


AI ADOPTION AND RISK REPORT

 cyberhaven labs

Q2 2025





AI ADOPTION AND RISK REPORT

Table of contents

-
- 01 Key findings
 - 02 Growth in AI usage
 - 03 The top 25 AI tools in the workplace
 - 04 AI adoption by industry
 - 05 Risk assessment of AI tools
 - 06 The top 10 AI tools by risk level
 - 07 DeepSeek usage by end users
 - 08 What corporate data workers put into AI tools
 - 09 How AI-generated content is used at work
 - 10 How employees of different seniority use AI at work
 - 11 How developers use AI to write code
 - 12 DeepSeek model usage in development projects
-

Introduction

Since ChatGPT launched in 2022, AI adoption has become one of the fastest growing workplace technologies in history. What began as individual employees experimenting with generative AI has shifted to these tools becoming embedded in core business processes across organizations of all sizes. The speed of this transition from novelty to necessity is unprecedented.

As AI becomes deeply integrated into critical business operations and adopted by increasing numbers of departments and employees, the volume and sensitivity of data flowing into these systems has grown exponentially. Companies now face a dual challenge: harnessing AI's potential while managing the substantial data risks it introduces.

The majority of current AI usage falls under what's called "shadow AI" – the use of AI tools unsanctioned by corporate IT departments. This grassroots adoption is predominantly driven by younger, mid-level employees and individual contributors, while adoption continues to lag among more senior employees who tend to be more cautious about new technologies.

THIS COMPREHENSIVE ANALYSIS FROM CYBERHAVEN LABS DRAWS ON ACTUAL AI USAGE PATTERNS OF 7 MILLION WORKERS, PROVIDING AN UNPRECEDENTED VIEW INTO THE ADOPTION PATTERNS AND SECURITY IMPLICATIONS OF AI IN THE CORPORATE ENVIRONMENT.

For forward-thinking organizations, a significant opportunity exists in understanding and leveraging this grassroots AI usage. By identifying how employees are successfully using AI to increase efficiency and drive business results, companies can strategically implement these tools and methodologies on a broader scale, capturing their benefits enterprise-wide.

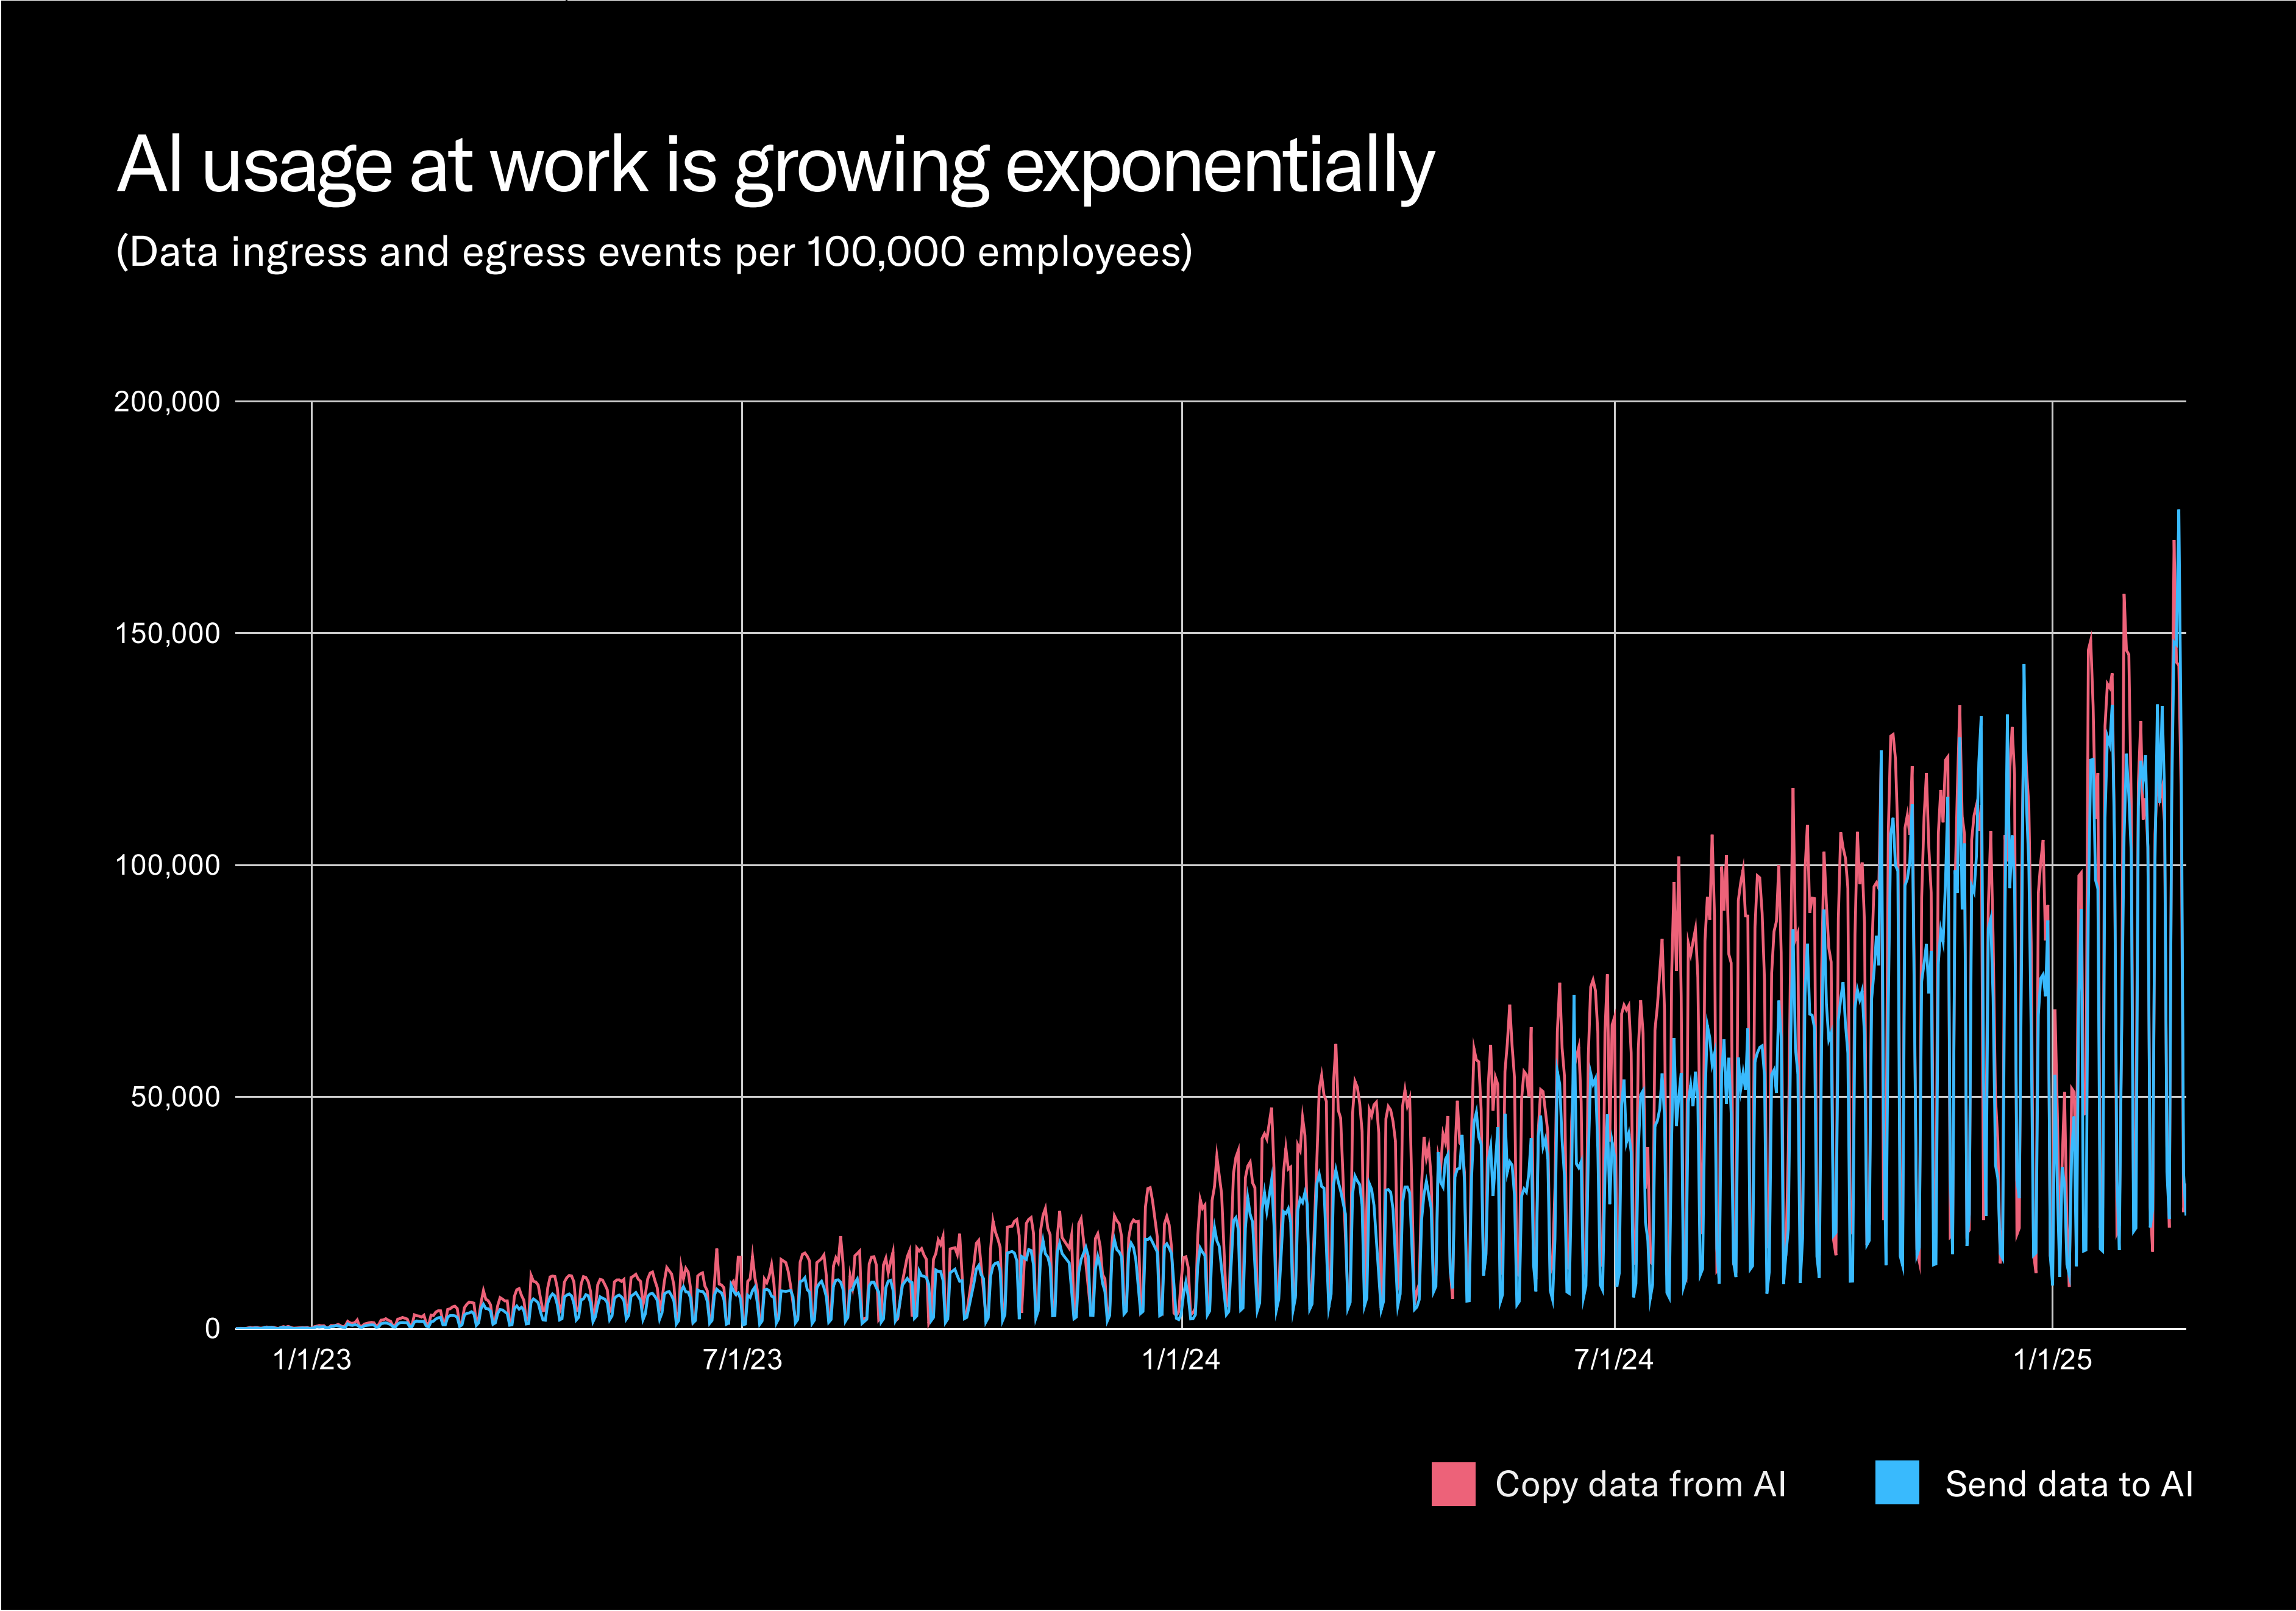
However, the risks to corporate data cannot be overlooked. Many AI tools incorporate user-provided data into their training models, potentially exposing sensitive information. This characteristic, among other risk factors, indicates that the majority of AI tools currently used in workplaces present significant data security risks. As organizations enable AI adoption, they must also implement robust guardrails to protect their most sensitive information assets.

