

February 2, 2025 Release

# Introducing deep research

An agent that uses reasoning to synthesize large amounts of online information and complete multi-step research tasks for you. Available to Pro users today, Plus and Team next.

[Try on ChatGPT ↗](#)

00:00 | 01:01

[Listen to article](#)

8:54

[Share](#)

# OpenAI

tens of minutes what would take a human many hours.

Deep research is OpenAI's next agent that can do work for you independently—you give it a prompt, and ChatGPT will find, analyze, and synthesize hundreds of online sources to create a comprehensive report at the level of a research analyst. Powered by a version of the upcoming OpenAI o3 model that's optimized for web browsing and data analysis, it leverages reasoning to search, interpret, and analyze massive amounts of text, images, and PDFs on the internet, pivoting as needed in reaction to information it encounters.

The ability to synthesize knowledge is a prerequisite for creating new knowledge. For this reason, deep research marks a significant step toward our broader goal of developing AGI, which we have long envisioned as capable of producing novel scientific research.

## Why we built deep research

Deep research is built for people who do intensive knowledge work in areas like finance, science, policy, and engineering and need thorough, precise, and reliable research. It can be equally useful for discerning shoppers looking for hyper-personalized recommendations on purchases that typically require careful research, like cars, appliances, and furniture. Every output is fully documented, with clear citations and a summary of its thinking, making it easy to reference and verify the information. It is particularly effective at finding niche, non-intuitive information that would require browsing numerous websites. Deep research frees up valuable time by allowing you to offload and expedite complex, time-intensive web research with just one query.

Deep research independently discovers, reasons about, and consolidates insights from across the web. To accomplish this, it was trained on real-world tasks requiring browser and Python tool use, using the same reinforcement learning methods behind OpenAI o1, our first reasoning model. While o1 demonstrates impressive capabilities in coding, math, and other technical domains, many real-world challenges demand

# OpenAI

on the types of problems people face in work and everyday life.

## How to use deep research

In ChatGPT, select ‘deep research’ in the message composer and enter your query. Tell ChatGPT what you need—whether it’s a competitive analysis on streaming platforms or a personalized report on the best commuter bike. You can attach files or spreadsheets to add context to your question. Once it starts running, a sidebar appears with a summary of the steps taken and sources used.

Deep research may take anywhere from 5 to 30 minutes to complete its work, taking the time needed to dive deep into the web. In the meantime, you can step away or work on other tasks—you’ll get a notification once the research is complete. The final output arrives as a report within the chat – in the next few weeks, we will also be adding embedded images, data visualizations, and other analytic outputs in these reports for additional clarity and context.


Compared to deep research, GPT-4o is ideal for real-time, multimodal conversations. For **multi-faceted, domain-specific** inquiries where depth and detail are critical, deep research’s ability to conduct extensive exploration and cite each claim is the difference between a quick summary and a well-documented, verified answer that can be usable as a work product.

- Business
- Needle in a Haystack
- Medical Research
- UX Design
- Shopping
- Genera

### GPT-4o

Help me find iOS and android adoption rates, % who want to learn another language, and change in mobile penetration, over the past 10 years, for top 10 developed and top 10 developing countries by GDP. Lay this info out in a table and separate stats into columns, and include


# OpenAI

-  Compiling comprehensive data on iOS and Android adoption rates, the percentage of populations interested in learning another language, and changes in mobile penetration over the past decade for the top 10 developed and developing countries by GDP is challenging due to data limitations. However, based on available information, we can provide insights and recommendations for targeting markets with a new iOS translation app powered by ChatGPT.

## Deep research

---

Help me find iOS and android adoption rates, % who want to learn another language, and change in mobile penetration, over the past 10 years, for top 10 developed and top 10 developing countries by GDP. Lay this info out in a table and separate stats into columns, and include recommendations on markets to target for a new iOS translation app from ChatGPT, focusing on markets ChatGPT is currently active in.

-  **Top 10 Developed Economies (by GDP) – Smartphone OS Share, Language-Learning Interest, and Mobile Penetration Growth**

Country	iOS Market Share	Android Market Share	Interest in New Language (% of population)	Mobile Penetration Change (2013→2023)
---------	------------------------	----------------------------	--	--

Deep research responds to the prompt in a highly detailed manner, providing side-by-side country-specific data for both top 10 developed and top 10 developing countries for easy reference and comparison. It uses that information to offer detailed market-entry recommendations that are informed and usable.

