

February 27, 2025 [Release](#) [Product](#)

# Introducing GPT-4.5

A research preview of our strongest GPT model.  
Available to Pro users and developers worldwide.

[Try in ChatGPT ↗](#)[▶ Listen to article](#) 7:08[🔗 Share](#)

We're releasing a research preview of GPT-4.5—our largest and best model for chat yet. GPT-4.5 is a step forward in scaling up pre-training and post-training. By scaling unsupervised learning, GPT-4.5 improves its ability to recognize patterns, draw connections, and generate creative insights without reasoning.

Early testing shows that interacting with GPT-4.5 feels more natural. Its broader knowledge base, improved ability to follow user intent, and greater “EQ” make it useful for tasks like improving writing, programming, and solving practical problems. We also expect it to hallucinate less.

We're sharing GPT-4.5 as a research preview to better understand its strengths and limitations. We're still exploring what it's capable of and are eager to see how people use it in ways we might not have expected.

## Scaling unsupervised learning

# OpenAI

1. Unsupervised learning increases world model accuracy and intuition. Models like GPT-3.5, GPT-4, and GPT-4.5 advance this paradigm.
2. Scaling reasoning, on the other hand, teaches models to think and produce a chain of thought before they respond, allowing them to tackle complex STEM or logic problems. Models like OpenAI o1 and OpenAI o3-mini advance this paradigm.

GPT-4.5 is an example of scaling unsupervised learning by scaling up compute and data, along with architecture and optimization innovations. GPT-4.5 was trained on Microsoft Azure AI supercomputers. The result is a model that has broader knowledge and a deeper understanding of the world, leading to reduced hallucinations and more reliability across a wide range of topics.

## Scaling the GPT paradigm

GPT-4.5 (2025) ✓

---

What was the first language?



We don't know exactly what the first language was, and scientists believe we'll probably never know for sure because spoken languages appeared long before written records existed.


Here's what we do know:

- **No single "first language":** Languages likely developed gradually in different groups of humans around the same time, tens or even hundreds of thousands of years ago.
- **Origin in Africa:** Most scientists agree that spoken language probably first emerged among humans in Africa, perhaps over 100,000 years ago.




## Deeper world knowledge

**SimpleQA Accuracy (higher is better)**



**SimpleQA Hallucination Rate (lower is better)**

# OpenAI



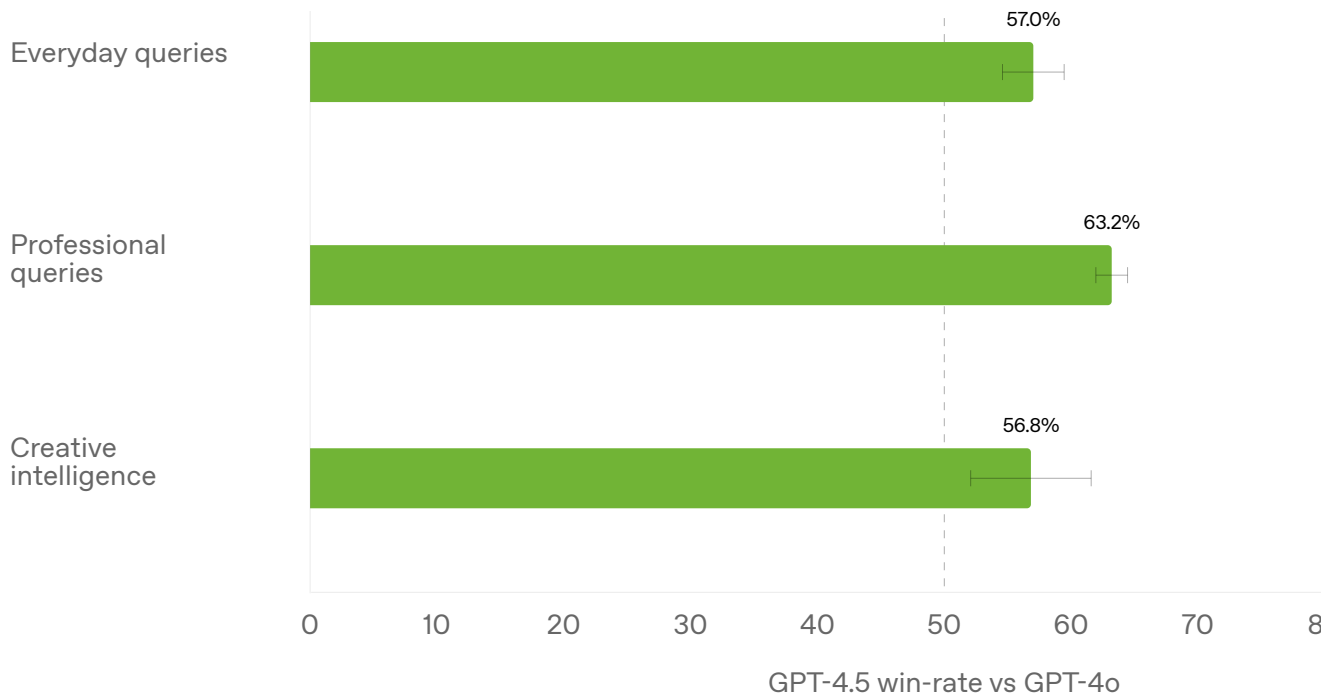
*SimpleQA measures LLM (large language model) factuality on straightforward but challenging knowledge questions.*