Key Findings

01	AI USAGE IN THE WORKPLACE IS GROWING EXPONENTIALLY	<p>The rate of AI usage frequency at work grew 4.6x in the past 12 months and an astonishing 61x over the past 24 months.</p> <p>AI adoption is highest among employees at tech companies (38.9% of employees). Employees using AI spread fastest in manufacturing (20x growth) and retail firms (24x growth).</p>
02	MOST AI TOOLS IN USE IN THE WORKPLACE ARE HIGH RISK	<p>71.7% of AI tools are high or critical risk, with 39.5% of AI tools inadvertently exposing user interaction/training data and 34.4% exposing user data.</p> <p>83.8% of enterprise data going to AI is going to risky AI tools, instead of enterprise-ready tools (low and very low risk).</p>
03	DEEPSEEK USAGE SURGED AND THEN PLATEAUED	<p>End user use of DeepSeek through deepseek.com surged in the 3 weeks after the R1 release, but soon plateaued ending the first 7 weeks at 672.8% growth.</p> <p>Use of DeepSeek in local application development by AI engineers peaked at 17.7% of AI developer usage in February 2025, second to Llama.</p>
04	AN INCREASING PERCENTAGE OF CORPORATE DATA GOING TO AI IS SENSITIVE	<p>34.8% of corporate data employees put into AI tools is sensitive, up from 27.4% a year ago and 10.7% two years ago.</p> <p>The most common types of sensitive data employees put into AI are source code (18.7% of sensitive data), R&D materials (17.1%) and Sales and Marketing data (10.7%).</p>
05	SOFTWARE ENGINEERS ARE LEVERAGING AI CODING TOOLS MORE THAN EVER	<p>When companies roll out AI developer tools like Cursor or Cline, usage grows by 400% in the first four months after rollout.</p> <p>At the same time, when companies roll out AI developer tools usage of traditional IDEs like VS Code, Xcode, PyCharm declines by 23.7%.</p>
06	AI ADOPTION IS HIGHEST AMONG YOUNGER, MID-LEVEL EMPLOYEES	<p>Mid-level employees like analysts and specialists use AI tools 3.5 times as much as the next-nearest cohort (manager-level employees).</p> <p>Among software engineers, the highest AI usage is among mid-level software engineers, who use AI 189% more than their more junior counterparts.</p>

Growth in AI usage

AI usage at work continues to grow at a remarkable pace. Growth in the past 12 months has been 4.6x and over the past 24 months AI usage has grown an astounding 61x. This represents one of the fastest rates of adoption for any new technology in the workplace, substantially outpacing even SaaS adoption, which took years to achieve similar penetration levels.



The periodic troughs visible in our usage data chart represent lower activity by workers on Saturday and Sunday using their work computers. Notably, AI usage on work machines during weekends is now comparable to the average weekday usage observed at the beginning of 2024, highlighting the rapid growth of AI tools for work.

Growth in the amount of data employees input into AI tools has outpaced the growth in data they extract from these systems. This trend suggests that AI is becoming more central to business processes, and as we'll explore in later sections, an increasing percentage of this data contains sensitive information.

