OpenAI

OpenAI Signals

Measuring AI adoption, protecting privacy, and empowering decisions

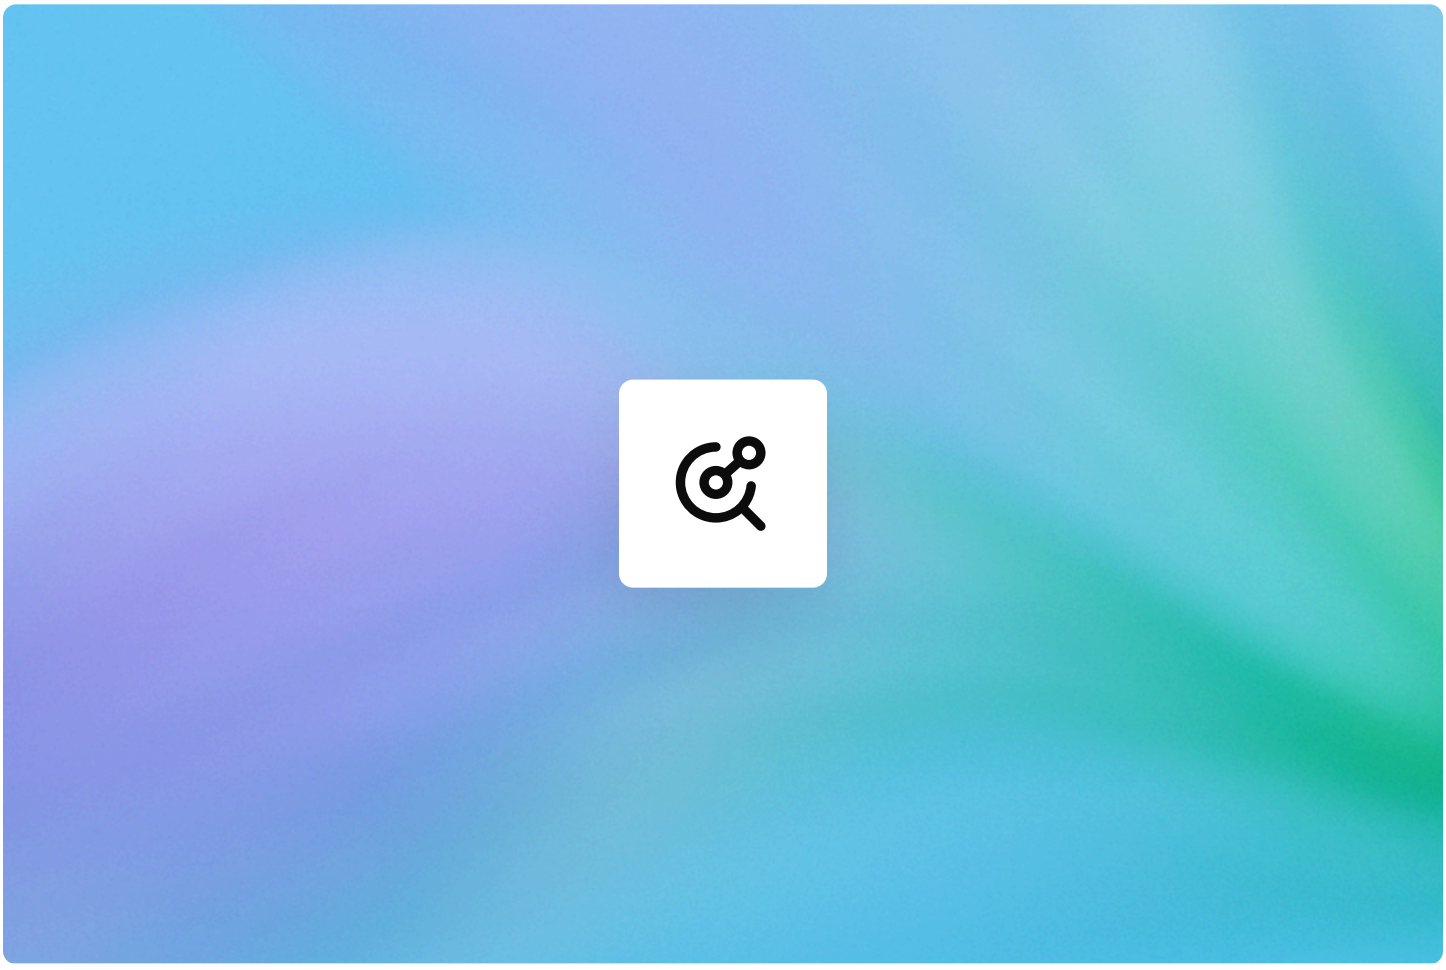

Global affairs

February 2026

Contents

Foreword

Artificial intelligence is moving from novelty to infrastructure. It is being used by students, small businesses, researchers, frontline workers, and entrepreneurs across the world. Adoption is accelerating even faster than traditional economic measurement systems were built to track.

That creates a challenge.

Governments are making decisions about workforce training. Employers are adapting workflows. Educators are redesigning instruction. Workers are deciding how to invest in new skills. These choices require reliable information about how AI is actually being used. Right now, these decision makers are flooded with speculation and anecdotes.

OpenAI Signals is designed to help meet that need.

Signals provides recurring, privacy-preserving data about how ChatGPT is being adopted and used across regions, age groups, and task types. It is built to complement official statistics and academic research. Its purpose is to add transparency and clarity to public discussion about AI's role in the economy.

At the same time, our commitment to protecting user privacy remains unwavering. Signals is engineered so that transparency does not come at the expense of individual protection. We use privacy-protecting technology to guarantee that message content is protected. We share patterns, not people.

As AI becomes part of everyday life, credible measurement becomes part of public infrastructure. Signals is our contribution to building that infrastructure responsibly.

Executive summary

OpenAI Signals is a public data initiative that provides regularly updated, aggregated insights into how ChatGPT is being used around the world.

It is designed to:

- Improve transparency around AI adoption
- Support evidence-based workforce decisions and economic policy
- Help researchers and journalists ground debate in credible data
- Build trust through strong privacy protections

Signals will equip decision makers across the board—from workers to policymakers—with data and insights to make the best decisions possible:

- Interactive dashboards
- Downloadable, de-identified datasets
- Ongoing research briefs and articles highlighting key trends and use cases
- Quarterly updates

All published data is aggregated to ensure quality and privacy safeguards.

Why Signals, and why now

