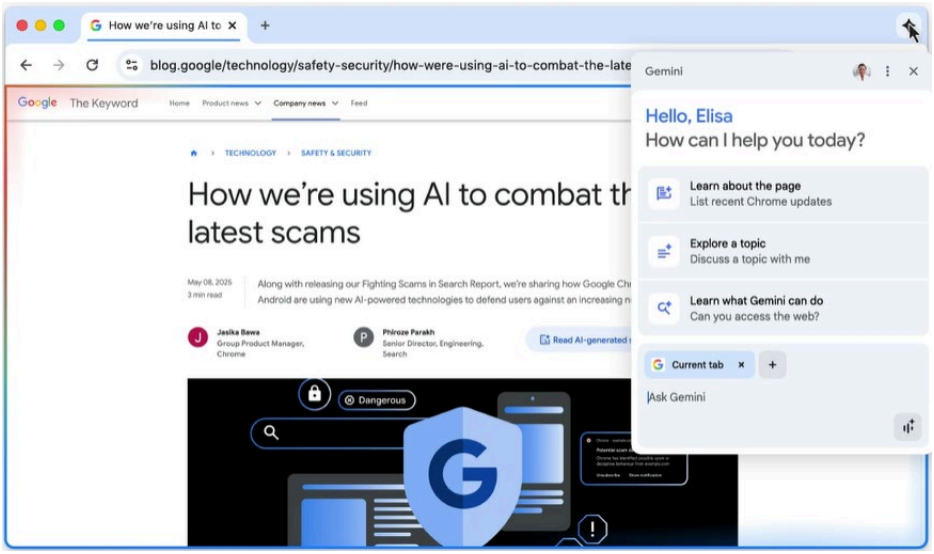

Sources: Google Chrome and William Blair Equity Research

Shifting From Search Engine Optimization (SEO) to Generative Engine Optimization (GEO)
Since AI Overviews are featuring more prominently among navigational and transactional queries, the relevance of traditional SEO strategies is eroding. Enter GEO—also referred to as AIO (artificial intelligence optimization) or AEO (artificial engine optimization)—an evolution of SEO that is tailored toward AI-generated, natural language responses, rather than a list of blue links.

- ***SEO.*** A foundational digital marketing tool for the traditional search landscape, professionalized in the late-20th century when Google released PageRank. The end-goal is to push a website as high up on the results page as possible in response to certain queries. SEO achieved this goal by focusing on keyword density, meta tags, and backlinks to rank search results, primarily focused on Google's 10 blue links format. As we have framed throughout this section, it is increasingly likely that this format evolves, at least for certain types of queries.

- ***GEO.*** An evolution of SEO focused on optimizing content for AI-driven search platforms (e.g., AI Overviews and ChatGPT). The strategy requires content emphasizing topic authority and contextual relevance as opposed to focusing on keyword density to influence AI-driven answers. As users increasingly adopt AI platforms for search, and trust a single answer as opposed to clicking through several blue links, the importance of GEO will increase, in our view. Brands will need to adapt marketing strategies to be pulled into the first answer from an AI chatbot, rather than appearing on the first page of Google's search results. The overall goal remains the same, but the techniques to achieve the goal differ. GEO requires an understanding of the underlying AI models and datasets that chatbots are pulling from, as well as an understanding of how they "choose" which data to pull. To date, marketers are struggling to understand this because of the complexity of these AI models, but it is clear that brands that can adopt and successfully integrate these strategies will benefit as more users shift from traditional search engines to AI agents.

For GEO to be a successful strategy, companies and marketers will need to develop an understanding of the types of content required to be featured in AI Overviews. To date, the sites being included the most within AI Overviews tend to feature visual-rich and structured content types, such as YouTube, Reddit, and Quora, with an emphasis on user-generated content. Other companies will need to evaluate and update entire web profiles to gain visibility in AI Overviews and other chatbots in its current form. Potential updates include refreshing websites to feature low-keyword-density, informational terms that are commercially valuable (i.e., "why does my back hurt after sitting at my desk all day?"), creating content that mirrors the way consumers/users are shifting to natural language queries, and study which competitors are currently visible within AI Overviews and structure websites/content similarly. Another suggested change for publishers may be adding more tables and data on websites so the LLMs can crawl and prioritize in chatbot results as opposed to graphics that LLMs may not prioritize.

One of the main challenges of this shift has been measuring brand visibility and sentiment in AI search responses because traditional search analytics do not apply. As an example, Google provides advertisers with concrete and actionable data through Google Search Console and Google Analytics. Google Search Console tells advertisers what specific keywords lead to traffic, CTR, number of impressions, average ranking position, and any potential issues limiting Google's ability to crawl or index specific webpages. Google Analytics goes a bit deeper in analyzing sources of traffic and tracking user actions on-site, giving advertisers a deeper understanding of how users are interacting with brands. As of now, LLM developers are not sharing this actionable data. Our industry participant confirmed this, noting that Google determines the display of AI Overviews independently of the ad stack, making it difficult to bid directly on those placements. Instead, our participant has found that the most effective strategy so far is to rank at the top of organic search results to increase the likelihood of appearing in the AI Overviews slot. Compared with the wealth of data Google provides for traditional search ads, this lack of transparency presents a challenge for advertisers.

As a result, brands are struggling to understand visibility within AI-generated responses. For example, in an August 2025 blog, Backlinko reported a decrease of 15% in total clicks to its website, while simultaneously reporting an increase in impressions of 54%. These results tell a story of brand decline from an analytics perspective but brand explosion from an influence perspective, indicating that users are still interacting with content, but fewer are clicking through. This makes it hard for advertisers to understand the efficacy of their marketing strategies. Therefore, as search continues to evolve toward an AI-dominated landscape, the analytics with which brands track marketing strategies will need to evolve as well, creating an opportunity for new players, in our view. Several companies are creating products to fill this need, including ***Profound***, ***Bluefish***, and ***Pattern.***

- ***Profound*** is a GEO platform that helps brands monitor and influence how they appear in responses from AI answer engines, helping brands optimize presence in AI-powered search and other "zero-click" consumer surfaces. Profound is based in New York City and raised a $35 million series B in August 2025. The company offers tools like Answer Engine Insights, which analyzes user interactions with AI chatbot interfaces and generates actionable insights based on those experiences. Agent Analytics represents Profound's answer to Google Analytics in the traditional search landscape. Profound installs server-side trackers that log and track known and verified AI crawlers and integrates with a wide range of cloud platforms, including AWS, Azure, Vercel, and Cloudflare. Using this tool, Profound can see every time an AI chatbot pings a brand's site in response to a user query, helping brands understand how LLMs are interacting with its content. Prompt Volumes tracks trending user queries on AI chatbots. According to its website, Profound licenses conversations from multiple, double-opt-in consumer panels of real answer engineers to analyze hundreds of millions of prompts per month, allowing the company to track frequency, intent, and sentiment of certain brands. Its shopping product

applies the same strategies specifically to ChatGPT shopping. Lastly, Profound builds on those insights to help brands create the content that is architected to be cited by AI answer engines through its Workflows product.

- ***Bluefish*** is another AI-marketing platform designed to help enterprise brands monitor, analyze, and optimize their visibility in responses generated by major LLMs and chatbots. The company is based in New York City and raised a $20 million series A in August 2025. Bluefish's solution is the AI Marketing Toolkit, which provides a wide range of marketing services for the AI internet. The company monitors brand sentiment and accuracy across major AI providers; provides actionable, data-driven insights surrounding consumer engagement; launches AI advertising campaigns built specifically for the AI internet; and builds custom AI agents that can represent brands on AI media channels. Bluefish represents a centralized marketing platform for the entire AI internet.

- ***Pattern***, a technology-driven e-commerce accelerator, created a GEO scorecard, displayed in exhibit 16, to measure how brands are showing up in AI-generated answers when users are researching and shopping for products. The GEO scorecard transforms data into actionable insights and enables brands to understand both visibility and sentiment in AI-generated responses compared to competitors.

Exhibit 16
Generative AI Report
Pattern GEO Scorecard

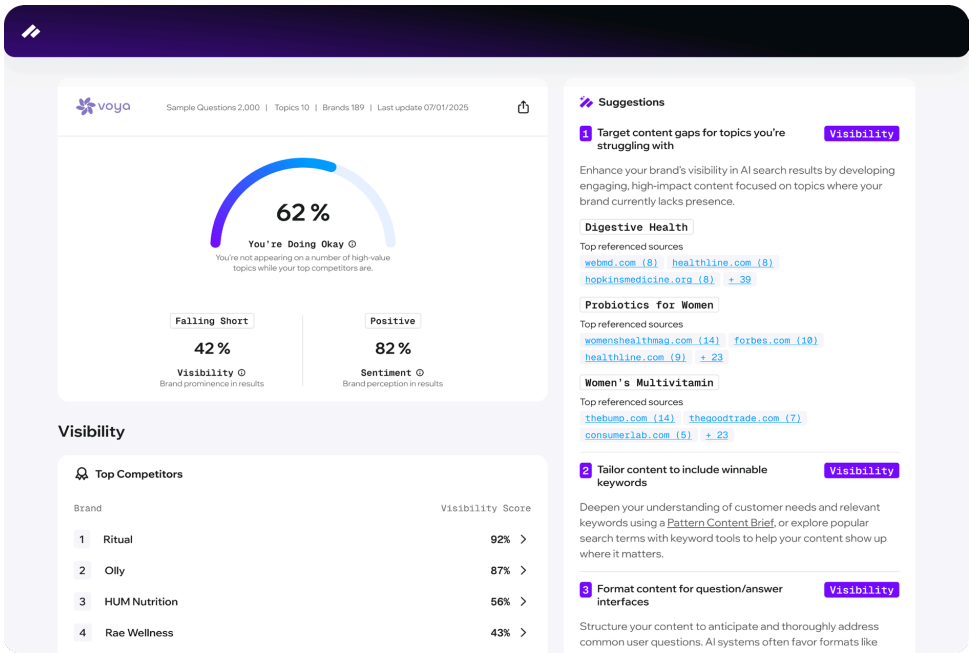

Sources: Pattern website and William Blair Equity Research

As AI search gains popularity, it is important to note that, as of a September 2025 blog posted by Ahrefs, 95% of Americans continue to use traditional search engines each month, and 86% are heavy users. In addition, there is reportedly a positive moderate correlation between ranking high on a search engine results page and being cited in LLM-generated responses. Therefore, at least for a period, there will be a place for both SEO and GEO monitoring solutions in analyzing brand

visibility. However, with trends indicating that LLM search will overtake traditional organic search in the next few years, many companies will turn to GEO-focused solutions and providers, such as Profound, Bluefish, and Pattern, to navigate this shift, in our view.