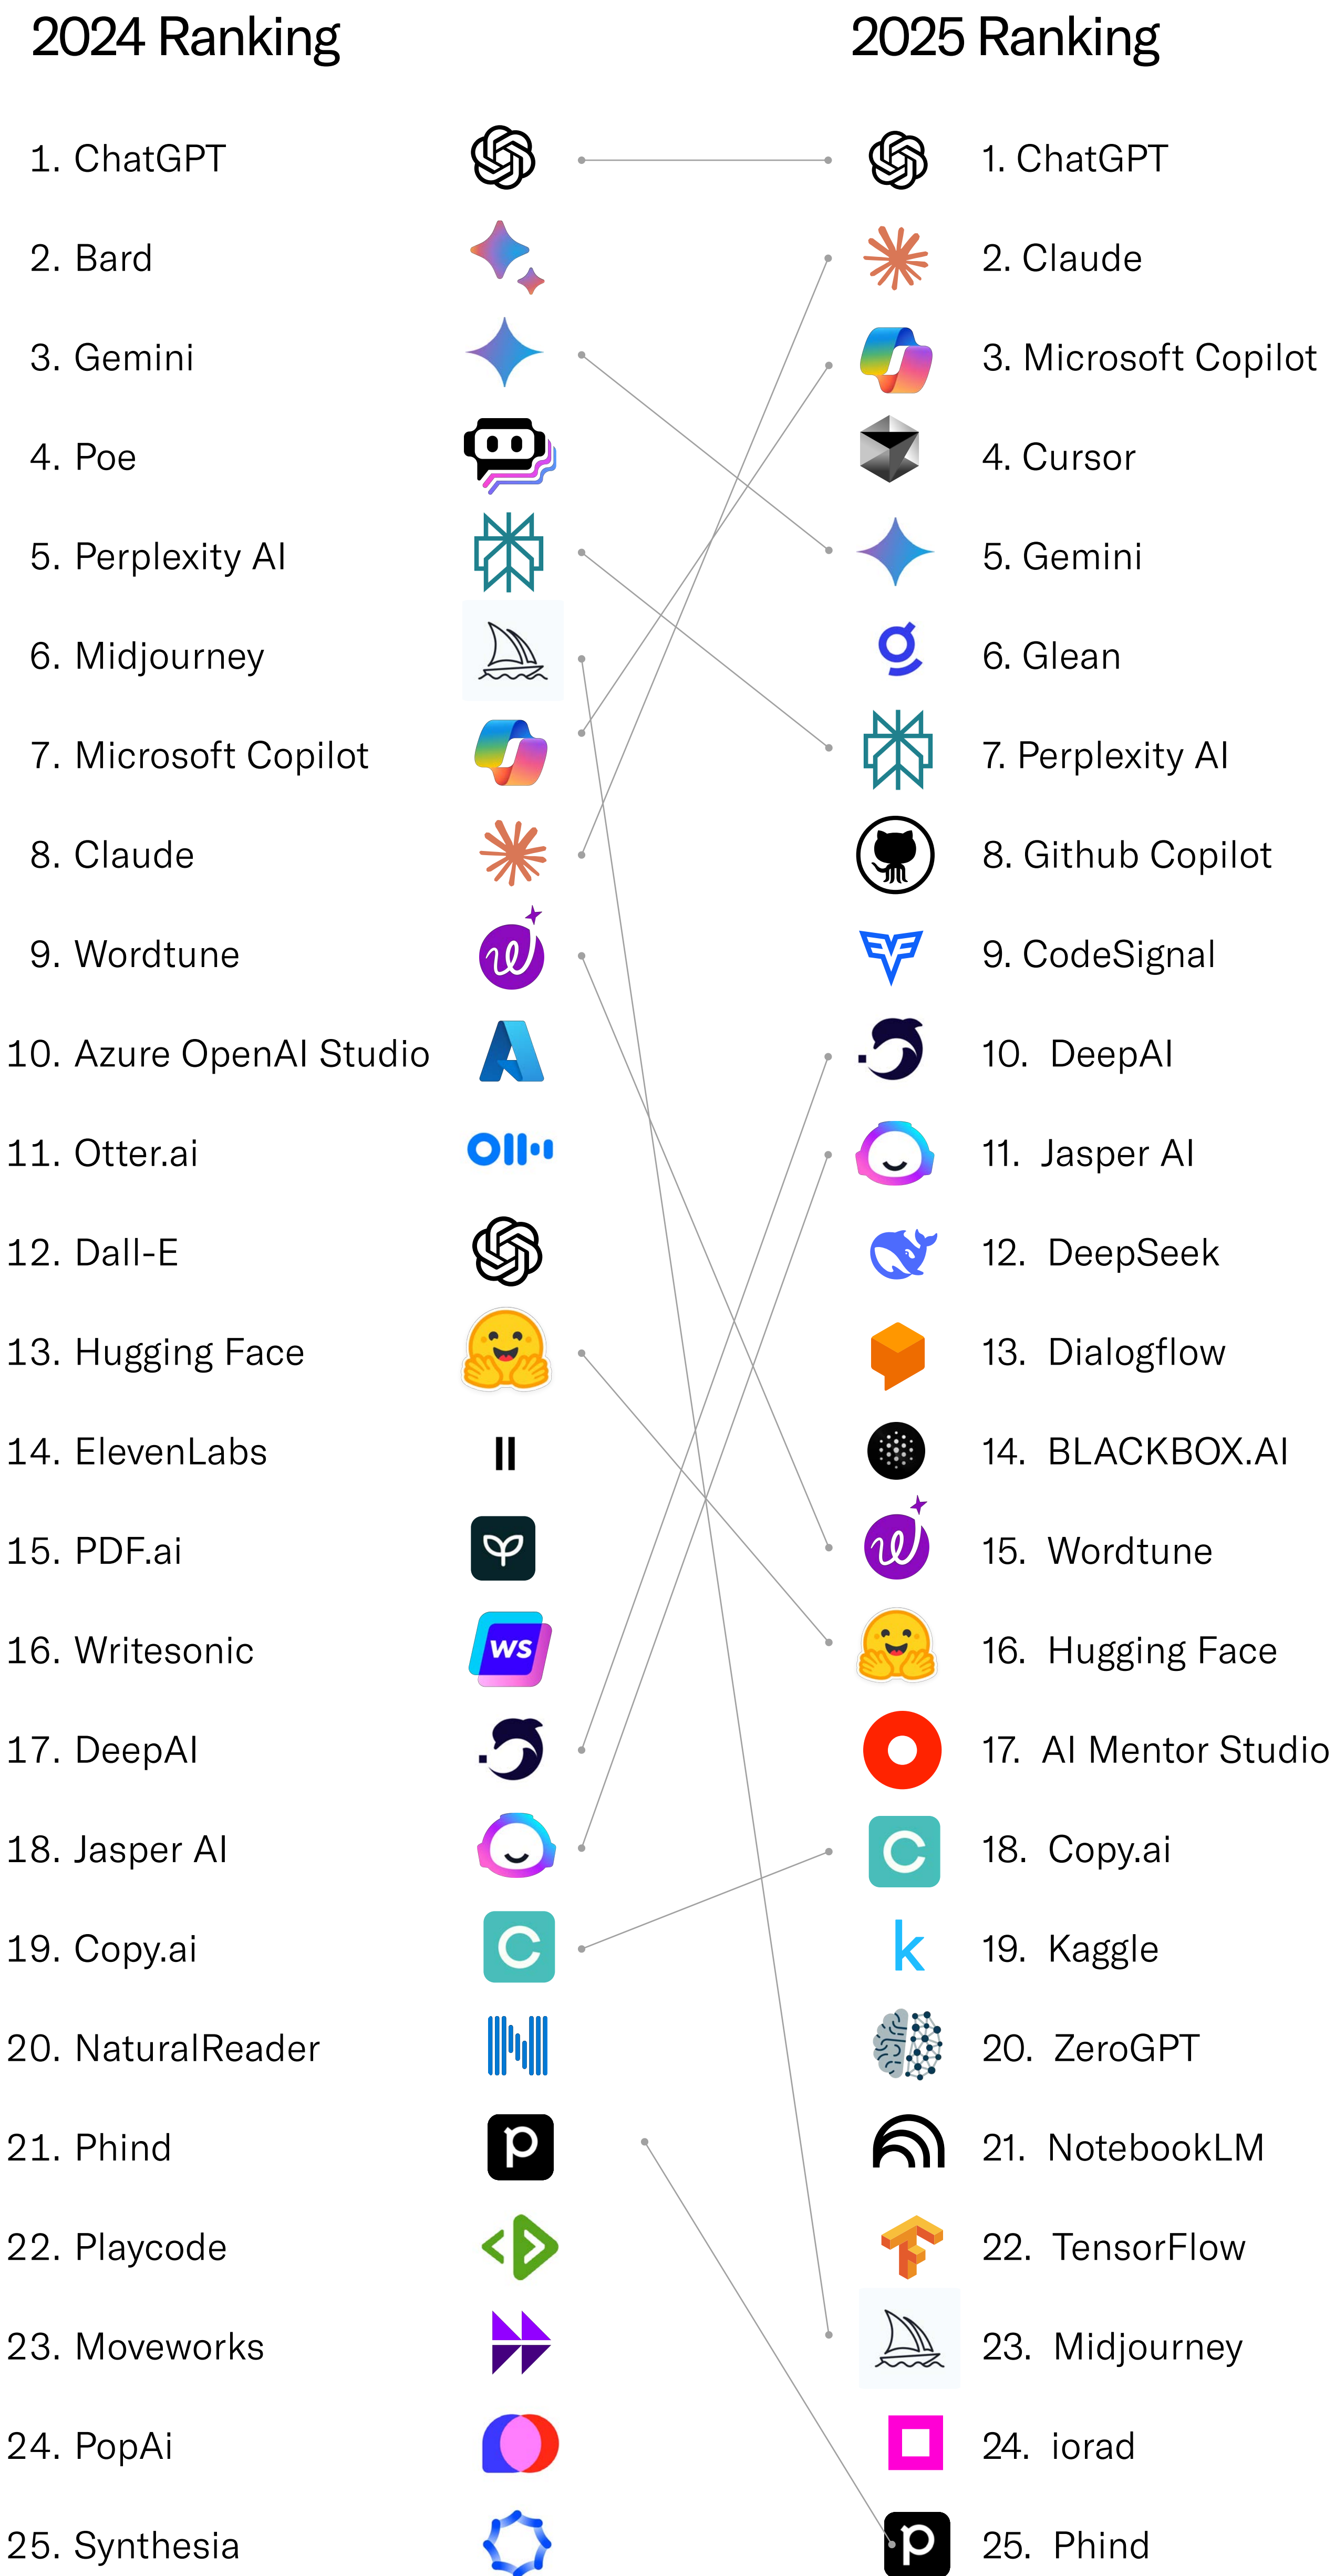
The top 25 AI tools in the workplace

Cyberhaven continuously tracks usage metrics for more than 700 leading AI platforms. Our analysis of the past three months reveals significant shifts in the AI landscape compared to last year's rankings, highlighting both established leaders and emerging challengers.