## Training for human collaboration

As we scale our models and they solve more complex problems, it becomes increasingly important to teach them a greater understanding of human needs and intent. For GPT-4.5, we developed new, scalable techniques that enable training larger and more powerful models with data derived from smaller models. These techniques improve GPT-4.5's steerability, understanding of nuance, and natural conversation.

### Comparative evaluations with human testers

# OpenAI



*Human preference measures the percentage of queries where testers preferred GPT-4.5 over GPT-4o.*

Combining deep understanding of the world with improved collaboration results in a model that integrates ideas naturally in warm and intuitive conversations that are more attuned to human collaboration. GPT-4.5 has a better understanding of what humans mean and interprets subtle cues or implicit expectations with greater nuance and “EQ”. GPT-4.5 also shows stronger aesthetic intuition and creativity. It excels at helping with writing and design.

## Use cases

Help through a tough time

Identify a painting

Space exploration

GPT-4.5

# OpenAI

I'm going through a tough time after failing a test



Aw, I'm really sorry to hear that. Failing a test can feel pretty tough and discouraging, but remember, it's just one moment—not a reflection of who you are or your capabilities. Want to talk about what happened, or do you just need a distraction? I'm here either way.

GPT-4o ▾

I'm going through a tough time after failing a test



I'm really sorry to hear that you're going through a tough time. Here are a few things you might consider:

1. **Reflect and Understand:** Try to analyze what might have gone wrong. Was it a lack of preparation, time management during the test, or perhaps misunderstanding the material? Understanding the root cause can help you address it in the future.

2. **Seek Support:** Talk to someone you trust—friends, family, or a school counselor. They can offer a different perspective and sometimes just talking about your feelings can be alleviating.

3. **Speak with Your Teacher or Professor:** They can provide insights into

GPT-4.5 shows greater “EQ” and knows when to invite further conversation versus when to provide the user with extensive information.

## Stronger reasoning on the horizon

GPT-4.5 doesn't think before it responds, which makes its strengths particularly different from reasoning models like OpenAI o1. Compared to OpenAI o1 and OpenAI o3-mini, GPT-4.5 is a more general-purpose, innately smarter model. We believe reasoning will be a core capability of future models, and that the two approaches to

# OpenAI

serve as an even stronger foundation for reasoning and tool-using agents.