The Digital Advertising Ecosystem Is Evolving Due to GenAI and Will Ultimately Benefit, in Our View

While we expect companies and brands to leverage GEO tools to organically appear in AI-generated search responses, we also expect companies and brands to increase direct spending on digital advertising channels, both to replace lost organic traffic from SEO degradation and because AI will transform the way businesses can leverage budgets. Our industry participant mentioned that SEO traffic will decrease about 10% annually, leading advertisers to increase spend to make up for lost traffic, resulting in a tailwind for the digital advertising ecosystem. Further, AI tools are improving efficiencies in digital advertising, particularly for SMBs with small budgets and marketing teams. Our industry participant also mentioned that there are SMBs advertising on YouTube without even realizing it because Performance Max uses nano-banana, Google's video and image generation model, to create video ads for them, unlocking advertising channels that were previously too cost-prohibitive. Overall, we view the continued evolution of GenAI as beneficial to the digital advertising ecosystem, especially from the point of view of SMBs whose capabilities will increase. As these tools increase return on advertising spend for businesses, we expect an increase in digital advertising dollars spent, representing an overall tailwind to the industry.

We Expect Google to Benefit From the Shifts Within Search and Digital Advertising

Overall, while a disruption of search and digital advertising could reasonably be viewed as a headwind for Google, we believe that Google is well positioned to emerge a beneficiary. We noted earlier that its scale, distribution, and data advantage will allow the company to compete with emerging AI browsers. However, another interesting data point from our industry participant sheds light on the near-term tailwinds associated with the degradation of traditional search and SEO strategies that we just described. With SEO traffic down about 10% year-over-year and SEO traffic generally accounting for 30% to 50% of total traffic, brands and advertisers are having to pay for that lost traffic by increasing digital advertising spend. According to our participant, this represents a 1.5% to 2% tailwind for Google in the near term, with the potential for 3% to 4% of annual revenue growth over the medium term.

Further, we believe this transformation of search will require an evolution of the search revenue model. Our conversations with industry participants also revealed that Google may be shifting from a click-based revenue model to a marketplace model as it prioritizes driving actual purchases over clicks. Since AI Overviews and other AI search methods are giving users more information before clicking into a website, the click-based model likely will transition. A marketplace model where Google captures a percentage of gross merchandise value might make more sense.

This would affect the auction bidding process. For example, Google might charge a hotel 15% of the booking value for a customer who booked through Google/Gemini rather than charging based on the number of clicks into the hotel's website effected through Google. This model would further prioritize driving high-intent-driven search experiences above clicks. Today, if a user wants to book a hotel through Gemini, the chatbot will return qualitative information followed by clickable links. However, those links lead to traditional Google search, as shown in exhibit 17. The user can then click through to the hotel website. In the future, we expect this intermediary step of traditional search to be unnecessary. Instead of Google monetizing the click to the hotel website, Google will likely take a percentage of the booking fee as compensation for directing high-intent traffic to the hotel website.

Exhibit 17
Generative AI Report
Current State of Google Hotel Search

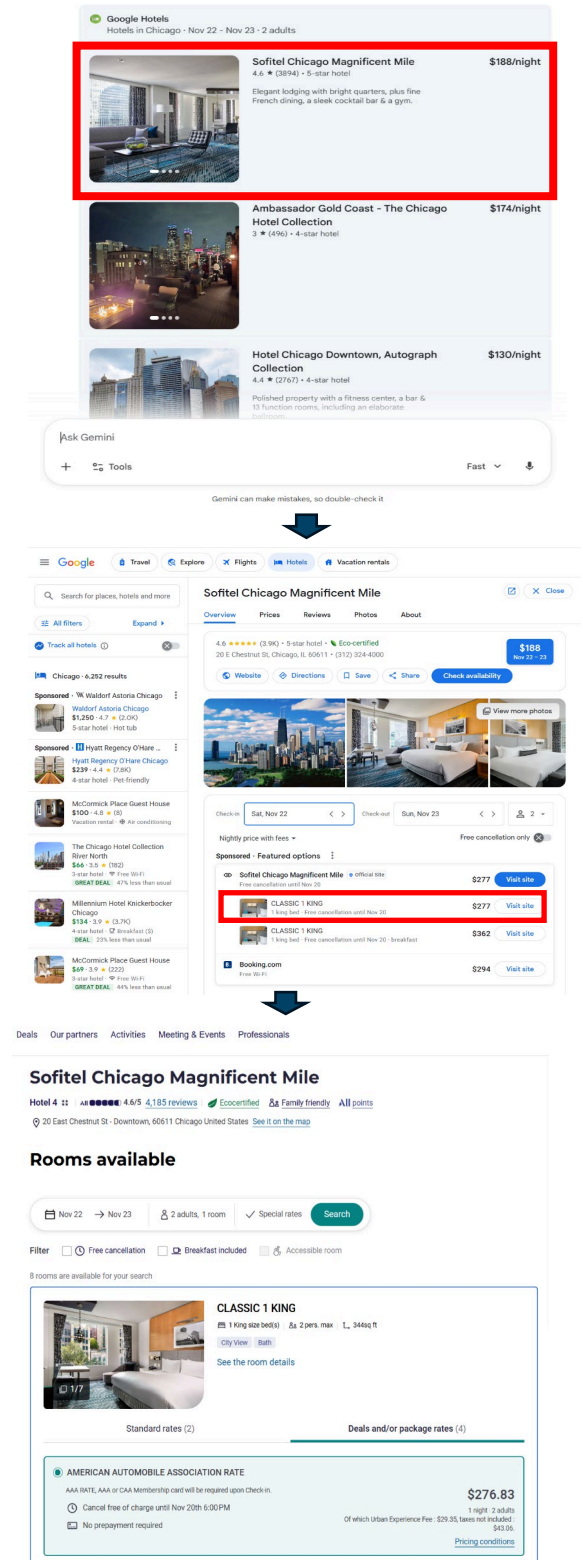

Sources: Gemini, Google Search, and William Blair Equity Research

Overall, we expect that Gemini will overtake traditional search within three to five years as consumer behavior continues to shift, leading to an evolution in Google's revenue model. However, given Google's scale, data advantage, and focus on innovation, we expect it to emerge as a winner.

Timeline
Experts seem to agree that 2026 will be a pivotal year in the space, and we believe clear winners will start to emerge in the latter half of 2026.

Key Predictions and Trends to Watch for Investors

1. GEO will replace SEO as the dominant strategy for digital visibility, creating demand for new marketing platforms and analytics tools tailored to AI-driven search.

2. AI-powered browsers will replace traditional browsers in one to two years.

3. Companies will increase digital advertising spend to replace degraded organic traffic, representing a tailwind to Google, Meta, and AppLovin, among others.

4. Google will shift toward a marketplace-based revenue model, monetizing transactions rather than clicks, which could reshape digital ad bidding strategies and margins.

5. Gemini will replace traditional Google search within three to five years.

Consumer Use-Cases

We conducted a supplemental, proprietary consumer AI survey to understand how users are interacting with the technology today. A discussion of the results can be found in the note titled First Quarter 2026 Proprietary Consumer AI Survey Highlights Monetization and Popular Use-Cases, published in conjunction with this report.

Conclusion

For consumer internet investors, the GenAI landscape presents a rare convergence of accelerating adoption, disruptive innovation, and early monetization. As the infrastructure buildout gives way to application-driven growth, companies are demonstrating the ability to scale revenues and capture market share faster than in previous technology cycles, often with creative, hybrid business models. The transformation of search, digital advertising, and consumer interfaces signals a fundamental shift in how value is created and captured, with agentic AI and voice modalities opening new verticals and deepening user engagement. While challenges around privacy, latency, and regulatory frameworks remain, the pace of innovation and the breadth of emerging use-cases suggest that the sector's flywheel is only gaining momentum. Investors who focus on companies combining cost efficiency, performance, and differentiated consumer experiences will be best positioned to capitalize on the next wave of AI-driven growth.

Appendix

The following exhibits highlight select private companies building consumer-focused products ranging from creative tools to health and wellness applications to recruiting assistants, among others.