There is currently a lack of consistent, comparable information about how AI tools are being used in the real economy. In that vacuum:

01	Public debate can be driven by isolated anecdotes
02	Policymakers lack consistent reference points

AI adoption data is increasingly shaping public narratives. Signals aims to provide a stable, transparent reference point grounded in rigorous methods and strong privacy protections.

This initiative builds on prior research explaining how people use ChatGPT and extends it into an ongoing public measurement system.

Data democratized for everyone from workers to policymakers

Building an AI-enabled economy that expands opportunity depends on something basic: clear, credible measurement. Right now, many decisions about training, education, business investment, and public policy are being made in an environment where evidence is thin and narratives move faster than facts. That creates uncertainty for workers, noise for the public, and real downside for leaders trying to allocate resources responsibly.

For workers and job seekers, the stakes are personal and immediate.

People want to know how AI is showing up in real jobs, whether it is helping them work faster or changing tasks entirely, and what skills will actually matter. They need more signals grounded in reality.

For educators and workforce organizations, the challenge is operational.

They are expected to redesign programs and credentials quickly, often without a reliable view of where adoption is accelerating, what kinds of tasks are changing, and whether AI is being used mainly for assistance or automation. Without measurement, it's easy to invest in the wrong things or to miss where support is most needed.

For researchers, journalists, and civil society, the challenge is interpretability and accountability.

Public debate about AI improves when it's grounded in consistent evidence. Signals is designed to provide a standardized, regularly updated reference point that can support independent analysis, scrutiny, and better public understanding.

For policymakers, measurement is a prerequisite for effective governance.

If AI adoption is rising unevenly across regions, that shapes where public investment should go. If work-related use is growing, that affects priorities in training, labor-market policy, and economic development. If usage patterns suggest augmentation rather than replacement in many settings, that shifts the focus toward productivity, mobility, and quality of work.