ChatGPT maintains its dominant position as the most-used AI tool in workplace environments, but the competition is intensifying. Claude has made a dramatic climb from eighth position to second place, while Microsoft Copilot has risen from seventh to third. These shifts reflect the market's consolidation around tools offering enterprise-grade capabilities and security features.

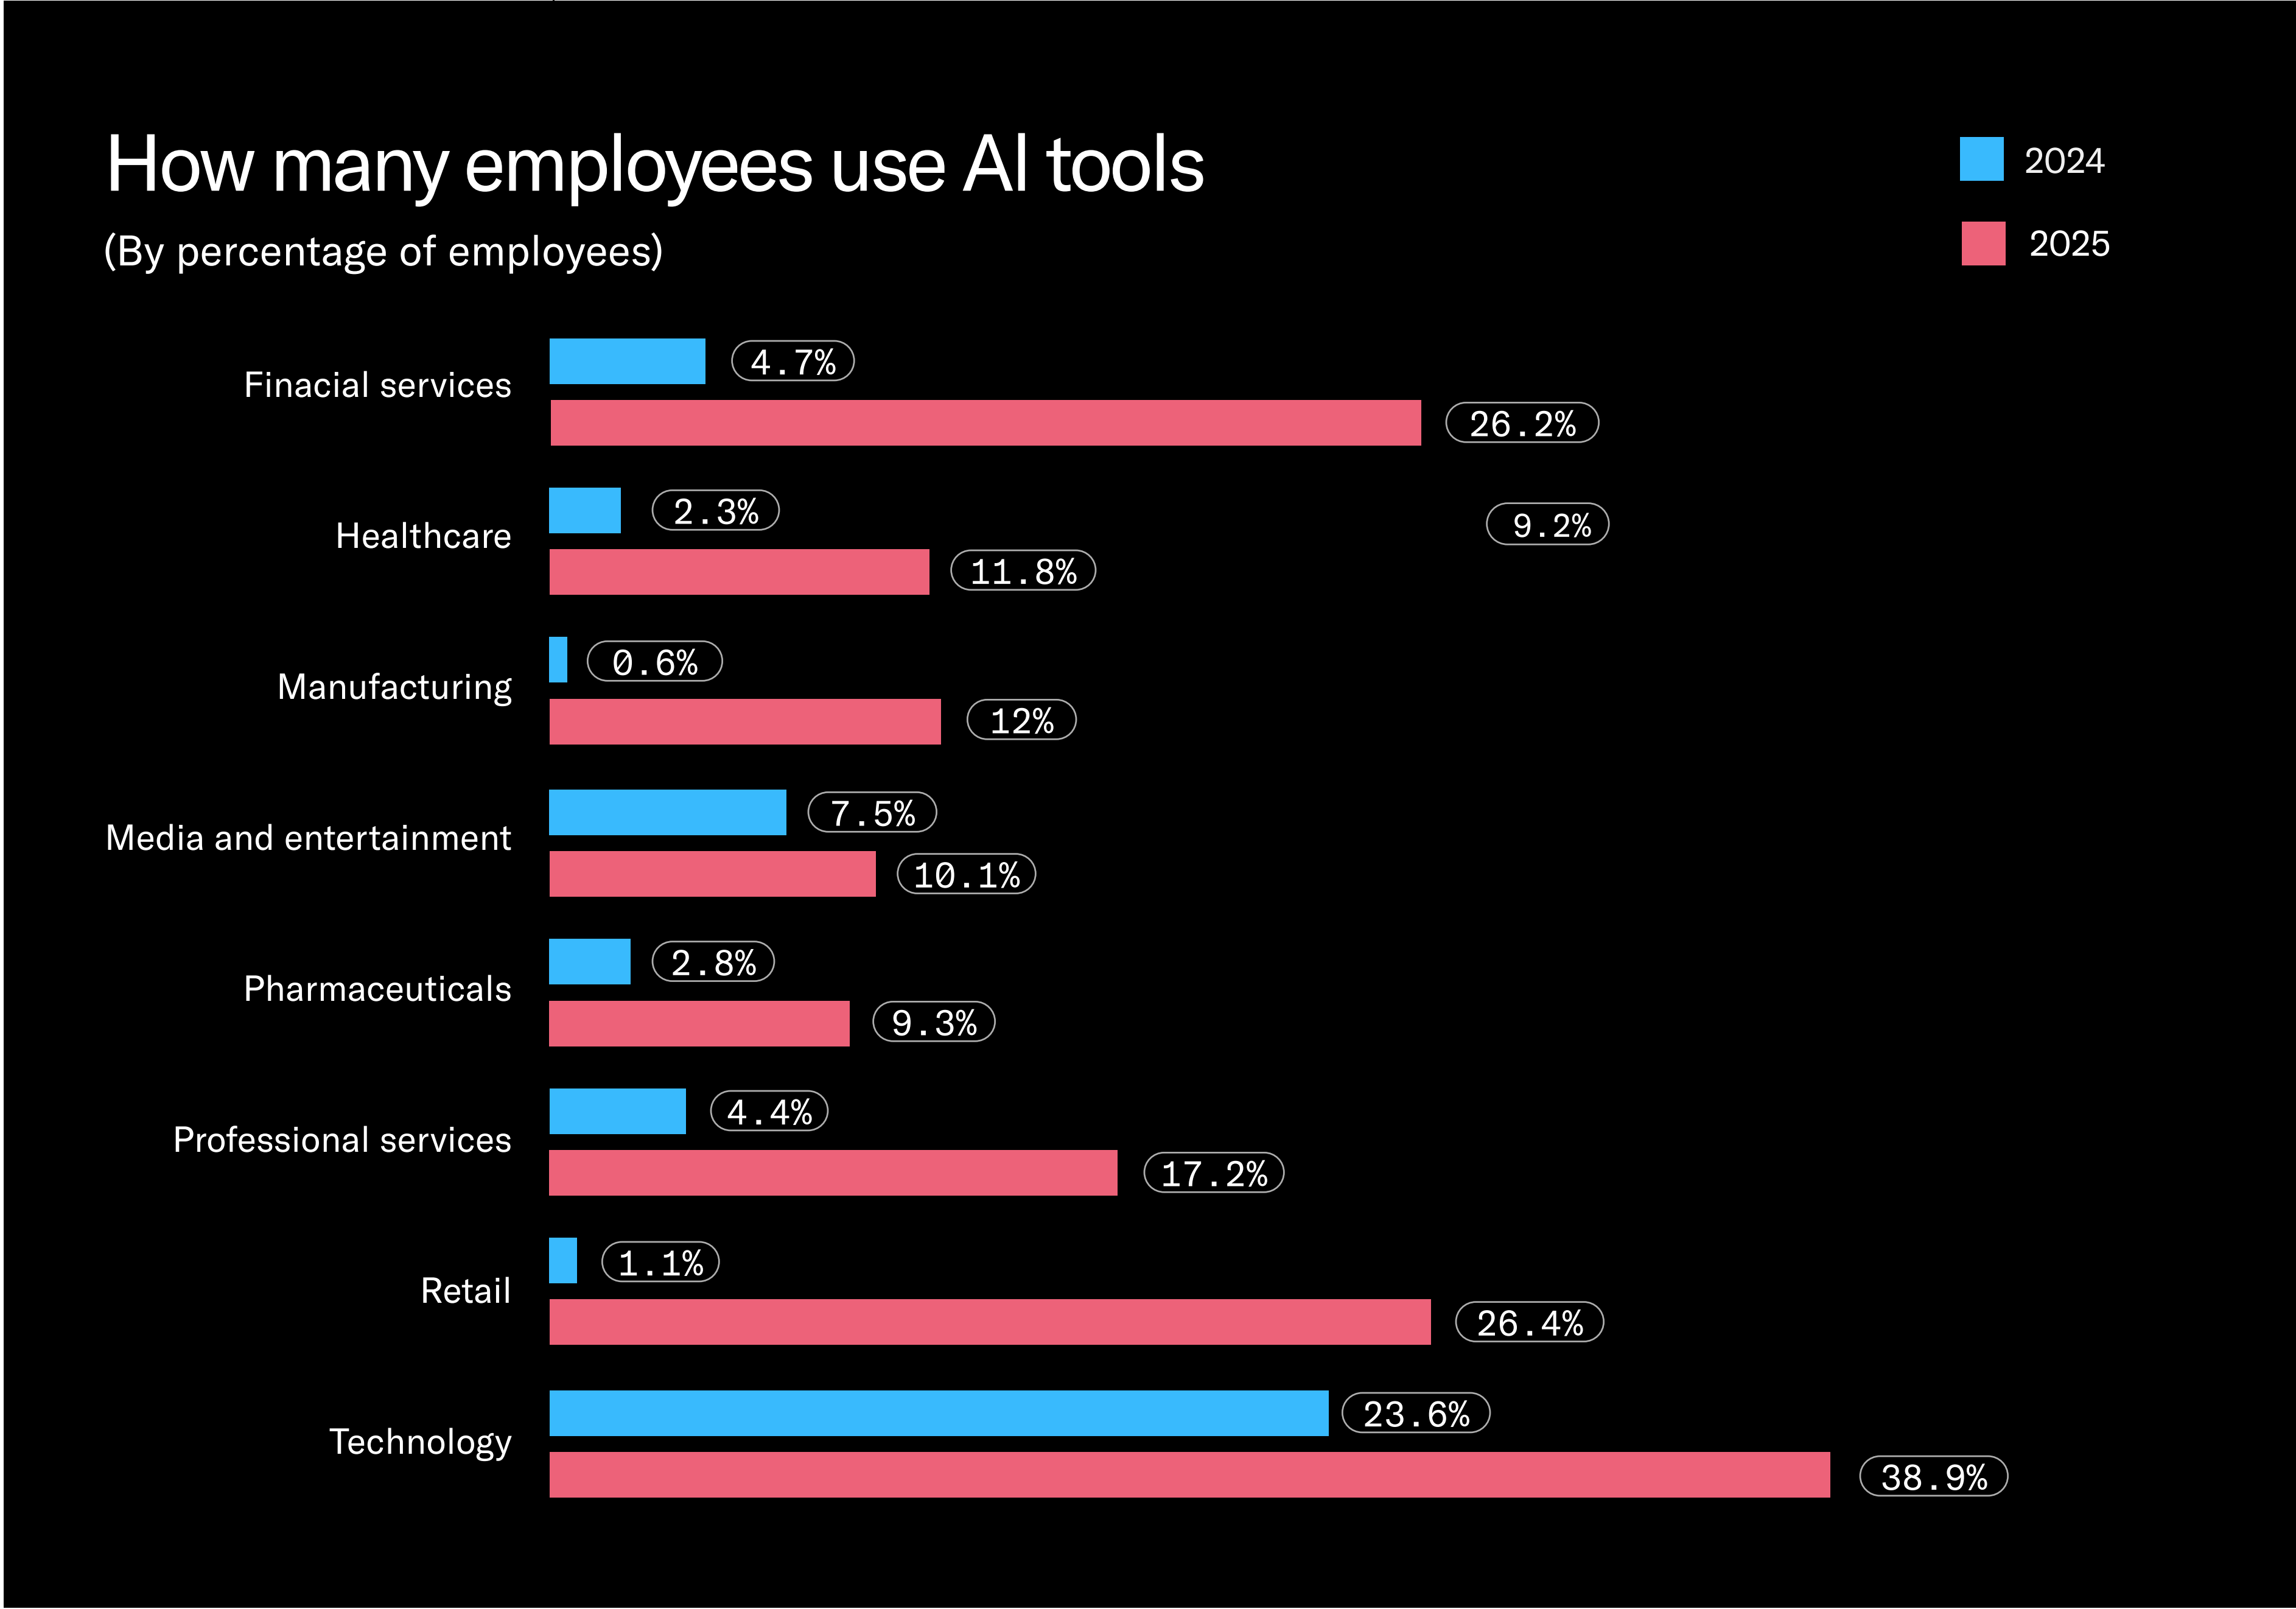
The competitive landscape continues to evolve rapidly, with some previously prominent tools falling from favor. Bard, which ranked second last year, has disappeared entirely from our top 25 list following Google's strategic pivot to Gemini. Similarly, Midjourney has slipped from sixth position to twenty-third as its relative growth has been outpaced by more business-focused tools.