## Safety

Each increase in model capabilities is also an opportunity to make the models safer. GPT-4.5 was trained with new techniques for supervision that are combined with traditional supervised fine-tuning (SFT) and reinforcement learning from human feedback (RLHF) methods like those used for GPT-4o. We hope this work will serve as a foundation for aligning even more capable future models.

To stress-test our improvements, we conducted a suite of safety tests before deployment, in accordance with our [Preparedness Framework](#). We found that scaling the GPT paradigm contributed to capability improvements across our evaluations. We are publishing the detailed results from these evaluations in the accompanying [system card](#).

## How to use GPT-4.5 in ChatGPT

Starting today, ChatGPT Pro users will be able to select GPT-4.5 in the model picker on web, mobile, and desktop. We will begin rolling out to Plus and Team users next week, then to Enterprise and Edu users the following week.

GPT-4.5 has access to the latest up-to-date information with search, supports file and image uploads, and can use canvas to work on writing and code. However, GPT-4.5 does not currently support multimodal features like Voice Mode, video, and screensharing in ChatGPT. In the future, we will work to simplify the user experience so AI “just works” for you.

## How to use GPT-4.5 in the API

# OpenAI

calling, Structured Outputs, streaming, and system messages. It also supports vision capabilities through image inputs.

Based on early testing, developers may find GPT-4.5 particularly useful for applications that benefit from its higher emotional intelligence and creativity—such as writing help, communication, learning, coaching, and brainstorming. It also shows strong capabilities in agentic planning and execution, including multi-step coding workflows and complex task automation.

GPT-4.5 is a very large and compute-intensive model, making it more expensive than and not a replacement for GPT-4o. Because of this, we're evaluating whether to continue serving it in the API long-term as we balance supporting current capabilities with building future models. We look forward to learning more about its strengths, capabilities, and potential applications in real-world settings. If GPT-4.5 delivers unique value for your use case, your feedback will play an important role in guiding our decision.

## Conclusion

With every new order of magnitude of compute comes novel capabilities. GPT-4.5 is a model at the frontier of what is possible in unsupervised learning. We continue to be surprised by the creativity of the community in uncovering new abilities and unexpected use cases. With GPT-4.5, we invite you to explore the frontier of unsupervised learning and uncover novel capabilities with us.



OpenAI



Appendix

Below, we provide GPT-4.5’s results on standard academic benchmarks to illustrate its current performance on tasks traditionally associated with reasoning. Even by purely scaling up unsupervised learning, GPT-4.5 shows meaningful improvements over previous models like GPT-4o. Still, we look forward to gaining a more complete picture of GPT-4.5’s capabilities through this release, because we recognize academic benchmarks don’t always reflect real-world usefulness.

Model evaluation scores

	GPT-4.5	GPT-4o	OpenAI o3-mini (high)
GPQA (science)	71.4%	53.6%	79.7%
AIME ‘24 (math)	36.7%	9.3%	87.3%
MMMLU (multilingual)	85.1%	81.5%	81.1%
MMMU (multimodal)	74.4%	69.1%	-

# OpenAI

	\$186,125	\$138,750	\$89,625
SWE-Bench Verified (coding)*	38.0%	30.7%	61.0%

\*Numbers shown represent best internal performance.

Authors

OpenAI

Foundational contributors

Alex Paino, Ali Kamali, Amin Tootoonchian, Andrew Tulloch, Ben Sokolowsky, Clemens Winter, Colin Wei, Daniel Kappler, Daniel Levy, Felipe Petroski Such, Geoff Salmon, Ian O’Connell, Jason Teplitz, Kai Chen, Nik Tezak, Prafulla Dhariwal, Rapha Gontijo Lopes, Sam Schoenholz, Youlong Cheng, Yujia Jin, Yunxing Dai