# OpenAI

Deep research was trained using end-to-end reinforcement learning on hard browsing and reasoning tasks across a range of domains. Through that training, it learned to plan and execute a multi-step trajectory to find the data it needs, backtracking and reacting to real-time information where necessary. The model is also able to browse over user uploaded files, plot and iterate on graphs using the python tool, embed both generated graphs and images from websites in its responses, and cite specific sentences or passages from its sources. As a result of this training, it reaches new highs on a number of public evaluations focused on real-world problems.

## Humanity's Last Exam

On Humanity’s Last Exam, a recently released evaluation that tests AI across a broad range of subjects on expert-level questions, the model powering deep research scores a new high at 26.6% accuracy. This test consists of over 3,000 multiple choice and short answer questions across more than 100 subjects from linguistics to rocket science, classics to ecology. Compared to OpenAI o1, the largest gains appeared in chemistry, humanities and social sciences, and mathematics. The model powering deep research showcased a human-like approach by effectively seeking out specialized information when necessary.

Model	Accuracy (%)
GPT-4o	3.3
Grok-2	3.8
Claude 3.5 Sonnet	4.3
Gemini Thinking	6.2
<b>OpenAI o1</b>	<b>9.1</b>
DeepSeek-R1*	9.4
OpenAI o3-mini (medium)*	10.5

# OpenAI

OpenAI o3-mini (high)*	13.0
<b>OpenAI deep research**</b>	<b>26.6</b>

\* Model is not multi-modal,  
evaluated on text-only subset.  
\*\*with browsing + python tools

## GAIA

On [GAIA](#)<sup>1</sup>, a public benchmark that evaluates AI on real-world questions, the model powering deep research reaches a new state of the art (SOTA), topping the external [leaderboard](#). Encompassing questions across three levels of difficulty, successful completion of these tasks requires abilities including reasoning, multi-modal fluency, web browsing, and tool-use proficiency.

GAIA				
	Level 1	Level 2	Level 3	Avg.
Previous SOTA ↗	67.92	67.44	42.31	63.64
Deep Research (pass@1)	74.29	69.06	47.6	67.36
<b>Deep Research (cons@64)</b>	<b>78.66</b>	<b>73.21</b>	<b>58.03</b>	<b>72.57</b>

### GAIA task examples

Level 1      Level 2      Level 3

# OpenAI

In Nature journal's Scientific Reports conference proceedings from 2012, in the article that did not mention plasmons or plasmonics, what nano-compound is studied? Don't use the prefix nano in your answer if there is one.



[The remaining Chain-of-Thought messages and the final answer are hidden to protect the benchmark's ground truth answer.]



## Clarifying the search

I'm piecing together the request for an article from 2012 in *Scientific Reports* that details a nano-compound, excluding any related to plasmonics. Likely, this involves filtering specific topics or compounds, like silver or graphene.

### Piecing together details

I'm sub-searching "Scientific Reports 2012 nanoparticle," likely seeking nano-compounds. Interested in filtering out plasmon-related topics, possibly focusing on silver, graphene, or quantum dots.



[search action]



## Navigating search filters

I'm digging into advanced search filters on nature.com and Scientific Reports for 2012, trying to pinpoint articles on amorphous carbon nanoparticles in food caramels.



## Identifying conference proceedings

The search results reveal a conference proceeding titled "2D quasiperiodic plasmonic crystals" from 2012, likely found under the "Conference Proceedings" category in *Scientific Reports*.



[browse action]



## Reflecting on potential issues

# OpenAI

🌐 [browse action]

## 🌀 Considering dynamic options

I'm thinking about a nature site with dynamic content or a cached version. Alternatively, trying a search for specific phrases might help.

🔍 [search action]

## 🌀 Pinpointing key features

OK, I'm identifying four conference proceedings from 2012 in Sci Rep. Notably, some mention "plasmonic" and should be excluded.

🌐 [browse action]

## 🌀 Piecing together insights

I'm gathering info on accessing cached versions or snippet expansions for specific titles to ensure comprehensive analysis.