OpenAI Signals is meant to support all of these audiences with a common foundation: timely, privacy-preserving evidence about how AI is actually being used, so decisions can focus on facts rather than guesswork.

What OpenAI Signals measures

OpenAI Signals is organized to answer a simple set of public-interest questions: Who is using AI? Where is adoption accelerating? And what are people using it for, especially in work and learning contexts? The goal is not to track individual usage, but to provide a consistent, aggregated view of how ChatGPT is showing up across the economy and everyday life.

Signals does this in four broad areas

Measuring aggregate adoption and engagement over time

This shows how usage changes across time, age, and other broad demographic indicators. It also tracks how use cases evolve, giving an early view of where interest and activity are shifting as the technology diffuses.

Share of messages by self-reported age group (28-day moving average)

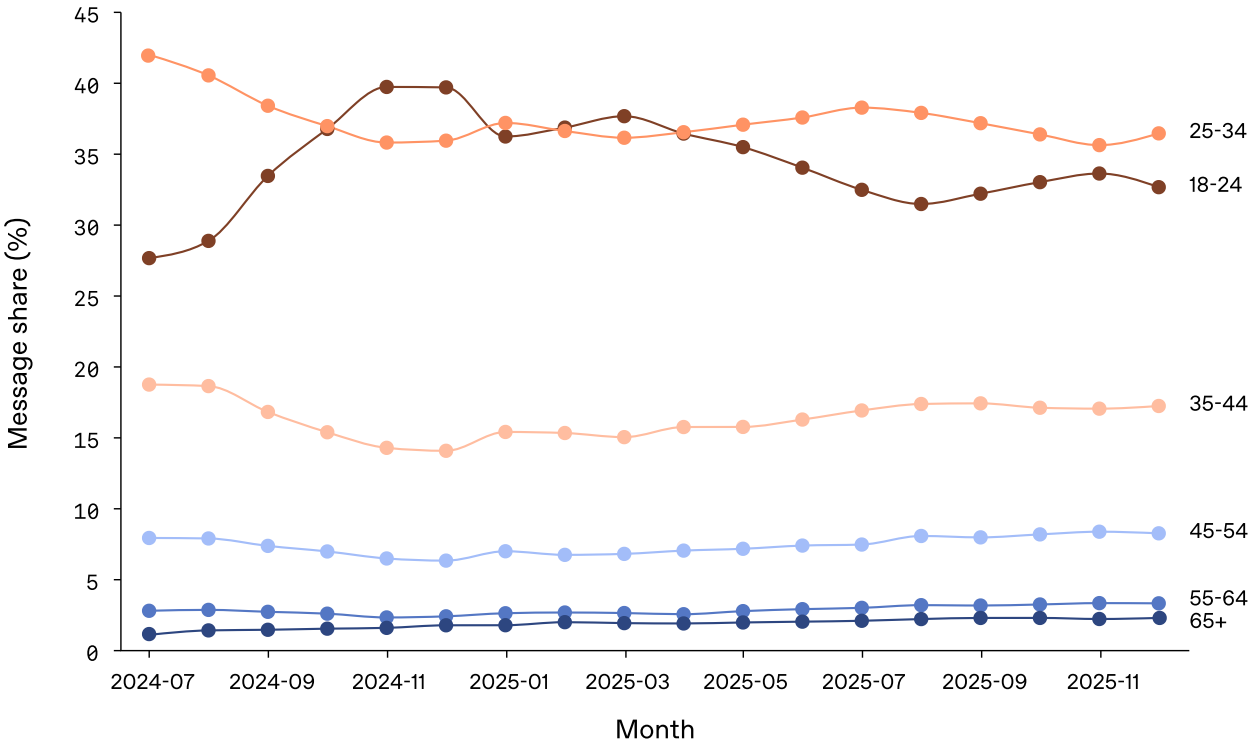

Measuring regional diffusion

Adoption is rarely uniform, and the public conversation is often distorted by assuming everyone is moving at the same pace. Signals helps show how usage differs across countries and, where privacy thresholds allow, across larger subnational regions, such as U.S. states. That makes it easier to see where uptake is higher, where it is low, and changes over time.