Perhaps most indicative of emerging trends is the strong debut of specialized tools like Cursor and Glean in our rankings. Cursor's appearance at fourth position – notably ahead of Github Copilot at eighth – signals a significant shift in developer preferences for AI-assisted coding platforms. The rapid adoption of these specialized tools suggests the AI market is entering a new phase of specialization after the initial dominance of general-purpose platforms.



AI adoption by industry

The past year has witnessed AI adoption extending deeper across all industry sectors, with notable growth in previously underrepresented verticals. Manufacturing and retail organizations, which had been AI adoption laggards, experienced the most dramatic growth rates, with manufacturing companies seeing 20x growth in employee AI adoption and retail firms achieving an even more impressive 24x increase.



This rapid expansion has reshaped the industry adoption landscape. Retail organizations, which previously showed minimal AI penetration, have surged to second place with 26.4% of employees now regularly using AI tools. This positions them behind only technology companies, where 38.9% of employees are active AI users.

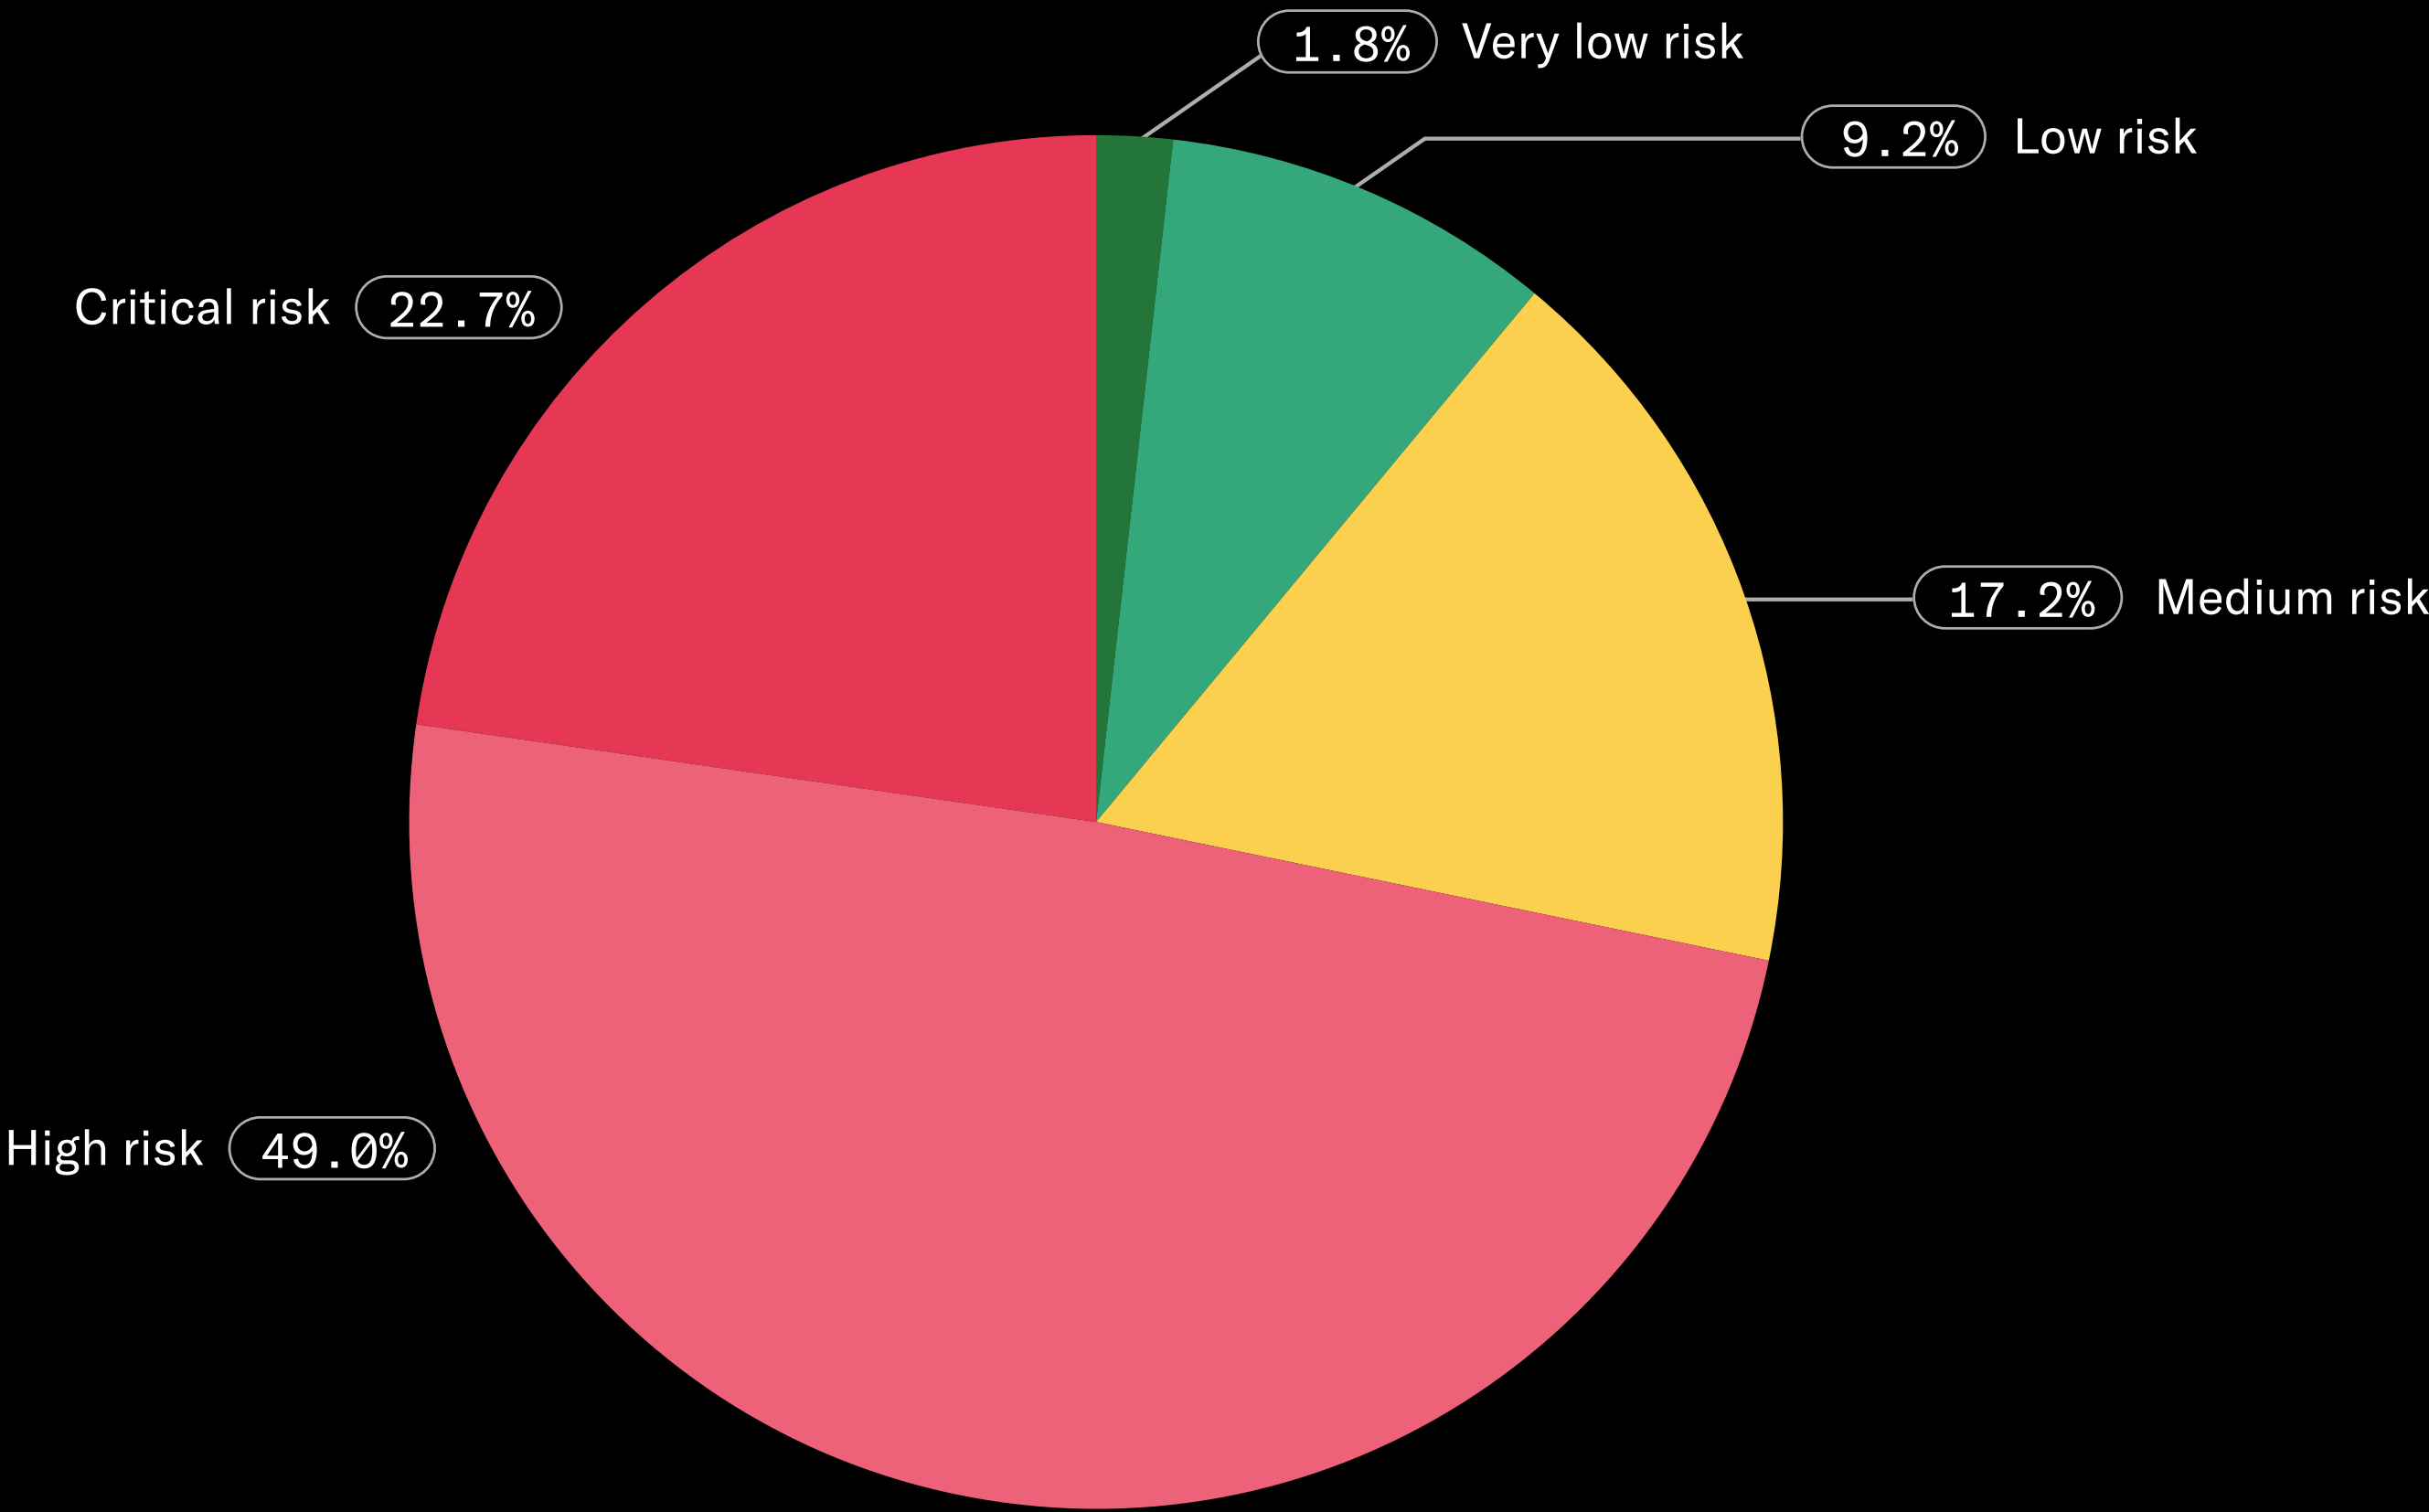
The data suggests AI adoption is following a pattern similar to previous technology waves – beginning in tech-forward sectors before rapidly diffusing to more traditional industries once clear use cases and value propositions emerge. What distinguishes the AI revolution, however, is the unprecedented speed of this cross-sector adoption.

Risk assessment of AI tools

Cyberhaven's comprehensive risk assessment methodology evaluates over 700 leading AI tools, categorizing them according to their security and data privacy characteristics. Our findings reveal substantial concerns about the current AI ecosystem, with 71.7% of tools falling into high or critical risk categories. Particularly concerning is that just 11% qualify for our low or very low risk classifications, indicating a significant security gap in the available tools.