Exhibit 18
Consumer AI
Creative & Media

Company/Product	Overview	Management	Headquarters	Year Founded
AI Gallery	Operates an online platform for creating and selling digital art generated by Artificial Intelligence.	Inderjeet Singh, Founder	Hong Kong	2022
Background Eraser	Provides camera and photo beautification applications, particularly focused on portrait retouching, makeup application, and AI enhancement.	Kenji Orito, Founder	Osaka, Japan	2016
BeautyCam	AI-powered mobile and PC app developed by Meitu for photo and video editing	Xinhong Wu, CEO Gary Ngan, CFO	Xiamen, China	2013
Black Forest Labs	Frontier AI research lab specializing in the development of state-of-the-art generative models for visual intelligence (images and video)	Robin Rambach, CEO & Cofounder	Freiburg, Germany	2024
CrushOn AI	Web-based platform that facilitates immersive, long-form conversations with AI characters.	Undisclosed	United States	2024
Cutout.pro	AI-driven platform that provides a comprehensive suite of tools for automated visual design and editing	Yong T, Founder & CEO	Hong Kong, China	2018
DeepAI	All-in-one Creative AI Platform (Generative AI for Image, Text, Video, Music)	Peter Griggs, Cofounder Kevin Baragona, Cofounder	Los Angeles, CA	2017
Descript	AI-powered video and podcast editing platform that is changing the media production workflow through text-based editing	Andrew Mason, Founder	San Francisco, CA, USA	2017
ElevenLabs	Develops advanced generative AI for audio to create realistic, contextually-aware, AI voices	Mati Staniszewski, CEO Piotr Dąbkowski, CTO	London, UK	2022
FaceApp	Creates photorealistic transformations that blend seamlessly into original images using deep learning and Generative Adversarial Networks (GANs)	Yaroslav Goncharov, CEO	Limassol, Cyprus	2014
Facemoji	Develops the Facemoji AI Emoji Keyboard application in mobile devices	Jiang Feng, Director of Baidu Global Business Unit	Sunnyvale, CA	2017
Flora	Provides tools for professional designers, marketers, and video editors who need speed, control, and consistency in their generative AI output	Weber Wong, Founder & CEO	New York, NY, USA	2024
Gamma	Focused on reinventing the presentation, document, and website creation process using generative AI as an AI-native alternative to traditional tools like PowerPoint, Google Slides, and Notion.	Grant Lee, Cofounder & CEO Jon Noronha, Cofounder James Fox, Cofounder	San Francisco, CA	2020
HeyGen	Enables businesses and individuals to create high-quality, professional videos at scale without the traditional costs and time associated with filming	Joshua Xu, Cofounder & CEO Wayne Liang, Cofounder & COO	Los Angeles, CA, USA	2020
Higgsfield	AI-powered video creation and editing company focused on giving creators and marketers cinematic control over their generative content, for scalable, professional video production	Alex Mashrabov, Founder & CEO	San Francisco, CA, USA	2023
Hypic	Develops a mobile application focused on providing an all-in-one, AI-powered solution for photo editing and visual creation	Kelly Zhang, CEO of ByteDance China	Singapore	2022
Ideogram	Creates accurate and legible text rendering within generated images	Mohammad Norouzi, Cofounder & CEO William Chan, Cofounder & CTO	Toronto, Canada	2022
invideo	Transforms text-based ideas or simple prompts into complete, professional-quality videos	Sanket Shah, Cofounder & CEO Anshul Khandelwal, Cofounder & CTO	San Francisco, CA, USA	2017
Janitor AI	Community-driven AI chatbot platform focused on creating and interacting with highly customizable, character-based conversations, primarily for entertainment, roleplaying, and creative writing	Jan Zoltkowski, Founder & CEO	San Francisco, CA, USA	2023
Kaiber	Streamlines creative workflows, enabling users to transform simple inputs (text, images, or audio) into dynamic visual narratives	Victor Wang, Cofounder & CEO Eric Gao, Cofounder & CTO	San Francisco, CA	2022
KlingAI	AI creative studio known for video and image generation capabilities	Su Hua, Founder & CEO of Kuaishou Technology	Beijing, China	2024
Krea AI	Generative AI platform designed for creating and enhancing high-quality digital content, including images, videos, and potentially 3D assets	Victor Perez, Cofounder & CEO Diego Rodriguez, Cofounder & CTO	San Francisco, CA	2022
Leonardo.ai	AI-powered visual asset generation (Images, textures, animation) with a focus on fine-tuned models and creator control	J.J. Fiasson, Cofounder & CEO Jachin Bhasme, Cofounder & CPO Ethan Smith, Cofounder	Sydney, Australia	2022
Luma AI	A generative AI company that works in multimodal general intelligence, with a primary focus on high-quality AI-powered video generation and 3D asset creation	Amit Jain, Cofounder & CEO Alex Yu, Cofounder & CTO	Palo Alto, CA	2021
Midjourney	A global generative AI company, known for its rapid development of high-quality, artistic text-to-image models	David Holz, Founder & CEO	San Francisco, CA	2021

Sources: Company websites

Exhibit 18 (Cont.)
Consumer AI
Creative & Media

Company/Product	Overview	Management	Headquarters	Year Founded
MiniMax	Develops powerful, multimodal foundation models and deploys them through popular consumer applications and enterprise APIs globally.	Yan Junjie, Founder & CEO Zhou Yucong, Cofounder Yang Bin, Cofounder	Shanghai, China	2022
Mirage	AI-powered creative studio that allows users to produce high-quality videos with minimal effort and no professional skills required, streamlining entire workflows from filming to final edit	Gaurav Misra, Cofounder & CEO Dwight Churchill, Cofounder	New York, NY	2021
Moescape AI	Provides AI-enabled tools for anime and Vtuber fans to generate digital art, create short videos, and interact with AI companions through immersive role-playing	Thiha Min, Cofounder & CTO Susan Huang, Cofounder & COO	Dover, DE	2021
Moonvalley	Builds production-grade, highly controllable, and legally safe video models specifically for professional filmmakers, studios, and enterprise clients	Naeem Talukdar, Cofounder & CEO	Toronto, Canada	2023
OpusClip	AI-powered video editing company that specializes in content repurposing	Young Zhao, Cofounder & CEO Grace Wang, Cofounder Jay Wu, Cofounder & CTO	Palo Alto, CA	2022
Photoroom	Develops solutions to make complex studio-quality image editing instant and accessible, particularly for small businesses and e-commerce sellers	Matthieu Rouif, Cofounder & CEO Eliot Andres, Cofounder & CTO	Paris, France	2019
Pika	AI-powered video generation platform developed by Pika Labs that enables users to create and edit short video clips from simple text prompts, images, or existing footage	Demi Guo, Cofounder & CEO Chenlin Meng, Cofounder & CTO	Palo Alto, CA	2023
PixAI	Creates high-quality, stylized digital art, particularly within the anime and character design space	Abdul Qaiyoom, Cofounder & CEO Abdul Qabiz, Cofounder & CTO	Singapore	2022
Pixelcut	Provides AI-powered photo editing solutions, primarily catering to the needs of e-commerce sellers, small businesses, and content creators	Dominique Yahyavi, Founder & CEO	Oakland, CA	2021
PixVerse	AI-driven platform that generates short, high-quality videos from text prompts and still images, focusing on speed and cinematic realism	Wang Changhu, Founder & CEO Xie Xuzhang, Cofounder	Singapore	2023
PlantNet	Science platform that utilizes advanced artificial intelligence and machine learning to facilitate plant identification and biodiversity monitoring	Alexis Joly, Lead Researcher Pierre Bonnet, Botanist	Montpellier, France	2009
Remaker	Operates as an all-in-one AI visual content transformation suite.	Alex Zhang, Founder & CEO	Beijing, China	2023
Remini	AI-powered photo and video enhancement application.	Francesco Patarnello, Cofounder Matteo Danieli, Cofounder	Milan, Italy	2019
Reve	Specializes in realism, aesthetic quality, and accurate typography rendering within images	Dheeraj Pandey, Cofounder & CEO Manoj Agarwal, Cofounder	Palo Alto, CA	2023
Riffusion	Specializes in real-time music creation from text prompts	Seth Forsgren, Cofounder & CEO Hayk Martiros, Cofounder & CTO	Menlo Park, CA	2022
Runway	AI-powered creative platform for artists and creators, offering intuitive tools to generate and edit video, images, and other multimedia content without extensive coding	Cristóbal Valenzuela, Cofounder & CEO Anastasis Germanidis, Cofounder & CTO Alejandro Matamala, Cofounder & Chief Design Officer	Brooklyn, NY	2018
SeaArt.ai	Creative studio and community platform that provides an all-in-one suite of AI tools, primarily centered on high-volume, accessible art and media generation	Fei Ma, Founder & CEO	Singapore	2023
Spicychat.ai	AI platform designed for adult roleplay, virtual companionship, and interactive storytelling with highly customized AI characters	Undisclosed	Montreal, Canada	2023
Stability AI	Specializes in developing open-source generative AI models for various media types, including images, video, audio, and 3D content	Prem Akkaraju, CEO Sean Parker, Executive Chairman	London, United Kingdom	2021
Suno	Creates original songs, including instrumentals and vocals, from simple text prompts	Mikey Shulman, Cofounder & CEO Georg Kucsko, Cofounder & CTO	Cambridge, MA	2022
Synthesia	Video generation platform that enables users to create professional, script-based videos using photorealistic digital avatars and synthetic voices, all without needing cameras, microphones, or traditional editing skills	Victor Riparbelli, Cofounder & CEO Steffen Tjerrild, Cofounder & COO	London, United Kingdom	2017
udio	Music platform that allows users to create complete, studio-quality songs, including sophisticated vocals, instrumentation, and lyrics, from simple text prompts	Andrew Sanchez, Cofounder	New York, NY	2023
UpFoto	Leverages deep learning to improve the quality of images, specializing in restoring old or low-resolution photos and generating new, stylized AI portraits	Undisclosed	Zhengzhou, China	2023
Veed	Online video editing platform that leverages AI to simplify and streamline video creation for a broad user base, from social media creators to large enterprises	Sabba Keynejad, Cofounder & CEO Timur Mamedov, Cofounder & CTO	London, United Kingdom	2018
VivaCut	Mobile application for professional-grade, multi-layer video editor for social media creators, integrating AI features to automate and enhance complex video production tasks.	Sheng Han, Cofounder & CEO	Hangzhou, China	2019
Voicemod	Provides real-time AI voice changing and soundboard technology, primarily targeting the gaming, streaming, and content creation communities	Jaime Bosch, Cofounder & CEO Fernando Bosch, Cofounder Juan Bosch, Cofounder	Valencia, Spain	2014
Wink	Provides an all-in-one suite of AI-powered tools for enhancing, retouching, and styling videos and photos, especially for social media content creators	Wu Zeyuan, Founder & CEO of Meitu	Xiamen, China	2021