Rank of countries by number of ChatGPT messages sent per capita

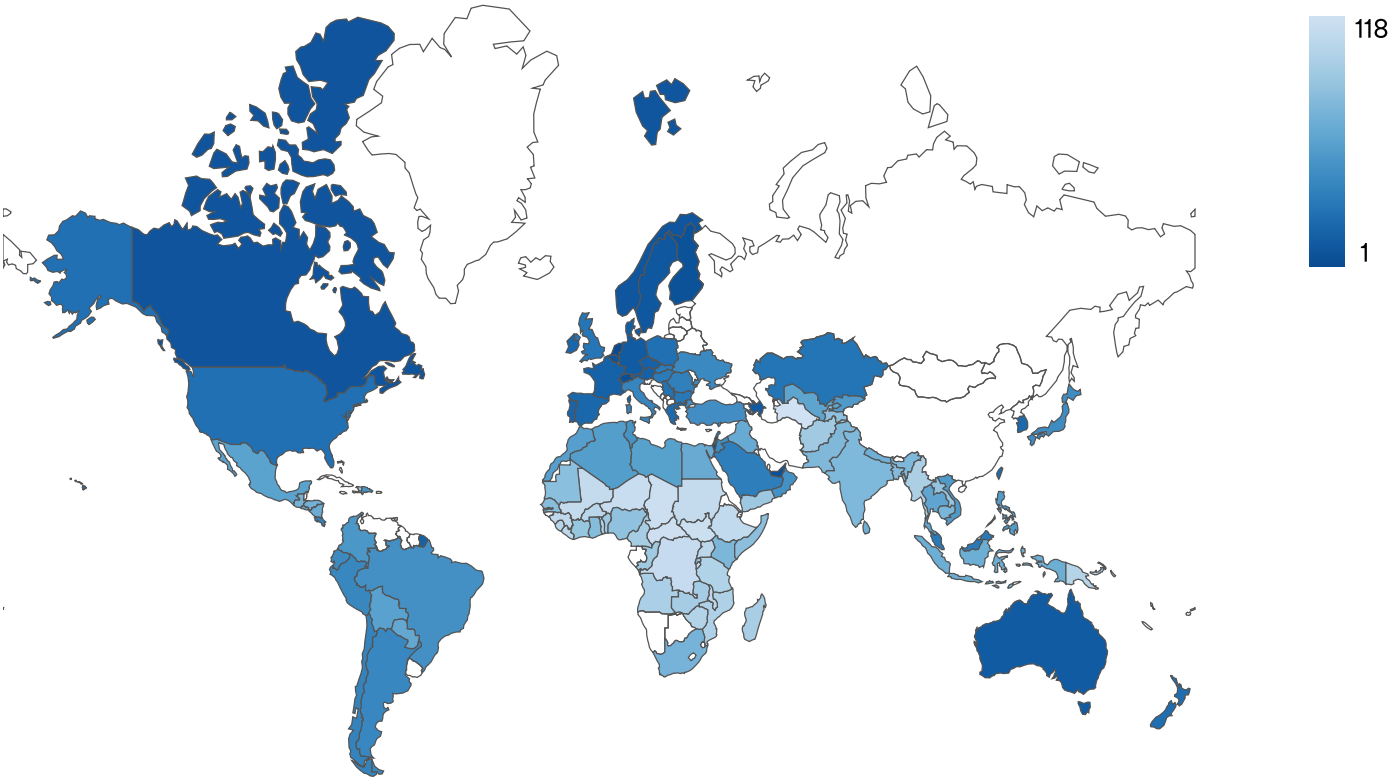

Specifically examining AI usage at work

Signals data estimates the share of usage that is work-related and describes the kinds of work tasks people are bringing to ChatGPT—using automated classification aligned, where appropriate, to established occupational and activity taxonomies (for example, activity categories commonly used in US labor market research).