🔍 [search action]

## 🌀 Pulling together the details

I'm noting the progression from the 2012 European Materials Research Society Spring Meeting to the scientific report, focusing on conference proceedings. This approach helps establish context and relevance.

🔍 [search action]

## 🌀 Tracking journal clues

Trying to connect conference proceedings from late 2012 with "Scientific Reports" journal, focusing on plasmonic 2D quasiperiodic materials. This approach feels methodical and precise.

🔍 [search action]

## 🌀 Analyzing event details

Noting that the search results [8] include a page likely summarizing the European Materials Research Society 2012 Spring Meeting, including



# OpenAI

🌐 [browse action]

## 🌀 Piecing together research

I'm examining a likely collection page with 4 articles, considering plasmonic and metamaterial topics, and identifying key references from the European Materials Research Society 2012 Spring Meeting.

🌐 [browse action]

## 🌀 Pinpointing the focus

I'm focusing on the 2012 conference proceedings in "Scientific Reports" from E-MRS, likely involving topics like "2D quasiperiodic plasmonic crystals" and "Layered plasmonic cloaks to tailor the optical scattering at the nanoscale."

🌐 [browse action]

## 🌀 Evaluating library sources

Thinking about special issue E-MRS 2012 Sci rep invited paper Monticone and metamaterials lab at News Archives – 2012.

🌐 [browse action]

## 🌀 Assessing accessibility

I'm gathering info on potential site access issues, considering outdated status or blockages. Also, evaluating the Monticone site for a list reference.

🌐 [browse action]

## 🌀 Piecing together E-MRS 2012

I'm trying to link "Layered Plasmonic Cloaks" to E-MRS 2012. Considering other possible topics like graphene and catalysis.

🌐 [browse action]

## 🌀 Identifying relevant topics

# OpenAI

leads me to focus on nanotechnology aspects more broadly.

🌐 [browse action]

## 🌀 **Adjusting my focus**

I'm broadening the search parameters by entering specific site and scholar identifiers to pinpoint articles about 2D quasiperiodic plasmonic crystals in Nature Scientific Reports from 2012.

🔍 [search action]

## 🌀 **Identifying non-plasmonic possibilities**

I'm narrowing down the search to identify articles from the 2012 E-MRS special issue in Scientific Reports. This involves checking titles for non-plasmonic topics, like graphene or functional oxides. Progress is evident through focused searches and exclusion of plasmonic articles.

🔍 [search action]

## 🌀 **Piecing together clues**

I'm analyzing interactions in a graphene superlattice, likely from a Sci Rep article. Considering titles by authors like Padooru, Yakovlev, Chen, and Alu.

🔍 [search action]

## 🌀 **Piecing together the puzzle**

I'm gathering info on E-MRS 2012 Spring Meeting's Sci Rep proceedings, focusing on invited papers, without plasmonic phenomena. Types under review are graphene, catalysis, energy materials, nanomedicine, and sensors.

🔍 [search action]

The more the model browses and thinks about what its browsing, the better it does, which is why giving it time to think is important.

## Expert-level task examples

Time saved on task: 4 hours

I am a researcher interested in bringing mixed-gas sorption capabilities to my lab. Please discuss the differences between pure- and mixed-gas sorption for glassy polymers, how the dual-mode sorption model can be used to predict mixed-gas sorption behavior in glassy polymers (include equations where applicable), and what challenges there are in accurately predicting pure- and mixed-gas sorption using the dual-mode sorption model. Only include open-source information as of October 1, 2024.



**Pure vs. mixed-gas sorption in glassy polymers:** In a glassy polymer (below  $T_g$ ), a single gas typically exhibits a concave sorption isotherm: uptake rises sublinearly with pressure, unlike the linear Henry's-law behavior above  $T_g$ . This is attributed to "dual-mode" sorption: some penetrant dissolves in the dense polymer matrix (obeying Henry's law) while additional penetrant saturates specific microvoid "sites" (a Langmuir-type adsorption in the polymer's nonequilibrium free volume). By contrast, when two or more gases are present, they **compete** for those limited Langmuir sorption sites. As a result, each component's sorbed concentration in a mixture is generally lower than it would be at the same partial pressure alone. In particular the less soluble (lower-affinity) gas tends to be more strongly suppressed, often causing the mixed-gas *solubility selectivity* (ratio of concentrations) to **exceed** the ideal value predicted from pure-gas isotherms. For example, in  $\text{CO}_2/\text{CH}_4$  mixtures both components sorb less than in their pure-gas cases (especially  $\text{CH}_4$ ), so the  $\text{CO}_2/\text{CH}_4$  sorption selectivity is higher under mixture conditions. In extreme cases, mixed-gas effects can even reverse performance trends observed with single gases. (We note that unlike competitive sorption, *synergistic* sorption – where one gas increases the sorption of another