Sources: Company websites

Exhibit 19
Consumer AI
Developer Tools & Hardware

Company/Product	Overview	Management	Headquarters	Year Founded
Anysphere	A research laboratory and software company that develops the Cursor artificial intelligence code editor to automate software development tasks through a unified interface that integrates multiple large language models	Michael Truell, Cofounder & CEO Aman Sanger, Cofounder & COO Sualeh Asif, Cofounder & CPO	San Francisco, CA	2022
Blackbox AI	Operates an artificial intelligence coding platform that provides a multi-model code editor and autonomous agents designed to assist developers with real-time code completion, bug fixing, and software documentation	Robert Rizk, Cofounder & CEO Roger Rizk, Cofounder & COO	San Francisco, CA	2021
StackBlitz	A software development platform (Bolt) that utilizes in-browser WebContainer technology to allow users to build, edit, and deploy full-stack web applications through natural language interactions	Eric Simons, Cofounder & CEO Albert Pai, Cofounder & CTO	San Francisco, CA	2024
Cobot	Operates a robotics and artificial intelligence laboratory that develops collaborative robots designed to integrate safely into human environments to automate tasks in manufacturing, logistics, and healthcare	Brad Porter, Founder & CEO	Santa Clara, CA	2022
EnCharge AI	Semiconductor design firm that develops analog in-memory computing hardware and software systems to run artificial intelligence models with higher energy efficiency and compute density than traditional digital processors	Naveen Verma, CEO Kailash Gopalakrishnan, CTO Echere Iroaga, COO	Santa Clara, CA	2022
Fauna Robotics	Designs lightweight and cost-effective robotic platforms intended for general-purpose use in unstructured, human-centered environments to facilitate large-scale data collection for embodied artificial intelligence	Rob Cochran, Cofounder & CEO Josh Merel, Cofounder & CTO	New York, NY	2024
Framer	Produces a web-based design and publishing platform that enables teams to build professional, interactive websites without writing code by using a visual canvas integrated with artificial intelligence	Koen Bok, Cofounder & CEO Jorn van Dijk, Cofounder	Amsterdam, Netherlands	2015
Lightricks	Develops a suite of artificial intelligence-powered video and image editing applications alongside open-source generative models to provide creators and businesses with high-fidelity visual storytelling tools	Zeev Farbman, Cofounder & CEO Yaron Inger, Cofounder & CTO	Jerusalem, Israel	2013
Lovable	Artificial intelligence development platform that utilizes vibe coding to convert natural language descriptions into fully functional, production-ready web applications and websites	Anton Osika, Cofounder & CEO Fabian Hedin, Cofounder	Stockholm, Sweden	2023
Matic	Engineers a fully autonomous floor-cleaning robot that utilizes five RGB cameras and neural networks to navigate home environments without the need for LIDAR or constant internet connectivity	Navneet Dalal, Cofounder & CEO Mehul Nariyawala, Cofounder	Mountain View, CA	2017
Mojo Vision	Develops a flexible, wafers-in, wafers-out micro-LED platform for use in AI glasses and optical interconnects	Nikhil Balram, CEO Mike Wiemer, CTO	Saratoga, CA	2015
Physical Intelligence	Develops foundational software for robotics that utilizes large-scale data and general-purpose artificial intelligence models to enable diverse hardware platforms to perform complex physical tasks	Karol Hausman, Cofounder & CEO Sergey Levine, Cofounder & Chief Scientist	San Francisco, CA	2024
Plaud.ai	Produces AI-native wearable devices and software that capture, transcribe, and summarize in-person and digital conversations to automate documentation and information management for professionals	Nathan Xu, Cofounder & CEO	San Francisco, CA	2021
Rabbit	Builds consumer electronics and artificial intelligence software centered on a proprietary "Large Action Model" that allows a handheld device to navigate application interfaces and execute digital tasks on behalf of the user	Jesse Lyu, Cofounder & CEO Dan Ferris, Cofounder & CSO	Santa Monica, CA	2020
Replit	Provides a browser-based integrated development environment and software creation platform that uses autonomous agents to allow users to build, test, and deploy applications using natural language	Amjad Masad, Cofounder & CEO Haya Odeh, Cofounder	San Francisco, CA	2016
Same	Builds and clones websites and web apps from simple text prompts, allowing users to create functional sites with React, APIs, and databases	Aiden Bai, Cofounder Nisarg Patel, Cofounder	San Francisco, CA	202
Sandbar	Manufactures a wearable smart ring and companion software that captures voice notes and utilizes artificial intelligence to organize thoughts into a searchable personal knowledge graph	Mina Fahmi, Cofounder & CEO Kirak Hong, Cofounder & CTO	New York, NY	2021
Skild AI	Develops a robot-agnostic foundation model and software platform that enables diverse hardware embodiments to perform physical tasks like object manipulation, navigation, and locomotion across unstructured environments	Deepak Pathak, Cofounder & CEO Abhinav Gupta, Cofounder & President	Pittsburgh, PA	2023
Turing	Maintains a platform that uses artificial intelligence to source, vet, and match remote software developers and domain experts with enterprises and AI labs to support model training and software engineering tasks	Jonathan Siddharth, Cofounder & CEO Vijay Krishnan, Cofounder & CTO Alok Bhushan, CFO	Palo Alto, CA	2018
Windsurf	Operates an agentic integrated development environment that converts natural language instructions into multistep coding actions by analyzing entire codebases to automate complex software engineering tasks	Varun Mohan, Founder & CTO	San Francisco, CA	2021

Sources: Company websites

Exhibit 20
Consumer AI
General & Platforms

Company/Product	Overview	Management	Headquarters	Year Founded
Adot	Decentralized AI-powered search network that enables users and developers to create personalized search engines by indexing on-chain and off-chain data from across the web	Wei Z., Founder & CEO	San Francisco, CA	2022
Andi	Operates a conversational, generative AI-powered search engine that provides direct answers and visual results while protecting user privacy and filtering out advertising and SEO spam	Angela Hoover, Cofounder & CEO Jed White, Cofounder & CTO	San Francisco, CA	2021
Anthropic	Artificial intelligence safety and research company that develops the Claude family of large language models designed to be reliable, interpretable, and steerable through a constitutional AI framework	Dario Amodei, Cofounder & CEO Daniela Amodei, Cofounder & President Jack Clark, Cofounder & Head of Policy	San Francisco, CA	2021
Celestial AI	Artificial intelligence infrastructure company that develops the Photonic Fabric, an optical interconnect technology platform designed to eliminate data bottlenecks and solve the memory wall challenge by using light-based data transfer for high-performance computing	David Lazovsky, Cofounder & CEO Preet Virk, Cofounder & COO	Santa Clara, CA	2020
Cici	Operates an artificial intelligence companion designed for productivity and emotional intelligence that functions as the international counterpart to the Chinese chatbot Doubao	Liang Rubo, Cofounder & CEO of ByteDance Zhang Yiming, Cofounder of ByteDance	Beijing, China	2023
DeepSeek	Research lab that develops open-source large language models and reasoning engines using an architectural framework designed to reduce training and inference expenses	Liang Wenfeng, Founder & CEO	Hangzhou, China	2023
Delphi	Digital cloning platform that enables users to create a digital version of their mind to scale knowledge and interactions through text, voice, and video	Dara Ladjevardian, Cofounder & CEO Samuel Spelsberg, Cofounder	San Francisco, CA	2022
Doubao	Operates a conversational artificial intelligence application that functions as a multimodal assistant for text generation, image creation, and audio processing within the ByteDance ecosystem	Liang Rubo, Cofounder & CEO of ByteDance Zhang Yiming, Cofounder of ByteDance	Beijing, China	2023
HuggingFace	Community-driven platform that provides a central repository for the machine learning community to build, share, and collaborate on open-source models and datasets	Clément Delangue, Cofounder & CEO Julien Chaumond, Cofounder & CTO Thomas Wolf, Cofounder & CSO	Brooklyn, NY	2016
LMArena	Crowdsourced benchmarking platform that evaluates large language models through anonymous pairwise comparisons to establish a public leaderboard based on human preference	Anastasios Angelopoulos, Cofounder & CEO Wei-Lin Chiang, Cofounder & CTO Ion Stoica, Cofounder & Executive Chairman	San Francisco, CA	2023
Lambda	Provides graphics processing unit cloud services and hardware designed for artificial intelligence training and inference workloads	Stephen Balaban, Cofounder & CEO Michael Balaban, Cofounder & CTO	San Jose, CA	2012
Lila Sciences	Develops a scientific superintelligence platform and autonomous laboratories designed to automate the scientific method across life, chemical, and materials sciences	Geoffrey von Maltzahn, Cofounder & CEO John Kim, Chairman & Co-President	Cambridge, MA	2023
Liner	Operates an artificial intelligence search engine that utilizes a database of human-curated highlights and citations to provide researchers and students with verified information	Luke Jinu Kim, Cofounder & CEO Brian Chanmin Woo, Cofounder & COO	San Francisco, CA	2012
Luzia	Artificial intelligence assistant integrated into messaging platforms like WhatsApp and Telegram that allows users to interact with large language models for tasks such as audio transcription, image generation, and translation	Álvaro Martínez Higes, Cofounder & CEO Javier Andrés, Cofounder Carlos Pérez, Cofounder & CTO	Madrid, Spain	2023
Manus	Develops a general-purpose autonomous agent designed to navigate web browsers and execute complex digital tasks independently	Xiao Hong, Cofounder & CEO Yichao "Peak" Ji, Cofounder & Chief Scientist	Singapore	2022
n8n	Workflow automation platform that utilizes a node-based visual interface to connect various applications and services for the deployment of artificial intelligence agents and automated business processes	Jan Oberhauser, Founder & CEO	Berlin, Germany	2019
Nexthop AI	Networking technology company that develops custom hardware and software infrastructure designed for cloud operators to manage large-scale artificial intelligence clusters	Anshul Sadana, Founder & CEO	Santa Clara, CA	2024
OpenAI	Research and deployment company that develops multimodal generative models and autonomous systems designed to assist with reasoning, creativity, and technical problem-solving	Sam Altman, Cofounder & CEO Greg Brockman, Cofounder & President	San Francisco, CA	2015
Perplexity	Offers an answer engine that synthesizes web content to provide conversational responses with cited sources	Aravind Srinivas, Cofounder & CEO Denis Yarats, Cofounder & CTO Johnny Ho, Cofounder & CSO	San Francisco, CA	2022
Poe (Quora)	Provides a unified interface for users to interact with multiple large language models and allows creators to build and monetize customized AI bots	Adam D'Angelo, Founder & CEO	Mountain View, CA	2022
Qwen3	Operates as a series of open-source large language models developed by Alibaba Group that features a hybrid reasoning engine allowing users to toggle between a step-by-step thinking mode for complex logic and a rapid non-thinking mode for general tasks	Eddie Wu, CEO of Alibaba Group Zhou Jingren, CTO of Alibaba Cloud	Hangzhou, China	2023
Reflection.AI	Research laboratory that develops open-source large language models and autonomous coding agents designed to automate complex software engineering tasks and move toward scientific superintelligence	Misha Laskin, Cofounder & CEO Ioannis Antonoglou, Cofounder & CTO	Brooklyn, NY	2024
Replicate	Cloud-based infrastructure platform that enables software engineers to run, fine-tune, and deploy open-source machine learning models through a standardized application programming interface	Ben Firshman, Cofounder & CEO Andreas Jansson, Cofounder & CTO	San Francisco, CA	2019
Snorkel AI	Data-centric artificial intelligence platform that utilizes programmatic labeling to automate the creation, curation, and evaluation of high-quality training datasets for enterprise machine learning models	Alexander Ratner, Cofounder & CEO Fred Sala, Chief Scientist	Palo Alto, CA	2019
TensorWave	Specialized cloud computing platform that provides bare-metal infrastructure optimized for large-scale artificial intelligence workloads using AMD Instinct graphics processing units	Darrick Horton, Cofounder & CEO Piotr Tomasik, Cofounder & COO Jeff Tatarchuk, Cofounder & CGO	Las Vegas, NV	2023
The Browser Company	Designs and maintains the Arc and Dia web browsers to provide a personalized, AI-integrated interface for internet navigation and task management	Josh Miller, Cofounder & CEO Hursh Agrawal, Cofounder & CTO	New York, NY	2019
Together AI	Operates a research-driven cloud platform that provides developers and enterprises with infrastructure for training, fine-tuning, and deploying open-source generative artificial intelligence models	Vipul Ved Prakash, Cofounder & CEO Ce Zhang, Cofounder & CTO	San Francisco, CA	2022
Tradepost	Digital marketplace platform that allows users to receive instant cash offers for collectibles and luxury goods by uploading screenshots or photos of their items	Tyler Mozeleski, Cofounder & CEO Alexander Perez, Cofounder & CTO	Chicago, IL	2024
xAI	Develops the Grok family of large language models and builds large-scale supercomputing infrastructure designed to accelerate scientific discovery and understand the universe	Elon Musk, Founder & CEO	Palo Alto, CA	2023
You.com	Provides a chat-focused productivity engine and search platform utilizing large language models to deliver cited answers and automate complex research tasks	Richard Socher, Cofounder & CEO Bryan McCann, Cofounder & CTO	Palo Alto, CA	2020