It also captures how people approach work with AI by categorizing messages into three categories: asking, doing, and expressing. “Asking” messages seek information or advice that will help the user be better informed or make better decisions, either at work, at school, or in their personal life. “Doing” messages request that ChatGPT perform tasks for the user. “Expressing” messages are neither asking for information, nor for the chatbot to perform a task. One important limitation is that the initial release reflects only consumer usage, which likely underrepresents enterprise work use.

Share of work-related and non-work related messages over time in consumer ChatGPT data

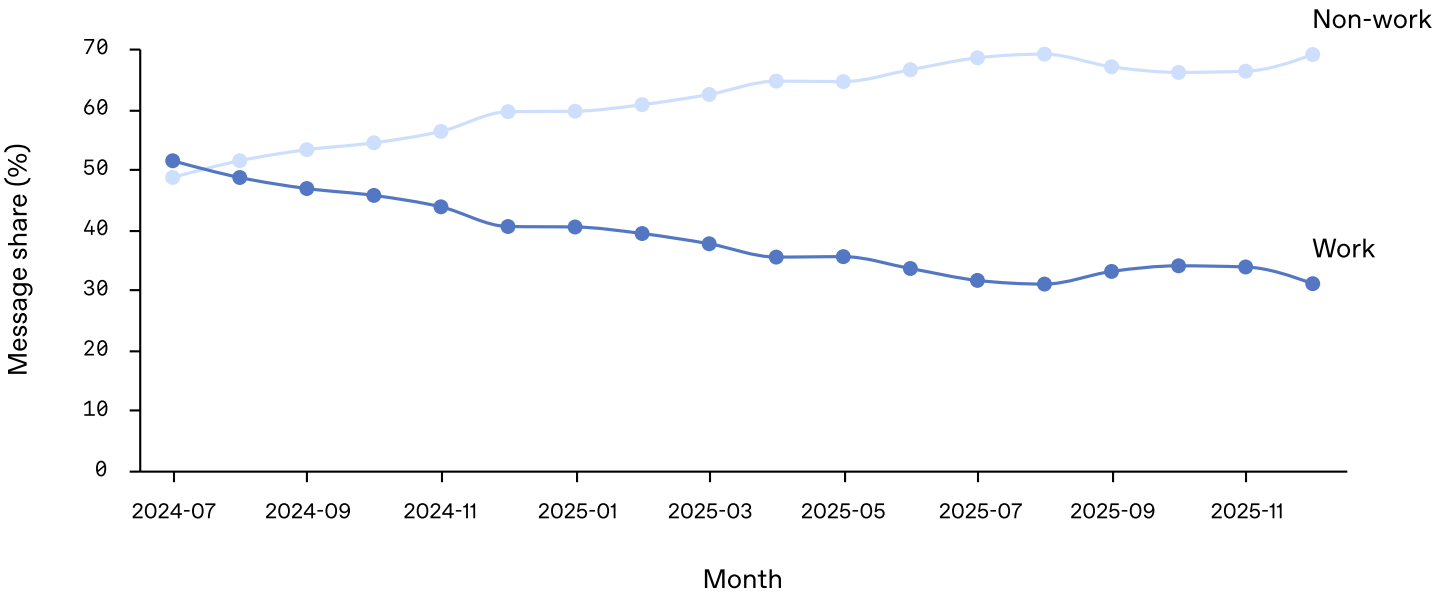

Measuring how AI is used outside of work, including learning and everyday life

This matters because AI is not only a workplace tool—it is also used for studying, personal productivity, creativity, and information seeking. Seeing those patterns over time helps the public and policymakers understand where AI is functioning beyond being a specialized productivity tool and more as a general-purpose assistant.