# OpenAI

**Dual-mode sorption model and mixed-gas prediction:** The dual-mode sorption (DMS) model quantitatively captures glassy polymer sorption by superposing a Henry's-law term and a Langmuir term. For a single gas, the sorbed concentration  $C$  (e.g. in  $\text{cm}^3(\text{STP})$  gas per  $\text{cm}^3$  polymer) is written as:

$$C(p) = k_D p + \frac{C'_H b p}{1 + b p}.$$

Here  $k_D$  is the Henry's law coefficient for dissolution,  $C'_H$  is the Langmuir capacity constant (proportional to the polymer's unrelaxed free-volume or "hole" capacity), and  $b$  is the Langmuir affinity constant. This empirical formula yields the characteristic concave isotherm of glassy polymers and reduces to simple Henry's law at the glass-transition (since  $C'_H \rightarrow 0$  as  $T \rightarrow T_g$ ). The same three parameters ( $k_D$ ,  $C'_H$ ,  $b$ ) are obtained by fitting pure-gas sorption data and are then used to predict mixture sorption. For a binary mixture of gases A and B (with partial pressures  $p_A$ ,  $p_B$ ), the DMS model assumes each gas obeys its own Henry's term and that both compete for the same Langmuir sites. One convenient form is to write the total sorbed concentration as the sum of contributions from each gas:

$$C_{\text{tot}} = k_{D,A} p_A + k_{D,B} p_B + \frac{C'_{H,A} b_A p_A + C'_{H,B} b_B p_B}{1 + b_A p_A + b_B p_B},$$

which implicitly means the concentration of component  $i$  in the mixture is  $C_i = k_{D,i} p_i + \frac{C'_{H,i} b_i p_i}{1 + \sum_j b_j p_j}$ . This reduces to the single-gas formula when the other component's pressure is zero. In practice, partial pressure in these equations is often replaced by fugacity for accuracy at high pressures, though using fugacities or pressures gives nearly the same results if applied consistently. The DMS mixed-gas model predicts the **competitive sorption** effect observed experimentally – namely, each gas's uptake is depressed relative to its pure-gas sorption isotherm.

**Challenges and limitations:** Although the dual-mode model is widely successful in correlating sorption data, it has notable limitations in

## OpenAI

polymer formation history, physical aging, and temperature/pressure range. This means a parameter set obtained under one set of conditions may lack predictive power outside its calibration range. For instance, sorption isotherms measured only over a limited pressure range can be fit well by DMS, but extrapolating beyond that (especially approaching plasticization pressures) may be unreliable. Indeed, multiple distinct combinations of  $k_D$ ,  $C'_H$ ,  $b$  can sometimes fit the same pure-gas isotherm almost equally well – yet they yield markedly different mixed-gas predictions. This sensitivity undermines confidence in mixture forecasts, as one study showed that using different pure-gas fit sets led to only *qualitative* agreement with some mixed-CO<sub>2</sub>/CH<sub>4</sub> data.

Furthermore, DMS inherently assumes a fixed population of Langmuir sites (microvoids) and ignores polymer relaxation or swelling induced by sorbate. In reality, penetrant-induced swelling (plasticization) can increase free volume and sorption capacity, especially for highly condensable gases at high activity – an effect the basic DMS model cannot capture. In binary sorption, this means any **synergistic** increase in sorption due to a second gas (or due to polymer dilation) is unaccounted for in the standard model. Conversely, the model also cannot describe sorption isotherm shapes that deviate from concave Langmuir-type behavior (e.g. sigmoidal isotherms caused by specific polymer-penetrant interactions or clustering). For example, glassy polymers absorbing polar molecules like alcohols can show sigmoidal sorption isotherms, which the simple dual-mode framework fails to represent. Finally, accurately measuring mixed-gas sorption itself is challenging, and discrepancies between idealized model assumptions (e.g. truly zero downstream pressure, no competitive diffusion limitations) and experimental realities can add to prediction error. Despite these challenges, the dual-mode sorption model – with appropriate care – remains a valuable first approximation for interpreting sorption in glassy polymers, and it highlights the key role of competitive site occupancy in mixed-gas environments. Researchers often turn to more elaborate models (e.g. non-equilibrium