Sources: Company websites

Exhibit 21
Consumer AI
Health & Wellness

Company/Product	Overview	Management	Headquarters	Year Founded
Abridge	Transforms clinical conversations into structured medical notes and billable documentation integrated directly into electronic health record systems	Dr. Shiv Rao, Cofounder & CEO Zack Lipton, Cofounder & CTO	Pittsburgh, PA	2018
Ada Health	AI-powered health platform that utilizes a sophisticated medical knowledge base and reasoning engine to provide users with clinical-grade symptom assessments and personalized care navigation	Daniel Nathrath, Cofounder & CEO Dr. Claire Novorol, Cofounder & CMO Dr. Martin Hirsch, Cofounder & Chief Scientific Officer	Berlin, Germany	2011
Alma	Membership-based platform that provides mental health clinicians with the financial and administrative infrastructure needed to accept insurance and manage private practices	Dr. Harry Ritter, Founder & CEO Chris Davis, CFO David Kaplan, CTO Erin Trimble, COO	Brooklyn, NY	2018
Cal AI	AI-powered nutrition application that allows users to track their caloric and nutrient intake by analyzing photographs of their meals through computer vision	Zach Yadegari, Cofounder & CEO Henry Langmack, Cofounder & CTO Blake W. Anderson, Cofounder Jake Castillo, Cofounder	Roslyn, NY	2024
Curai Health	Operates an AI-driven virtual clinic that utilizes machine learning and natural language processing to provide text-based primary care and clinical decision support	Neal Khosla, Cofounder & CEO	San Francisco, CA	2017
Eight Sleep	Sleep fitness technology company that develops smart mattress systems and AI-powered software designed to optimize recovery through dynamic temperature regulation and biometric monitoring	Matteo Franceschetti, Cofounder & CEO Massimo Andreasi Bassi, Cofounder & CTO Alexandra Zatarain, Cofounder & VP of Brand & Marketing	New York, NY	2014
Finch	Self-care application that utilizes a virtual pet to gamify mental health habits and encourage users to complete daily wellness activities such as journaling, goal setting, and mood tracking	Thomas Budi, Cofounder & CEO Stephanie Yuan, Cofounder	Santa Clara, CA	2021
Granola	AI-powered notepad and meeting assistant designed to enhance productivity by turning raw meeting notes and transcripts into organized, actionable insights	Chris Pedregal, Cofounder & CEO Sam Stephenson, Cofounder	London, UK	2023
Hippocratic AI	Safety-focused generative artificial intelligence platform that develops non-diagnostic, patient-facing healthcare agents designed to address the global shortage of healthcare workers by automating tasks such as chronic care management and post-discharge follow-ups	Munjal Shah, Cofounder & CEO Vishal Parikh, Cofounder & Chief Product Officer Dr. Meenesh Bhimani, Cofounder & Chief Medical Officer	Palo Alto, CA	2022
K Health	Operates an AI-driven virtual healthcare platform that utilizes a clinical reasoning engine and medical database to provide on-demand symptom assessments, personalized treatment plans, and virtual primary care visits	Allon Bloch, Cofounder & CEO Ran Shaul, Cofounder & Chief Product Officer	New York, NY	2016
Ollie	Personalized family health and wellness platform that utilizes generative artificial intelligence to automate meal planning, grocery list creation, and household management tasks	Bill Lennon, Cofounder & CEO Max Fergus, Cofounder Christine Shannon, Cofounder & CMO	San Diego, CA	2023
Oura	Health technology company that develops the Oura Ring, a wearable smart device that utilizes advanced sensors to provide personalized insights into sleep, recovery, and overall physiological health	Tom Hale, CEO Petteri Lahtela, Cofounder Kari Kivelä, Cofounder	Oulu, Finland	2013
Roon	Digital health platform that provides patients and caregivers with medically vetted video content and expert answers to navigate complex medical conditions like ALS, glioblastoma, and menopause	Vikram Bhaskaran, Cofounder & CEO Arun Ranganathan, Cofounder & CTO Dr. Rohan Ramakrishna, Cofounder & Chief Medical Officer	New York, NY	2021
SandboxAQ	Enterprise SaaS platform that develops large quantitative models and AI-native software to solve complex challenges in cybersecurity, drug discovery, and navigation by combining artificial intelligence with quantum physics	Jack Hidary, Founder & CEO Eric Schmidt, Chairman David Silverman, CFO	Palo Alto, CA	2022
Slingshot AI (Ash)	AI research lab that develops Ash, a foundation model specifically for psychology designed to provide personalized and clinically rigorous mental health support through a voice and text interface	Daniel Reid Cahn, Cofounder & CEO Neil Parikh, Cofounder & President Dr. Derrick Hull, R&D Lead	New York, NY	2024
Superpower	Preventive health platform that provides members with comprehensive biomarker testing and a personalized health dashboard to manage longevity and primary care through data-driven insights	Jacob Peters, Cofounder & CEO Max Marchione, Cofounder Kevin Unkrich, Cofounder & CTO	San Francisco, CA	2023
Whoop	Develops a screenless wearable device and an AI-powered coaching platform designed to provide continuous physiological monitoring and personalized health guidance through generative artificial intelligence	Will Ahmed, Founder & CEO Carlos Famadas, CFO Johan Liden, Chief Creative Officer	Boston, MA	2012
Woebot Health	Digital health platform that provides evidence-based, AI-powered mental health support through a relational agent designed to deliver cognitive behavioral therapy and monitor symptoms of anxiety and depression	Dr. Alison Darcy, Founder & CEO	San Francisco, CA	2017

Sources: Company websites

Exhibit 22
Consumer AI
Industry-Specific (Marketing, Legal, Defense)

Company/Product	Overview	Management	Headquarters	Year Founded
Bluefish	Enables brands to monitor, optimize, and influence how they are portrayed across various artificial intelligence assistants and search models	Alex Sherman, Cofounder & CEO Jing Feng, Cofounder & COO Andrei Dunca, Cofounder & CTO	New York, NY	2024
Eden AI	Operates an AI-as-a-service aggregator platform that provides a single API for developers to access and compare multiple artificial intelligence models for tasks such as text analysis, translation, and computer vision	Taha Zemmouri, Cofounder & CEO Samy Melaine, Cofounder & CTO	Lyon, France	2021
Eudia	Augmented intelligence platform designed for *Fortune* 500 legal departments that combines AI agents with human expertise to automate complex workflows and digitize institutional legal knowledge	Omar Haroun, Cofounder & CEO Ashish Agrawal, Cofounder & CTO David Van Reyk, Cofounder & COO	Palo Alto, CA	2023
Harvey	Provides law firms and professional service providers with customized large language models for document analysis, legal research, and litigation support	Winston Weinberg, Cofounder & CEO Gabriel Pereyra, Cofounder & President Siva Gurumurthy, CTO	San Francisco, CA	2022
Koah	Provides a native advertising network specifically designed for conversational AI and chat interfaces that utilizes an SDK to provide contextually relevant monetization for third-party developers	Nic Baird, Cofounder & CEO Herrick Fang, Cofounder Mike Choi, Cofounder	San Francisco, CA	2024
Profound	Analytics and optimization platform that helps enterprise brands monitor and improve their visibility within AI search engines and large language models	James Cadwallader, Cofounder & CEO Dylan Babbs, Cofounder & CTO	New York, NY	2024
Shield AI	Defense technology platform that develops the Hivemind AI pilot to enable autonomous flight and mission execution for drones and fighter jets in GPS-denied and communication-degraded environments	Gary Steele, CEO Ryan Tseng, Cofounder & President	San Diego, CA	2015