Understanding what makes an AI tool risky provides insight into these statistics. Some key risk factors we identified include inadvertent exposure of user interactions and training data (present in 39.5% of AI tools) and user data being accessible to third parties without adequate controls (found in 34.4% of tools). These vulnerabilities create substantial data exfiltration risks that organizations must address as AI adoption accelerates.

The majority of AI apps are risky
(Numbers of apps by Cyberhaven AI Risk IQ level)













Our analysis of the past three months ranks the ten most-used AI tools within each risk category, providing organizations with actionable intelligence for security planning and tool selection.

In the low and very low risk categories, Claude leads the rankings, followed by Microsoft Copilot. For this analysis, we've separated ChatGPT's personal and enterprise editions due to their different security profiles – ChatGPT Enterprise ranks third among low-risk tools with its enhanced data handling guarantees and customizable data retention policies.











The top 10 AI tools by risk level.

At the other end of the spectrum, DeepAI tops our critical risk category by usage, followed closely by DeepSeek. The prevalence of high-usage tools in our critical risk category underscores the security challenges organizations face as employees gravitate toward tools that may prioritize functionality and accessibility over enterprise-grade security controls.











Low and very low risk

-  Claude
-  Microsoft Copilot
-  ChatGPT Enterprise
-  Glean
-  Dialogflow
-  Hugging Face
-  NotebookLM
-  Synthesia
-  Murf AI
-  Google Recorder











Medium risk

-  ChatGPT Personal
-  Gemini
-  Perplexity AI
-  Jasper AI
-  Copy.ai
-  Kaggle
-  TensorFlow
-  Otter.ai
-  Horizon3.ai
-  ElevenLabs

High risk

-  CodeSignal
-  BLACKBOX.AI
-  Wordtune
-  AI Mentor Studio
-  Midjourney
-  iorad
-  Phind
-  TurboScribe
-  Leonardo AI
-  NaturalReader

Critical risk

-  DeepAI
-  DeepSeek
-  ZeroGPT
-  FreeTTS
-  Playcode
-  Toolsaday
-  TTSMaker
-  ChatGLM
-  EmojiCopy
-  Undetectable AI


DeepSeek usage by end users

The January 2025 release of DeepSeek's R1 model generated significant attention, primarily due to its advanced capabilities but also because of security and privacy concerns. DeepSeek's data collection practices encompass personal information, keystrokes, and device data – all of which are transmitted and stored in China.