Share of messages by purpose and conversation topics in consumer ChatGPT data

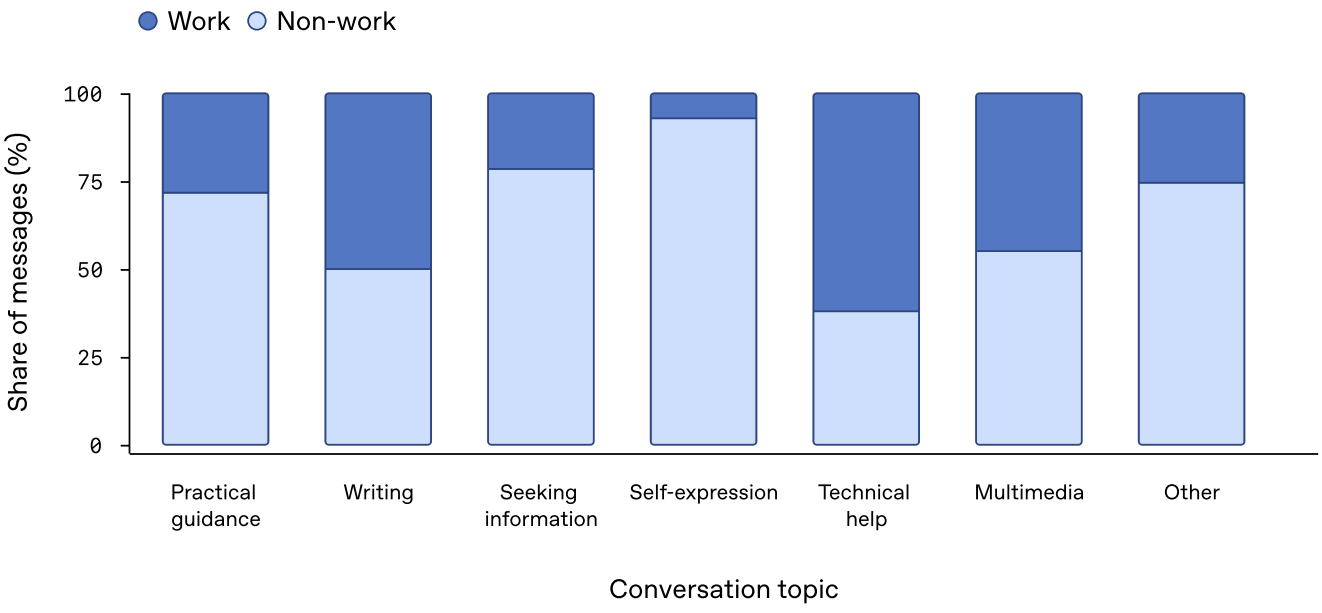

Across all four areas, the emphasis is on trends that are useful for public understanding and policy—published only at levels of aggregation that protect privacy while allowing meaningful comparison.

Privacy by design

Signals is designed to provide transparency into AI usage without weakening user privacy. The platform is designed to publish useful, population-level insights without exposing any individual's activity, identity, or conversations. That commitment is enforced through how the data is processed, what is excluded, what is aggregated, and how results are released.

01	**Signals never publishes information about individual conversations.**	No raw message text is released. Instead, Signals reports only aggregated metrics that describe broad patterns—like the share of usage that appears work-related or how topic distributions change over time.
02	**Signals relies on automated classification.**	Content is processed using automated systems applied to de-identified data that has been stripped of all personally identifying information. This is a core safeguard: the analysis is designed so that people do not read private user conversations to produce the public metrics.
03	**Signals uses strict aggregation thresholds to prevent reporting on small groups.**	Metrics are only published when they reflect sufficiently large populations, and geographic reporting is limited to levels at which anonymity can be preserved.
04	**Signals employs differential privacy to add an additional layer of protection.**	Differential privacy refers to technique to allow researchers to add carefully-calibrated amounts of noise to statistics to guarantee that the underlying data cannot be reverse-engineered.
05	**Additional user protections**	Signals includes additional guardrails such as excluding users with a declared age under 18 and those who have opted out of training data usage.[1]

Ongoing updates

Signals is designed as durable measurement infrastructure and as such will be updated regularly. We will include data refreshes, research and analyses, and user stories. Over time, this creates a consistent time series for researchers, policymakers, and the public.

Conclusion

AI is reshaping how work is done, how people learn, and how businesses operate. Public debate about its impact should be grounded in credible evidence. OpenAI Signals is built to support that goal with regular updates, transparent methods, and strong privacy protections including differential privacy.

Workers, policymakers, and the public deserve clarity. Signals is one step toward providing it.