Sources: Company websites

Exhibit 23
Consumer AI
Learning & Education

Company/Product	Overview	Management	Headquarters	Year Founded
Atypical	Provides personalized tutoring and classroom management tools that adapt to individual student learning styles	Bethanie Drake-Maples, Founder & CEO Kevin Wiggen, CTO Nate Edwards, Chief Design Officer	Austin, TX	2023
Brainly	Peer-to-peer learning platform that uses AI to help students and parents find verified answers to homework questions and academic problems	Michał Borkowski, Cofounder & CEO Bill Salak, CTO & COO	New York, NY	2009
Class Companion	Enables teachers to provide students with instant, personalized feedback on written assignments	Avery Pan, Cofounder & CEO Jack Forbes, Cofounder	San Francisco, CA	2023
ELSA Speak	Operates an artificial intelligence platform that provides personalized English pronunciation and speech training through a mobile application	Vu Van, Cofounder & CEO	San Francisco, CA	2015
Gauth	Provides step-by-step solutions for STEM and social science subjects through a "snap-to-solve" mobile interface	Gao Hao, Director Wu Shiqi, Director	Singapore	2020
Gizmo	Operates an AI-powered learning platform that uses gamification and cognitive science principles to transform study materials into interactive quizzes and flashcards	Petros Christodoulou, Cofounder & CEO Paul Evangelou, Cofounder Robin Jack, Cofounder	London, UK	2019
Hi Translate	Cross-application translation platform that provides real-time text, voice, and image translation to facilitate communication across different mobile apps	Li Yang, Cofounder & CEO	Hong Kong	2018
Magic School	Provides educators with specialized tools for lesson planning, grading, and individualized education program development	Adeel Khan, Founder & CEO	Denver, CO	2023
MasterClass	Streaming and interactive learning platform that combines expert-led video courses with AI-powered coaching tools and roleplay features to provide personalized professional development	David Rogier, Cofounder & CEO Mark Williamson, COO Len Small, Chief Content Officer	San Francisco, CA	2012
Qanda	AI-driven learning platform that utilizes proprietary OCR technology and a specialized math LLM to provide students with instant solutions and personalized tutoring	Jongheun (Ray) Lee, Cofounder & CEO Yongjae (Jake) Lee, Cofounder & CEO	Seoul, South Korea	2015
Speak	Provides users with a virtual tutor for real-time conversational practice and instant feedback on pronunciation and fluency	Connor Zwick, Cofounder & CEO Andrew Hsu, Cofounder & CTO	San Francisco, CA	2016
StudyX	AI learning assistant that integrates multiple language models and a massive community database to provide instant homework solutions, step-by-step explanations, and writing support	Alex Muz, Founder	Middletown, DE	2023
Super Teacher	Tutoring application that provides voice-interactive and adaptive lessons across various subjects for elementary-aged children	Tim Novikoff, Founder & CEO Krzysztof Kulewski, Cofounder	New York, NY	2022

Sources: Company websites

Exhibit 24
Consumer AI
Productivity & Office

Company/Product	Overview	Management	Headquarters	Year Founded
AGI, Inc.	AI-focused productivity start-up developing AGI-0, a proactive virtual coworker designed to execute complex tasks across computers and smartphones by navigating web interfaces and mobile applications autonomously	Div Garg, Cofounder & CEO Naman Garg, Cofounder	San Francisco, CA	2024
Cleo	AI-powered financial assistant tailored for Gen Z and millennials that uses a conversational interface to help users track spending, build credit, and manage budgets through a mix of humor and proactive alerts	Barney Hussey-Yeo, Founder & CEO Aleksandra Wozniak, Cofounder Srivalli Nistala, Cofounder Chitra Akileswaran, Cofounder	London, UK	2016
Cluely	Desktop assistant that provides real-time information and suggested responses through an undetectable screen overlay during virtual meetings, interviews, and sales calls	Chungin Roy Lee, Cofounder & CEO Neel Shanmugam, Cofounder & COO Alex Chen, Cofounder & CTO	San Francisco, CA	2025
DoNotPay	Legal services platform often described as the world's first robot lawyer that automates tasks such as contesting parking tickets, canceling subscriptions, and filing small claims lawsuits	Joshua Browder, Founder & CEO Andrew Kim, VP of Product & AI Tai Nguyen, Head of Engineering	New York, NY	2015
Duckbill	Personal life management platform that combines generative AI with human copilots to automate and execute everyday administrative tasks such as booking appointments, managing household repairs, and researching travel options	Meghan Joyce, Cofounder & CEO Emilia Gonzalez, Cofounder	Boston, MA, USA	2022
Fireflies.ai	Meeting intelligence platform that automates recording, transcription, and summarization across all major video conferencing apps	Krish Ramineni, Cofounder & CEO Sam Udotong, Cofounder & CTO	Pleasanton, CA	2016
Genspark	AI-powered agentic workspace that automates end-to-end business tasks such as creating financial models, board presentations, and full-stack applications	Eric Jing, Cofounder & CEO Kay Zhu, Cofounder & CTO Wen Sang, Cofounder & COO	Palo Alto, CA	2023
Glean	Provides enterprise-grade search and generative AI assistance by connecting a company's entire internal knowledge base across hundreds of SaaS applications to help employees find information and automate workflows	Arvind Jain, Cofounder & CEO Vishwanath T R, Cofounder & CTO Tony Gentilcore, Cofounder & Engineering Lead	Palo Alto, CA	2019
Monarch	Subscription-based personal finance platform that provides a unified dashboard for tracking net worth, spending, and investments without advertisements	Val Agostino, Cofounder & CEO Jon Sutherland, Cofounder & Head of Design Ozzie Osman, Cofounder & Head of Engineering	Covina, CA	2018
Moonshot AI	Leading artificial intelligence lab in China specializing in large language models with ultra-long context windows. The company is widely recognized for its Kimi chatbot and the recently released Kimi K2 Thinking model, which uses a 1-trillion-parameter architecture to solve complex analytical and multistep reasoning tasks.	Yang Zhilin, Founder & CEO Zhou Xinyu, Cofounder Wu Yuxin, Cofounder	Beijing, China	2023
Ohai.ai	AI-powered household assistant designed to manage family logistics by automating school calendar syncing, coordinating caregiver schedules, and extracting tasks from emails or images via a conversational interface	Sheila Lirio Marcelo, Founder & CEO Kevin Yang, Cofounder Lauren Tornow, Cofounder	New York, NY	2021
Otter.ai	Meeting intelligence platform that provides real-time transcription, automated summaries, and action items for virtual and in-person meetings	Sam Liang, Cofounder & CEO Yun Fu, Cofounder & CTO Darius Contractor, Chief Growth Officer	Mountain View, CA	2016
Prosper	Healthcare-focused artificial intelligence platform that develops voice agents to automate administrative phone tasks such as patient scheduling, insurance verification, and claims follow-up	Xavier de Gracia, Cofounder Josep Mingot, Cofounder	New York, NY	2023
QuillBot	Operates a writing platform that offers tools for paraphrasing, grammar checking, and summarization to help users improve the clarity and fluency of their writing	Andrew Grauer, CEO Rohan Gupta, Cofounder & Chief Strategist Stephen Van Horne, COO & CFO	Chicago, IL, USA	2017
Speechify	Text-to-speech platform that uses artificial intelligence to convert written text into high-quality audio across various devices and formats	Cliff Weitzman, Founder & CEO Tyler Weitzman, Cofounder & Head of AI	Miami, FL	2017
Superhuman	AI-native productivity platform that provides a high-speed email experience through integrated tasks, communication tools, and workflow automations	Rahul Vohra, Founder & CEO Vivek Sodera, Cofounder Conrad Irwin, Cofounder & CTO	San Francisco, CA	2014
TurboScribe	Transcription service that converts audio and video files into text using GPU-accelerated speech-to-text technology	Leif Foged, Founder	United States	2023