# OpenAI

additional parameters or complexity. In summary, dual-mode sorption provides a useful baseline for pure- and mixed-gas sorption in glassy polymers, but one must be mindful of its empirical nature and the above limitations when predicting multicomponent sorption behavior in practice.

**Sources:** Dual-mode sorption model and mixed-gas extension; competitive sorption effects in mixtures; model equations and parameters; limitations of DMS model; parameter sensitivity and predictivity issues.

## Understanding sorption models

I'm gathering info on pure and mixed-gas sorption in glassy polymers, focusing on the dual-mode sorption model and addressing challenges in modeling predictions.

Q [search action]

## Assessing open access

OK, let me see. The result from MDPI on mixed-gas sorption in glassy polymers is probably open access.

⊕ [browse action]

⊕ [browse action]

## Q Examining sections

I'm pulling together insights from the Dual Mode Sorption Model and its extended analysis within the MDPI paper.

⊕ [browse action]

## Clarifying key properties

I'm examining the lab report to understand Henry's law constant, Langmuir affinity constant, and Langmuir capacity constant, which

# OpenAI

🌐 [browse action]

## 🌀 Piecing together

I'm working through the possibility that "Equation (3)" and "Equation (4)" might be images or formula displays and considering consulting the PDF or searching for earlier references to understand them better.

🌐 [browse action]

## 🌀 Confirming title relevance

I'm verifying if the document from Freeman and Koros (1980) specifically matches the reference "Model for sorption of mixed gases in glassy polymers" by checking its abstract and title.

🔍 [search action]

## 🌀 Examining article connections

I noticed that the result from "pmc.ncbi.nlm.nih.gov" is likely the same MDPI article, and I'm curious to see if it presents the text or a simplified version.

🌐 [browse action]

🌐 [browse action]

🌐 [browse action]

## 🌀 Gathering info

I'm gathering information about pure and mixed gas sorption in glassy polymers. This is a critical step in understanding the material's behavior under different gas environments.

🔍 [search action]

## 🌀 Comparing gas sorption

I noticed that carbon dioxide's presence in a mixture affects methane sorption in polymer, reducing methane's sorption compared to a pure gas scenario.

# OpenAI

## **Linking concepts**

I'm thinking through how the term "extended dual mode equation mixture" might relate to dual mode superconducting cable equations. This shows potential for connecting concepts across fields.

 [search action]

## **Requesting review details**

I'm noting the request to review "Theories of Sorption and Transport in Polymer Membrane."

 [browse action]

## **Checking PDF link**

I'm curious about Result [2] and exploring the snippet, which involves extending a dual-mode sorption model to a binary mixture and expressing concentration of A in terms of partial pressures and constants.

 [search action]

## **Predicting sorption behavior**

The dual-mode model anticipates that adding a second gas reduces sorption at a constant partial pressure, unlike pure gas scenarios.

 [browse action]

## **Fetching a PDF**

I'm gathering info on a nature.com PDF, possibly an older Yampol'ski article, but thinking about alternative sources like the University of Texas at Austin repository for a text snippet.

 [browse action]

## **Evaluating repository options**

I've discovered that the UT repository result might not display correctly. Searching within it is uncertain. Alternatively, result [3] contains a snippet



# OpenAI

Q [search action]

Q **Refining the focus**

I'm adjusting my approach to ensure the search targets the right context.

Q [search action]

🌀 **Understanding sorption models**

I'm detailing the dual-mode sorption model for pure gases: Henry dissolving and Langmuir microvoid filling. For mixed gases, competitive sorption and non-idealities are emphasized.