To provide clarity, our analysis distinguishes between the website version (deepseek.com) and the open-source version that developers can deploy locally. While we examine the latter separately in this report, the consumer-facing website version has shown remarkable adoption.

DeepSeek usage grew quickly after R1 but growth plateaued

(Events per 100,000 users by week)



End user engagement with DeepSeek through its web interface surged dramatically in the initial three weeks post-release, but this growth wasn't sustained. Usage plateaued by the end of the first seven weeks, settling at 672.8% growth relative to pre-release baselines.

For context, the growth of DeepSeek following the R1 release substantially outpaces that of other major AI model releases. Gemini usage increased 171.9% in the seven weeks following its 2.0 release, while Claude usage rose 136.1% after version 3.5 launched. DeepSeek's dramatically higher growth rate, even with its subsequent plateau, suggests substantial market interest in emerging models regardless of their security profile.

What corporate data workers put into AI tools

As AI moves from experimental to operational status within organizations, we're witnessing a concerning trend in the sensitivity of data being processed by these systems. Currently, 34.8% of all corporate data that employees input into AI tools is classified as sensitive – a substantial increase from 27.4% a year ago and more than triple the 10.7% observed two years ago.

This progression reflects AI's deepening integration into business-critical functions and workflows that inherently involve sensitive information. As AI proves its value in one context, organizations naturally expand its use to more sensitive applications, often without correspondingly enhanced security controls.


Examining the specific categories of sensitive data reveals concerning patterns. The most common types of sensitive data employees put into AI are source code (18.7% of sensitive data), R&D materials (17.1%), highlighting AI's growing role in product development processes. Sales and marketing data constitutes another 10.7%, including marketing plans and confidential data about customers.

Perhaps most alarming are the healthcare-related findings. Health records comprise 7.4% of sensitive data going into AI, such as when medical professionals use AI to draft communications with insurers or summarize patient visits. Similarly, HR and employee records account for 4.8% of sensitive data, with AI increasingly used to draft performance reviews and handle confidential personnel matters.

The risk implications become even more apparent when examining where this sensitive data is going. A staggering 83.8% of enterprise data input into AI tools flows to platforms classified as medium, high, or critical risk – with just 16.2% destined for enterprise-ready, low-risk alternatives. This imbalance highlights the urgent need for organizations to implement more robust AI governance controls.


Amount of corporate data sent to AI that is sensitive

(By volume of data)




Sensitive corporate data sent to AI by type

(By volume of data)



Most enterprise data is going to risky AI tools

(Percentage of corporate data going to AI tools by app risk level)



How AI-generated content is used at work

Understanding how employees use AI-generated content represents an opportunity for organizations seeking to scale AI's benefits. By identifying successful usage patterns in early adopters, companies can strategically extend these approaches across the organization.

Our analysis shows that 35.9% of AI-generated content flows into email and messaging platforms, making communication the dominant use case.

Cloud documents receive 18.0% of AI-generated content, spanning everything from summarizing strategic planning documents to formulas used in spreadsheet calculations. This widespread usage in collaborative workspaces indicates significant value in knowledge work.

Technical use cases show promising adoption, with 10.8% of AI material entering source code management


systems. Organizations can leverage these patterns to implement enterprise-grade AI coding assistants that provide similar benefits with appropriate guardrails.

IT and security functions are embracing AI-generated content as well, with 5.5% of outputs appearing in infrastructure and security tools. This typically manifests as automation scripts and configuration templates, demonstrating AI's value in streamlining technical operations.

While understanding these usage patterns is valuable for scaling AI benefits, organizations should also consider potential risks. These include accuracy concerns in business communications, intellectual property considerations in creative content, and security implications in technical implementations.


Top destinations of AI generated content in the enterprise

(Percentage of AI-generated data by destination)



How employees of different seniority use AI at work

AI adoption follows distinct patterns across organizational hierarchies, with mid-level employees emerging as the most enthusiastic adopters. Analysts, specialists, and similar mid-tier roles use AI tools 3.5 times more frequently than the next-highest cohort (manager-level employees), suggesting a sweet spot where employees have both the autonomy to adopt new tools and the practical knowledge needed to apply these tools to increase their own productivity.



How developers use AI to write code

Software development represents one of the most transformative areas of AI adoption in the enterprise. Our data shows that coding tasks generate some of the highest volumes of both inbound and outbound AI data flows across all business functions, highlighting the fit between AI capabilities and developer workflows.


While developers typically begin their AI journey through grassroots experimentation, the impact becomes more significant once organizations formally support these tools. When companies officially deploy specialized AI development environments like Cursor or Cline, grows by 400% in the first four months after rollout, quickly becoming integral to development processes.

This formal adoption reshapes established development workflows. Traditional integrated development environments (IDEs) such as VS Code, Xcode, and PyCharm experience a 23.7% decline in usage when AI alternatives become available. This shift demonstrates that AI-powered coding tools aren't merely supplements to existing processes but are actively transforming how software gets built.

The seniority pattern in AI usage is also present within technical teams. Among software engineers, mid-level professionals (typically Senior Software Engineers) demonstrate the highest AI usage rates, outpacing both entry-level developers and higher-seniority Staff Software Engineering leaders. This suggests that some baseline technical proficiency is beneficial for effectively leveraging AI capabilities, but the highest-ranking technical experts may be more cautious or have less direct need for AI assistance.


The introduction of AI development tooling leads to a decrease in traditional developer tools

(Events per 1,000 users per day before and after company-wide rollout of AI developer tool)



AI usage by software engineer seniority level

(Average number of daily events per employee)



DeepSeek model usage in development projects


Within software development projects, open-source AI models have gained significant traction over the past year. Llama has established a dominant position, consistently accounting for at least 50% of local model development over the past twelve months as developers build custom AI applications and services.

However, the January 2025 release of DeepSeek R1 disrupted the market. Developer adoption of DeepSeek surged rapidly, reaching 17.7% of AI development activity by February – firmly establishing it as the second-most utilized model behind Llama. However, this initial enthusiasm partially subsided by March 2025, with usage settling at 11.0% of developer activity.

Looking at the geographic origins of models reveals clear patterns in developer preferences. AI engineers overwhelmingly build applications using models developed in the United States, with open-source alternatives from China forming a solid second tier. European AI models significantly trail these leaders, suggesting an innovation gap in this strategically important technology area.


Local model usage by AI developers

(By percentage of usage events)