Sources: Company websites

Exhibit 25
Consumer AI
Social & Companions

Company/Product	Overview	Management	Headquarters	Year Founded
Bible Chat	Christian companion platform that offers personalized Bible study, spiritual counseling, and prayer guidance by grounding its conversational AI in scripture and theological research	Laurențiu Bălașa, Cofounder & CEO Marius Iordache, Cofounder & CTO	Galati, Romania	2019 (Pivoted to Bible Chat in 2023)
Boardy.ai	Operates an AI "super-connector" platform that uses a conversational voice agent to interview professionals about their goals and projects, subsequently facilitating high-signal, double-opt-in introductions within a global network of founders and investors	Andrew D'Souza, Cofounder & CEO	San Francisco, CA	2024
Candy AI	Provides a generative AI platform focused on virtual companionship and romantic roleplay, allowing users to interact with customizable AI characters through text and voice, with support for high-fidelity image generation and adaptive personalities	Alexis Soulopoulos, Cofounder	Santa Vanera, Malta	2023
Character.ai	Conversational AI platform for creating and interacting with personalized AI characters for entertainment, roleplay, and productivity	Karandeep Anand, CEO Dominic Perella, Chief Legal Officer	Menlo Park, CA	2021
Civitai	Provides an open-source platform and repository for sharing and discovering custom AI art models, particularly those built on Stable Diffusion architecture	Justin Maier, Cofounder & CEO Maxfield Hulker, Cofounder	Boise, ID	2022
Co-Star	AI-powered social astrology app that uses NASA satellite data and natural language generation to provide personalized horoscopes and compatibility insights	Banu Guler, Cofounder & CEO Anna Kopp, Cofounder Ben Weitzman, Cofounder	New York, NY	2017
DittoAI	AI-powered matching platform that eliminates swiping and texting by using conversational agents to vet users and facilitate immediate in-person dates	Allen Wang, Cofounder & CEO Eric Liu, Cofounder	San Francisco, CA	2024
Gigi	AI-powered professional networking platform that analyzes calendar data and real-world interactions to map social capital and facilitate high-signal career introductions	Clara Gold, Founder & CEO	San Francisco, CA	2023
Hallow	Catholic prayer and meditation app that offers audio-guided sessions, including the Rosary, daily Gospel readings, and sleep stories, while using AI for search tools and personalized content recommendations	Alex Jones, Cofounder & CEO Erich Kerekes, Cofounder & CTO Alessandro DiSanto, Cofounder	Chicago, IL	2018
Keeper	Matchmaking platform that uses deep psychometric data and relationship science to automate the search for long-term partners, specifically targeting marriage-minded users	Jake Kozloski, Founder & CEO	New York, NY	2022
Kindroid	Provides a consumer-focused AI companion platform that allows users to create highly personalized virtual beings with persistent long-term memory, customizable backstories, and multimodal features like real-time voice calls and generated selfies	Jerry Meng, Founder & CEO	Los Angeles, CA	2023
Mivi	Manufactures consumer electronics and audio products, recently launching a proprietary conversational AI platform integrated into wearable devices to enable human-like voice interactions, personalized coaching, and contextual awareness	Viswanadh Kandula, Cofounder & CEO Midhula Devabhaktuni, Cofounder & CMO	Hyderabad, India	2025
OurDream AI	Provides a generative AI playground for adults, allowing users to create and interact with customizable virtual companions through uncensored text chat, voice calls, and high-fidelity image and video generation	Not specified	USA	2023
Replika	Allows users to create personalized 3D avatars for emotional support, roleplay, and digital friendship	Dmytro Klochko, CEO Eugenia Kuyda, Founder	San Francisco, CA	2013
Series	Operates an AI-powered social and professional network that uses conversational agents to facilitate merit-based introductions and match users for dating or hiring	Nathaneo Johnson, Cofounder & CEO Sean Hargrow, Cofounder & COO Jack Roberts, Cofounder & CTO	New York, NY	2023
Sesame	Develops lifelike conversational AI companions and "fashion-forward" smart glasses designed to replace traditional screens with natural voice interaction	Brendan Iribe, Cofounder & CEO Ankit Kumar, Cofounder & CTO Nate Mitchell, Chief Product Officer	San Francisco, CA	2023
Stitch	Operates as an experimental AI-powered design tool that generates full user interface layouts and production-ready front-end code from text prompts or uploaded sketches. Stitch is a Google Labs project.	Vincent Nallatamby, Product Manager Arnaud Benard, Research Scientist Sam El-Husseini, Software Engineer	Mountain View, CA	2025
Tolan	Provides an AI-powered "embodied companion" in the form of a Pixar-like alien character designed to offer emotional support, help with daily tasks, and reduce social overwhelm for Gen Z users	Quinten Farmer, Cofounder & CEO Evan Goldschmidt, Cofounder & CTO Ajay Mehta, Cofounder & President	San Francisco, CA	2023
Viggle	Provides a controllable video generation platform that uses generative AI to animate static images or mix character movements from one video into another with high consistency	Hang Chu, Founder & CEO	Toronto, Canada	2022
Weights	Operates a consumer AI platform for creative media synthesis where users train and share custom voice models, create AI song covers, and generate images or chat characters	Not publicly disclosed	an Francisco, CA, US	2024

Sources: Company websites

Exhibit 26
Consumer AI
Talent & Hiring

Company/Product	Overview	Management	Headquarters	Year Founded
Alex	Staffing platform that utilizes voice and video AI agents to autonomously conduct interviews, screen resumes, and manage follow-up scheduling to streamline the recruitment process	Aaron Wang, Cofounder & CEO John Rytel, Cofounder & CTO	San Francisco, CA	2023
Handshake	Career network for students and graduates that has expanded to include a high-skill AI data labeling business, leveraging its network of millions of Ph.D.s and subject-matter experts to provide human feedback and domain-specific training data for frontier AI labs	Garrett Lord, Cofounder & CEO Jonathan Stull, President & COO Sahil Bhaiwala, Chief Strategy & Innovation Officer (Handshake AI)	San Francisco, CA	2014 (Handshake AI launched in 2025)
Lightscreen	Interview platform featuring a voice and video agent that conducts technical screenings and collaborative pair-programming sessions to assess candidate skills while using multi-modal cheating detection	Prachie Banthia, Cofounder & CEO Gavin Saldanha, Cofounder & CTO	Sunnyvale, CA	2024
Mercor	Recruitment and talent platform that automates the hiring process by using AI agents to conduct interviews and screen resumes, while also providing high-skilled human experts to AI labs for data annotation and model fine-tuning	Brendan Foody, Cofounder & CEO Adarsh Hiremath, Cofounder & CTO Surya Midha, Cofounder & Chairman	San Francisco, CA	2023
Wayfaster	Recruitment platform that utilizes voice and video agents to conduct automated screening interviews, generate real-time evaluations, and integrate with existing applicant tracking systems to accelerate hiring for retail and BPO sectors	Varun Khurana, Cofounder	Monmouth Junction, NJ	2022

Sources: Company websites

Exhibit 27
Consumer AI
Travel & Lifestyle

Company/Product	Overview	Management	Headquarters	Year Founded
Alta	Styling application that creates a digital closet for users to generate personalized outfit recommendations, virtual try-ons, and shopping suggestions based on their existing wardrobe, lifestyle, and local weather	Jenny Wang, Founder	New York, NY	2023
Doji	Fashion platform that enables users to create realistic, full-body digital avatars to virtually try on clothing and explore personalized styles	Dorian Dargan, Cofounder & CEO Jim Winkens, Cofounder & CTO	New York, NY	2024
OneOff	Fashion discovery platform that uses computer vision and large language models to identify clothing worn by celebrities and influencers, allowing users to shop exact or similar items through a unified checkout	Bobby Maylack, Cofounder Emir Talu, Cofounder	Miami Beach, FL	2024
Layla	Provides a personal styling and digital wardrobe platform that allows users to catalog their existing clothes, receive AI-generated outfit combinations, and discover new pieces that complement their current collection	Saad Saeed, CEO Sardar Bali, Cofounder & CPO	Berlin, Germany	2024
Mindtrip	Develops an AI-native travel platform that combines conversational AI with a proprietary knowledge base to help users discover, plan, and book entire personalized itineraries in a single interface	Andy Moss, Cofounder & CEO Garrick Toubassi, Cofounder Trey Matteson, Cofounder	San Francisco, CA	2023
Wanderboat	Provides an AI-powered local discovery and trip planning platform that analyzes social media content and videos to offer personalized recommendations for activities, dining, and itineraries	You Wu, Cofounder & CEO Xiaochuan Ni, Cofounder Ying Yang, Head of Growth	Sunnyvale, CA	2023
Wanderlog	Trip planning application that allows users to build detailed itineraries, organize travel documents, and map out routes with real-time syncing across devices	Harry Yu, Cofounder & CEO Peter Xu, Cofounder	San Francisco, CA	2019

Sources: Company websites

The prices of the common stock of other public companies mentioned in this report follow:

Alphabet, Inc. (Outperform)	$322.16
Apple Inc.	$246.70
The Bank of New York Mellon Corporation	$120.68
Cloudflare, Inc. (Outperform)	$177.42
Dick's Sporting Goods, Inc.	$206.25
Duolingo, Inc. (Outperform)	$148.58
Meta Platforms, Inc. (Outperform)	$604.12
Microsoft Corporation (Outperform)	$454.52
NerdWallet, Inc. (Outperform)	$12.81
Netflix, Inc. (Outperform)	$87.26
NVIDIA Corporation (Outperform)	$178.07
Pattern Group Inc. (Outperform)	$13.46
Spotify Technology SA	$511.33
Tesla, Inc. (Market Perform)	$419.25
Walmart Inc.	$118.71
Warby Parker, Inc. (Outperform)	$26.79
ZipRecruiter, Inc. (Outperform)	$2.56

IMPORTANT DISCLOSURES

One or more of the research analysts responsible for covering Airbnb, Inc., Alphabet, Inc., Meta Platforms, Inc. and Uber Technologies, Inc. mentioned in this report owns shares of the security.

William Blair or an affiliate was a manager or co-manager of a public offering of equity securities for Pattern Group Inc. within the prior 12 months.

William Blair or an affiliate is a market maker in the security of 1stdibs.com, Inc., Airbnb, Inc., Alphabet, Inc., AppLovin Corp., CarGurus Inc., Dolby Laboratories, Inc., DoorDash Inc., Duolingo, Inc., EverQuote, Inc., Meta Platforms, Inc., NerdWallet, Inc., Netflix, Inc., Pattern Group Inc., Roku, Inc., The Trade Desk, Inc., Uber Technologies, Inc., Vivid Seats Inc. and ZipRecruiter, Inc.

William Blair or an affiliate expects to receive or intends to seek compensation for investment banking services from 1stdibs.com, Inc., Airbnb, Inc., Alphabet, Inc., AppLovin Corp., CarGurus Inc., Dolby Laboratories, Inc., DoorDash Inc., Duolingo, Inc., EverQuote, Inc., Meta Platforms, Inc., NerdWallet, Inc., Netflix, Inc., Pattern Group Inc., Roku, Inc., The Trade Desk, Inc., Uber Technologies, Inc., Vivid Seats Inc. and ZipRecruiter, Inc. or an affiliate within the next three months.

William Blair or an affiliate received compensation for investment banking services or non-investment-banking services from Pattern Group Inc. within the last 12 months. Pattern Group Inc. is or was, within the last 12 months, an investment banking client of William Blair & Company and/or one or more of its affiliates.

Officers and employees of William Blair or its affiliates (other than research analysts) may have a financial interest in the securities of 1stdibs.com, Inc., Airbnb, Inc., Alphabet, Inc., AppLovin Corp., CarGurus Inc., Dolby Laboratories, Inc., DoorDash Inc., Duolingo, Inc., EverQuote, Inc., Meta Platforms, Inc., NerdWallet, Inc., Netflix, Inc., Pattern Group Inc., Roku, Inc., The Trade Desk, Inc., Uber Technologies, Inc., Vivid Seats Inc. and ZipRecruiter, Inc.

This report is available in electronic form to registered users via R*Docs™ at https://williamblairlibrary.bluematrix.com or www.williamblair.com.

Please contact us at +1 312 236 1600 or consult https://www.williamblair.com/equity-research/coverage for all disclosures.

Ralph Schackart attests that 1) all of the views expressed in this research report accurately reflect his/her personal views about any and all of the securities and companies covered by this report, and 2) no part of his/her compensation was, is, or will be related, directly or indirectly, to the specific recommendations or views expressed by him/her in this report. We seek to update our research as appropriate. Other than certain periodical industry reports, the majority of reports are published at irregular intervals as deemed appropriate by the research analyst.