Research

Core contributors

Aiden Low, Alec Radford, Alex Carney, Alex Nichol, Alexis Conneau, Ananya Kumar, Ben Wang, Charlotte Cole , Elizabeth Yang, Gabriel Goh, Hadi Salman, Haitang Hu, Heewoo Jun, Ian Sohl, Ishaan Gulrajani, Jacob Coxon, James Betker, Jamie Kiros, Jessica Landon, Kyle Luther, Lia Guy, Lukas Kondraciuk, Lyric Doshi, Mikhail Pavlov, Qiming Yuan, Reimar Leike, Rowan Zellers, Sean Metzger, Shengjia Zhao, Spencer Papay, Tao Wang

Contributors

Adam Lerer, Aidan McLaughlin, Alexander Prokofiev, Alexandra Barr, Allan Jabri, Ananya Kumar, Andrew Gibiansky, Andrew Schmidt, Casey Chu, Chak Li, Chelsea Voss, Chris Hallacy, Chris Koch, Christine McLeavey, David Mely, Dimitris Tsipras, Eric Sigler, Erin Kavanaugh, Farzad Khorasani, Huiwen Chang, Ilya Kostrikov, Ishaan Singal, Ji Lin, Jiahui Yu, Jing Yu Zhang, John Rizzo, Jong Wook Kim, Joyce Lee, Juntang Zhuang, Leo Liu, Li Jing, Long Ouyang, Louis Feuvrier, Mo Bavarian, Nick Stathas, Nitish Keskar, Oleg Murk, Preston Bowman, Scottie Yan, SQ Mah, Tao Xu, Taylor Gordon, Valerie Qi, Wenda Zhou, Yu Zhang

# OpenAI

## Scaling

### Core contributors

Adam Goucher, Alex Chow, Alex Renzin, Aleksandra Spyra, Avi Nayak, Ben Leimberger, Christopher Hesse, Duc Phong Nguyen, Dinghua Li, Eric Peterson, Francis Zhang, Gene Oden, Kai Fricke, Kai Hayashi, Larry Lv, Leqi Zou, Lin Yang, Madeleine Thompson, Michael Petrov, Miguel Castro, Natalia Gimelshein, Phil Tillet, Reza Zamani, Ryan Cheu Stanley Hsieh, Steve Lee, Stewart Hall, Thomas Raoux, Tianhao Zheng, Vishal Kuo, Yongjik Kim, Yuchen Zhang, Zhuoran Liu

### Contributors

Alvin Wan, Andrew Cann, Antoine Pelisse, Anuj Kalia, Aaron Hurst, Avital Oliver, Brad Barnes, Brian Hsu, Chen Ding, Chen Shen, Cheng Chang, Christian Gibson, Duncan Findlay, Fan Wang, Fangyuan Li, Gianluca Borello, Heather Schmidt, Henrique Ponde de Oliveira Pinto, Ikai Lan, Jiayi Weng, James Crooks, Jos Kraaijeveld, Junru Shao, Kenny Hsu, Kenny Nguyen, Kevin King, Leah Burkhardt, Leo Chen, Linden Li, Lu Zhang, Mahmoud Eariby, Marat Dukhan, Mateusz Litwin, Miki Habryn, Natan LaFontaine, Pavel Belov, Peng Su, Prasad Chakka, Rachel Lim, Rajkumar Samuel, Renaud Gaubert, Rory Carmichael, Sarah Dong, Shantanu Jain, Stephen Logsdon, Todd Underwood, Weixing Zhang, Will Sheu, Weiye Zheng, Yinghai Lu, Yunqiao Zhang

## Safety Systems

Andrea Vallone, Andy Applebaum, Cameron Raymond, Chong Zhang, Dan Mossing, Elizabeth Proehl, Eric Wallace, Evan Mays, Grace Zhao, Ian Kivlichan, Irina Kofman, Joel Parish, Kevin Liu, Keren Gu-Lemberg, Kristen Ying, Lama Ahmad, Lilian Weng, Leon Maksin, Leyton Ho, Meghan Shah, Michael Lampe, Michele Wang, Miles Wang, Olivia Watkins, Phillip Guo, Samuel Miserendino, Sam Toizer, Sandhini Agarwal, Tejal Patwardhan, Tom Dupré la Tour, Tong Mu, Tyna Eloundou, Yunyun Wang