Q [search action]

Q **Exploring hero.epa.gov**

Thinking through hero.epa.gov's potential abstract or content.

🌐 [browse action]

Q **Refining the model**

I'm gathering insights on pure and mixed gas sorption, considering factors like competition, plasticization, and parameter stability. Incorporating mixed-gas data for more accurate predictions.

🌐 [browse action]

🌀 **Addressing dual sorption model challenges**

Parameter coupling and non-uniqueness impact predictions. Absence of polymer condition changes might lead to oversights.

🌐 [browse action]

🌀 **Quantifying deviation**

I'm quantifying nonideal CO<sub>2</sub>-CH<sub>4</sub> solubility in 6FDA-mPDA, showing mixed gas effects on solubility. Competitive sorption isn't the sole cause of permeability selectivity decline.

# OpenAI

## **Incorporating references**

I'm gathering citations, focusing on open-source materials like MDPI 2019 and Vopicka 2014, to support the discussion on differences and challenges in the dual-mode model for gas sorption in polymeric membranes.

 [browse action]

## **Proposing sections**

I'm organizing sections on pure vs mixed-gas sorption, exploring the dual-mode sorption model, and detailing challenges like parameter uncertainty and mixed gas prediction.

### **Organizing content**

I'm focusing on the dual-mode sorption model, detailing pure and mixed-gas sorption behaviors, and highlighting key differences and challenges with bullet points and structured sections for clarity.

 [search action]

## **Choosing citations**

OK, let me see. I'm leaning on [3], [20], [22], and [21] for their clear insights into dual-mode sorption and gas mixtures.

### **Mapping out parameters**

I'm piecing together parameter names to avoid confusion. Sticking with (k), (C<sub>H</sub>'), and (b) for clarity. The equations for pure gas and mixture are detailed and clearly formatted.

### **Addressing challenges**

I'm tackling multi-solution parameter fits, polymer plasticization, and temporal effects. These intricacies make predicting pure and mixed gas sorption tricky.

## **Piecing together**

I'm structuring sections on sorption differences, the dual-mode model for pure gases, its extension to mixed gases, and the associated challenges. This progression seems logical and effective.

# OpenAI

00:00 | 00:00

## Limitations

Deep research unlocks significant new capabilities, but it's still early and has limitations. It can sometimes hallucinate facts in responses or make incorrect inferences, though at a notably lower rate than existing ChatGPT models, according to internal evaluations. It may struggle with distinguishing authoritative information from rumors, and currently shows weakness in confidence calibration, often failing to convey uncertainty accurately. At launch, there may be minor formatting errors in reports and citations, and tasks may take longer to kick off. We expect all these issues to quickly improve with more usage and time.

## Access

Deep research in ChatGPT is currently very compute intensive. The longer it takes to research a query, the more inference compute is required. We are starting with a version optimized for **Pro users** today, with up to 100 queries per month. **Plus and**

# OpenAI

## Economic Area.

All paid users will soon get significantly higher rate limits when we release a faster, more cost-effective version of deep research powered by a smaller model that still provides high quality results.

In the coming weeks and months, we'll be working on the technical infrastructure, closely monitoring the current release, and conducting even more rigorous testing. This aligns with our principle of iterative deployment. If all safety checks continue to meet our release standards, we anticipate releasing deep research to Plus users in about a month.

## What's next

Deep research is available today on ChatGPT web, and will be rolled out to mobile and desktop apps within the month. Currently, deep research can access the open web and any uploaded files. In the future, you'll be able to connect to more specialized data sources—expanding its access to subscription-based or internal resources—to make its output even more robust and personalized.

Looking further ahead, we envision agentic experiences coming together in ChatGPT for asynchronous, real-world research and execution. The combination of deep research, which can perform asynchronous online investigation, and Operator, which can take real-world action, will enable ChatGPT to carry out increasingly sophisticated tasks for you.