DOW JONES: 48488.60
S&P 500: 6796.86
NASDAQ: 22954.30

Additional information is available upon request.

Current Rating Distribution (as of January 22, 2026):

Coverage Universe	Percent	Inv. Banking Relationships *	Percent
Outperform (Buy)	72	Outperform (Buy)	11
Market Perform (Hold)	28	Market Perform (Hold)	3
Underperform (Sell)	1	Underperform (Sell)	0

*Percentage of companies in each rating category that are investment banking clients, defined as companies for which William Blair has received compensation for investment banking services within the past 12 months.

The compensation of the research analyst is based on a variety of factors, including performance of his or her stock recommendations; contributions to all of the firm's departments, including asset management, corporate finance, institutional sales, and retail brokerage; firm profitability; and competitive factors.

OTHER IMPORTANT DISCLOSURES

Stock ratings and valuation methodologies: William Blair & Company, L.L.C. uses a three-point system to rate stocks. Individual ratings reflect the expected performance of the stock relative to the broader market (generally the S&P 500, unless otherwise indicated) over the next 12 months. The assessment of expected performance is a function of near-, intermediate-, and long-term company fundamentals, industry outlook, confidence in earnings estimates, valuation (and our valuation methodology), and other factors. Outperform (O) - stock expected to outperform the broader market over the next 12 months; Market Perform (M) - stock expected to perform approximately in line with the broader market over the next 12 months; Underperform (U) - stock expected to underperform the broader market over the next 12 months; not rated (NR) - the stock is not currently rated. The valuation methodologies include (but are not limited to) price-to-earnings multiple (P/E), relative P/E (compared with the relevant market), P/E-to-growth-rate (PEG) ratio, market capitalization/revenue multiple, enterprise value/EBITDA ratio, discounted cash flow, and others. Stock ratings and valuation methodologies should not be used or relied upon as investment advice. Past performance is not necessarily a guide to future performance.

The ratings and valuation methodologies reflect the opinion of the individual analyst and are subject to change at any time.

Our salespeople, traders, and other professionals may provide oral or written market commentary, short-term trade ideas, or trading strategies-to our clients, prospective clients, and our trading desks-that are contrary to opinions expressed in this research report. Certain outstanding research reports may contain discussions or investment opinions relating to securities, financial instruments and/or issuers that are no longer current. Investing in securities involves risks. This report does not contain all the material information necessary for an investment decision. Always refer to the most recent report on a company or issuer. Our asset management and trading desks may make investment decisions that are inconsistent with recommendations or views expressed in this report. We will from time to time have long or short positions in, act as principal in, and buy or sell the securities referred to in this report. Our research is disseminated primarily electronically, and in some instances in printed form. Research is simultaneously available to all clients. This research report is for our clients only. No part of this material may be copied or duplicated in any form by any means or redistributed without the prior written consent of William Blair & Company, L.L.C.

This is not in any sense an offer or solicitation for the purchase or sale of a security or financial instrument.

The factual statements herein have been taken from sources we believe to be reliable, but such statements are made without any representation as to accuracy or completeness or otherwise, except with respect to any disclosures relative to William Blair or its research analysts. Opinions expressed are our own unless otherwise stated and are subject to change without notice. Prices shown are approximate.

This report or any portion hereof may not be copied, reprinted, sold, or redistributed or disclosed by the recipient to any third party, by content scraping or extraction, automated processing, or any other form or means, without the prior written consent of William Blair. Any unauthorized use is prohibited.

If the recipient received this research report pursuant to terms of service for, or a contract with William Blair for, the provision of research services for a separate fee, and in connection with the delivery of such research services we may be deemed to be acting as an investment adviser, then such investment adviser status relates, if at all, only to the recipient with whom we have contracted directly and does not extend beyond the delivery of this report (unless otherwise agreed specifically in writing). If such recipient uses these research services in connection with the sale or purchase of a security referred to herein, William Blair may act as principal for our own account or as riskless principal or agent for another party. William Blair is and continues to act solely as a broker-dealer in connection with the execution of any transactions, including transactions in any securities referred to herein.

For important disclosures, please visit our website at williamblair.com.

This material is distributed in the United Kingdom and the European Economic Area (EEA) by William Blair International, Ltd., authorised and regulated by the Financial Conduct Authority (FCA). William Blair International, Limited is a limited liability company registered in England and Wales with company number 03619027. This material is only directed and issued to persons regarded as Professional investors or equivalent in their home jurisdiction, or persons falling within articles 19 (5), 38, 47, and 49 of the Financial Services and Markets Act of 2000 (Financial Promotion) Order 2005 (all such persons being referred to as "relevant persons"). This document must not be acted on or relied on by persons who are not "relevant persons."

This report is being furnished in Brazil on a confidential basis and is addressed to the addressee personally, and for its sole benefit. This does not constitute an offer or solicitation for the purchase or sale of a security by any means that would constitute a public offering in Brazil under the regulations of the Brazilian Securities and Exchange Commission (*Comissão de Valores Mobiliários*) or an unauthorized distribution under Brazilian laws and regulations. The securities are authorized for trading on non-Brazilian securities markets, and this report and all the information herein is intended solely for professional investors (as defined by the applicable Brazilian regulation) who may only acquire these securities through a non-Brazilian account, with settlement outside Brazil in a non-Brazilian currency.

William Blair & Company, L.L.C. licenses and applies the SASB Materiality Map® and SICSTM in our work.

William Blair

Equity Research Directory

John Kreger, Partner Director of Research +1 312 364 8612
Kyle Harris, CFA, Partner Operations Manager +1 312 364 8230

Scott Hansen, Partner Associate Director of Research +1 212 245 6526

CONSUMER

Sharon Zackfia, CFA, Partner +1 312 364 5386
Group Head–Consumer
Lifestyle and Leisure Brands, Restaurants, Automotive/E-commerce

Jon Andersen, CFA, Partner +1 312 364 8697
Consumer Products

Phillip Blee, CPA +1 312 801 7874
Home and Outdoor, Automotive Parts and Services, Discount and Convenience

Dylan Carden +1 312 801 7857
E-commerce, Specialty Retail

ECONOMICS

Richard de Chazal, CFA +44 20 7868 4489

ENERGY AND POWER TECHNOLOGIES

Jed Dorsheimer +1 617 235 7555
Group Head–Energy and Power Technologies
Generation, Efficiency, Storage

Neal Dingmann +1 312 801 7835
Oil and Gas, Rare Earth

Tim Mulrooney, Partner +1 312 364 8123
Energy and Environmental Services

FINANCIAL SERVICES AND TECHNOLOGY

Adam Klauber, CFA, Partner +1 312 364 8232
Group Head–Financial Services and Technology
Financial Analytic Service Providers, Insurance Brokers, Property & Casualty Insurance

Andrew W. Jeffrey, CFA +1 415 796 6896
Fintech

Cristopher Kennedy, CFA +1 312 364 8596
Fintech, Specialty Finance

Jeff Schmitt +1 312 364 8106
Wealthtech, Wealth Management, Capital Markets Technology

GLOBAL SERVICES

Tim Mulrooney, Partner +1 312 364 8123
Group Head–Global Services
Commercial and Residential Services

Andrew Nicholas, CPA +1 312 364 8689
Consulting, HR Technology, Information Services

Trevor Romeo, CFA +1 312 801 7854
Staffing, Waste and Recycling

HEALTHCARE

Biotechnology

Matt Phipps, Ph.D., Partner +1 312 364 8602
Group Head–Biotechnology

Sami Corwin, Ph.D. +1 312 801 7783

Lachlan Hanbury-Brown +1 312 364 8125

Andy T. Hsieh, Ph.D., Partner +1 312 364 5051

Myles R. Minter, Ph.D., Partner +1 617 235 7534

Healthcare Technology and Services

Ryan S. Daniels, CFA, Partner +1 312 364 8418
Group Head–Healthcare Technology and Services
Healthcare Technology, Healthcare Services

Brandon Vazquez, CFA +1 212 237 2776
Dental, Animal Health, Medical Technology

Life Sciences

Matt Larew, Partner +1 312 801 7795
Life Science Tools, Bioprocessing, Healthcare Delivery

Andrew F. Brackmann, CFA +1 312 364 8776
Diagnostics

Max Smock, CFA +1 312 364 8336
Pharmaceutical Outsourcing and Services

INDUSTRIALS

Brian Drab, CFA, Partner +1 312 364 8280
Co-Group Head–Industrials
Advanced Manufacturing, Industrial Technology

Ryan Merkel, CFA , Partner +1 312 364 8603
Co-Group Head–Industrials
Building Products, Specialty Distribution

Louie DiPalma, CFA +1 312 364 5437
Aerospace and Defense, Smart Infrastructure

Ross Sparenblek +1 312 364 8361
Diversified Industrials, Robotics, and Automation

TECHNOLOGY, MEDIA, AND COMMUNICATIONS

Jason Ader, CFA, Partner +1 617 235 7519
Co-Group Head–Technology, Media, and Communications
Infrastructure Software

Arjun Bhatia, Partner +1 312 364 5696
Co-Group Head–Technology, Media, and Communications
Software

Dylan Becker, CFA +1 312 364 8938
Software

Louie DiPalma, CFA +1 312 364 5437
Government Technology

Jonathan Ho, Partner +1 312 364 8276
Cybersecurity, Security Technology

Sebastien Naji +1 212 245 6508
Infrastructure Software, Semiconductor and Infrastructure Systems

Maggie Nolan, CPA, Partner +1 312 364 5090
IT Services

Jake Roberge +1 312 364 8056
Software

Ralph Schackart III, CFA, Partner +1 312 364 8753
Internet and Digital Media

Stephen Sheldon, CFA, CPA, Partner +1 312 364 5167
Vertical Technology – Real Estate, Education, Restaurant/Hospitality

EDITORIAL AND SUPERVISORY ANALYSTS

Steve Goldsmith, Head Editor and SA +1 312 364 8540
Katie Anderson, Editor and SA +44 20 7868 4451
Audrey Majors, Editor and SA +1 312 364 8992
Beth Pekol Porto, Editor and SA +1 312 364 8924
Lisa Zurcher, Editor and SA +44 20 7868 4549