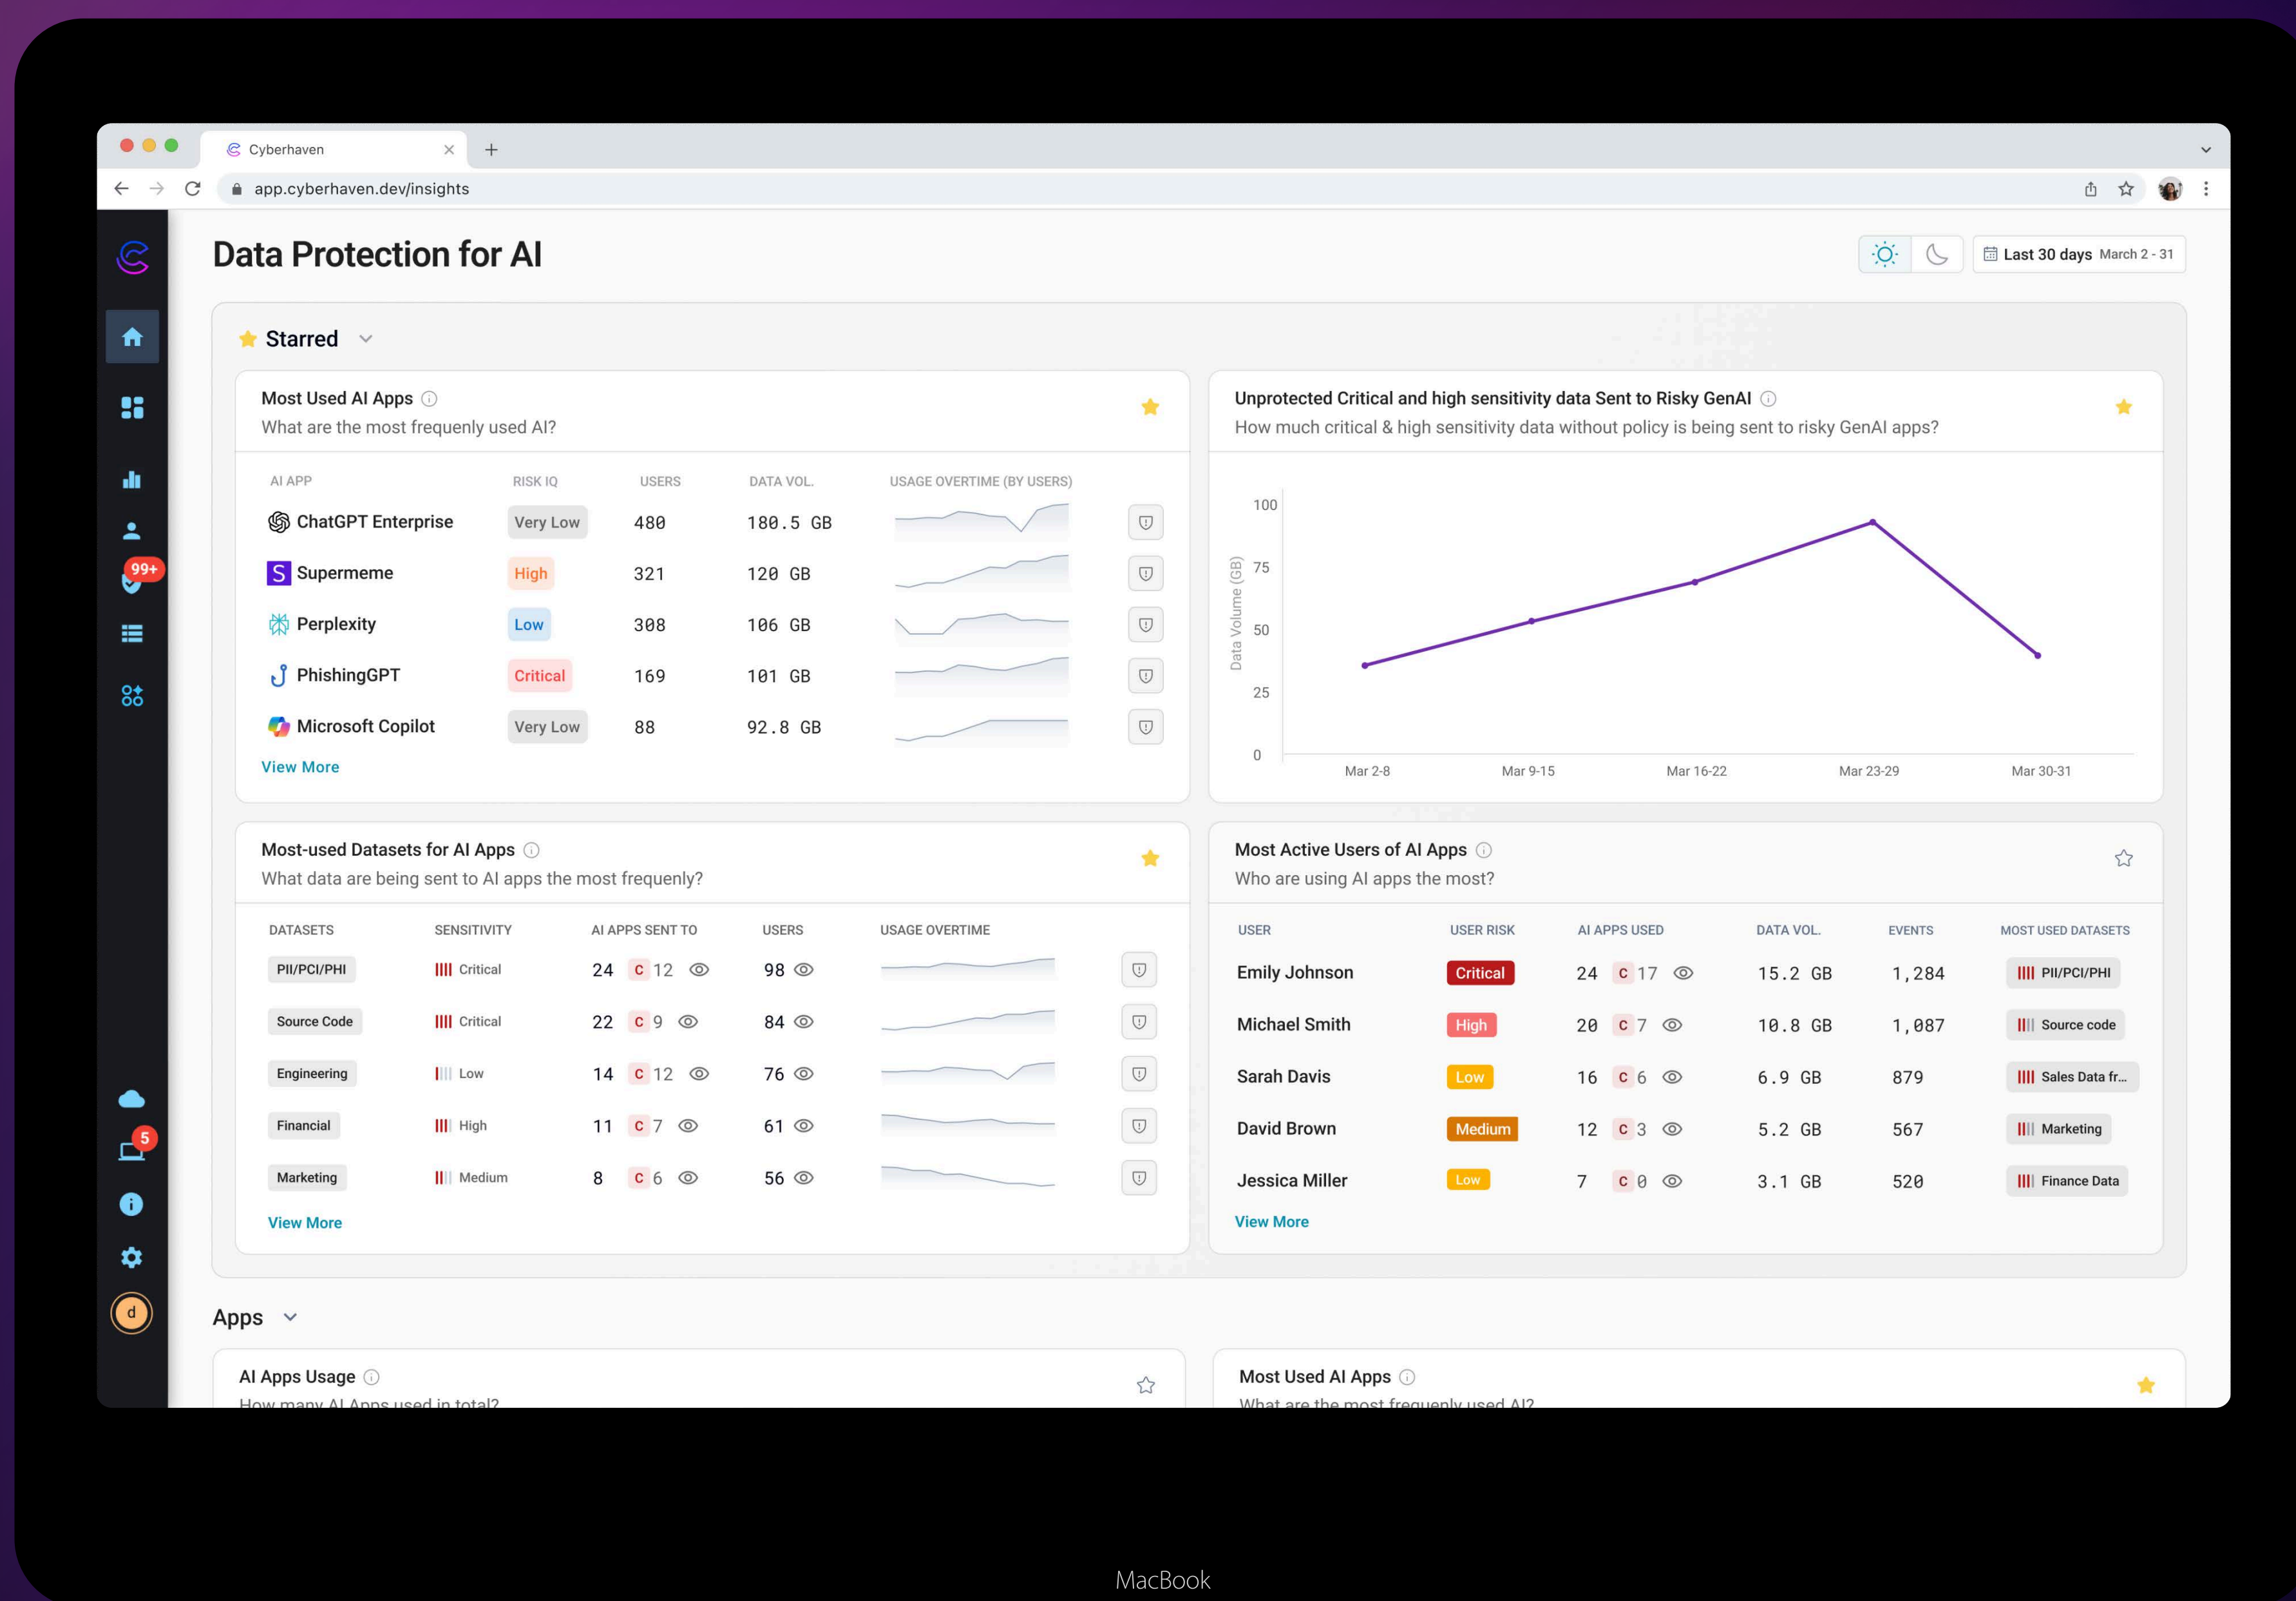


AI engineers prefer building with models from the US and China

(User count per 1,000 AI engineers in the past 3 months)




DISCOVER YOUR ORGANIZATION'S AI USAGE AND RISK



If you found the insights in this report insightful, let Cyberhaven audit your AI usage and reveal how your employee base is adopting AI and how your data flows to and from AI tools.

[Schedule a Demo](#)



Cyberhaven is the AI-powered data security company revolutionizing how companies detect and stop the most critical insider threats to their most important data. Until now, data security products relied on manual rulesets and pre-defined policies that looked for keywords and specific user actions. Our AI technology analyzes billions of workflows to understand every piece of data within an organization, when it's at risk, and what's needed to protect it. It's like nothing that's come before and protects data like nothing else.

To learn more, visit cyberhaven.com