## Deployment

Adam Brandon, Adam Perelman, Adele Li, Akshay Nathan, Alan Hayes, Alfred Xue, Alison Ben, Alec Gorge, Alex Guziel, Alex Iftimie, Ally Bennett, Andrew Chen, Andy Wang, Andy Wood, Angad Singh, Anoop Kotha, Antonia Woodford, Anuj Saharan, Ashley Tyra, Atty Eleti, Ben Schneider, Bessie Ji, Beth Hoover, Bill Chen, Blake Samic, Britney Smith, Brian Yu, Caleb Wang, Cary Bassin, Cary Hudson, Charlie Jatt, Chengdu Huang, Chris Beaumont, Christina Huang, Cristina Scheau, Dana Palmie, Daniel Levine, Daryl Neubieser, Dave Cummings, David Sasaki, Dibya Bhattacharjee, Dylan Hunn, Edwin Arbus, Elaine Ya Le, Enis Sert, Eric Kramer, Fred von Lohmann, Gaby Janatpour, Garrett McGrath, Garrett Ollinger, Gary Yang, Hao Sheng, Harold Hotelling, Janardhanan Vembunarayanan, Jeff Harris, Jeffrey Sabin Matsumoto, Jennifer Robinson, Jessica Liang, Jessica Shieh, Jiacheng Yang, Joel Morris, Joseph Florencio, Josh Kaplan, Kan Wu, Karan Sharma, Karen Li,

# OpenAI

Manolis Liodakis, Matt Nichols, Michelle Fradin, Minal Khan, Mingxuan Wang, Nacho Soto, Natalie Staudacher, Nikunj Handa, Niko Felix, Ning Liu, Olivier Godement, Oona Gleeson, Philip Pronin, Raymond Li, Reah Miyara, Rohan Nuttall, R.J. Marsan, Sara Culver, Scott Ethersmith, Sean Fitzgerald, Shamez Hemani, Sherwin Wu, Shiao Lee, Shuyang Cheng, Siyuan Fu, Spug Golden, Steve Coffey, Steven Heidel, Sundeep Tirumalareddy, Tabarak Khan, Thomas Degry, Thomas Dimson, Tom Stasi, Tomo Hiratsuka, Trevor Creech, Uzair Navid Iftikhar, Victoria Chernova, Victoria Spiegel, Wanning Jiang, Wenlei Xie, Yaming Lin, Yara Khakbaz, Yilei Qian, Yilong Qin, Yo Shavit, Zhi Bie

## Executive Leadership

Bob McGrew, Greg Brockman, Hannah Wong, Jakub Pachocki, Johannes Heidecke, Joanne Jang, Kate Rouch, Kevin Weil, Lauren Itow, Liam Fedus, Mark Chen, Mia Glaese, Mira Murati, Nick Ryder, Sam Altman, Srinivas Narayanan, Tal Broda

Our Research	ChatGPT	For Business	Terms & Policies
Research Index	Explore ChatGPT	Overview	Terms of Use
Research Overview	Team	Company	Privacy Policy
Research Residency	Enterprise		Security
Latest Advancements	Education		Other Policies
	Pricing		
	Download	Our Charter	
OpenAI o1		Careers	
OpenAI o1-mini		Brand	
GPT-4o	Sora	More	
GPT-4o mini	Sora Overview	News	
	Features		

# OpenAI

Safety

Safety Approach

Security & Privacy

Sora log in ↗

API Platform

Platform Overview

Pricing

API log in ↗

Documentation ↗

Developer Forum ↗



OpenAI © 2015–2025  
[Manage Cookies](#)

English United States