---

*Update on February 5, 2025: Deep research is now available to Pro users in the United Kingdom, Switzerland, and the European Economic Area.*

---

# OpenAI

*research, identifying it as Medium risk. We also ran additional safety testing to better understand incremental risks associated with deep research's ability to browse the web, and we have added new mitigations. We will continue to thoroughly test and closely monitor the current limited release. We will share our safety insights and safeguards for deep research in a system card when we widen access to Plus users.*



## Footnotes

- 1 We found that the ground-truth answers for this dataset were widely leaked online and have blocked several websites or URLs accordingly to ensure a fair evaluation of the model. ↩

## Authors

OpenAI

## Research Leads

# OpenAI

## Foundational Contributors

Alex Tachard Passos, Alexandra Barr, Allison Tam, Charlotte Cole, Hyung Won Chung, Jason Wei, Jon Blackman, Scott Mayer McKinney, Valerie Qi

## Core Contributors

### Research

Elaine Ya Le, Eric Mitchell, Eric Wallace, Hyung Won Chung, Ignasi Clavera, Leo Liu, Lorenz Kuhn, Louis Feuvrier, Max Schwarzer, Saachi Jain, Scottie Yan, Shunyu Yao, Vitchyr Pong

### Deployment

Carpus Chang, Harry Zhao, Joseph Trasatti, Joshua Dickens, Matt Kaufer, Mike Trpcic, Minnia Feng, Neel Ajjarapu, Peter Vidani, Sean Fitzgerald

## Contributors

### Research

Ahmed El-Kishky, AJ Ostrow, Alexander Wei, Andrei Gheorghe, Andrew Kondrich, Andrey Mishchenko, Anuj Nair, Behrooz Ghorbani, Brydon Eastman, Chak Li, Foivos Tsimpourlas, Francis Song, Giambattista Parascandolo, Gildas Chabot, Hessam Bagherinezhad, Haitang Hu, Hongyu Ren, Henry Aspegren, Hunter Lightman, Ilya Kostrikov, Ilge Akkaya, James Lennon, Jean Harb, Jonathan Ward, Kai Chen, Katy Shi, Kevin Liu, Kevin Yu, Manuka Stratta, Marvin Zhang, Mengyuan Yan, Mostafa Rohaninejad, Noam Brown, Phoebe Thacker, Raz Goan, Reah Miyara, Spencer Papay, Taylor Gordon, Wenda Zhou, Wenlei Xie, Yash Patil, Yann Dubois, Youlong Cheng, Yushi Wang, Wyatt Thompson

+ all the contributors to o3.

### Safety Systems

Adam Kalai, Alex Beutel, Andrea Vallone, Andy Applebaum, David Robinson, Elizabeth Proehl, Evan Mays, Grace Zhao, Irina Kofman, Jason Phang, Joaquin Quinonero Candela, Joel Parish, Kevin Liu, Kristen Ying, Lama Ahmad, Leon Maksin, Leyton Ho, Meghan Shah, Michele Wang, Miles Wang, Phillip Guo, Olivia Watkins, Owen Campbell-Moore, Patrick Chao, Sam Toizer, Samuel Miserendino, Sandhini Agarwal, Tejal Patwardhan, Tina Sriskandarajah, Troy Peterson, Yaodong Yu, Yunyun Wang

### Deployment

# OpenAI

Hudson, Cory Decareaux, Cristina Scheau, Derek Chen, Dibya Bhattacharjee, Drea Lopez, Eric Antonow, Eric Burke, Filippo Raso, Fotis Chantzis, Freddie Sulit, Harris Cohen, Heather Whitney, Jay Dixit, Jeffrey Han, Jen Robinson, Jessica Shieh, Joel Parish, Kan Wu, Kevin Gladstone, Kshitij Wadhwa, Leo Vandriel, Leyton Ho, Liang Chen, Madeline Christian, Mamie Rheingold, Matt Jones, Michelle Fradin, Mike McClay, Mingxuan Wang, Nacho Soto, Niko Felix, Patrick Delaney, Paul McMillan, Philip Pronin, Rodrigo Rianza Perez, Samuel Miserendino, Scott Ethersmith, Steven Baldwin, Thomas Dimson, Tomo Hiratsuka, Yaming Lin, Yara Khakbaz, Yining Chen

## Leadership

Akshay Nathan, Greg Brockman, Hannah Wong, Jakub Pachocki, Jerry Tworek, Johannes Heidecke, Josh Tobin, Liam Fedus, Mark Chen, Mia Glaese, Nick Turley, Sam Altman, Wojciech Zaremba

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		
	Features		
Sora		News	
		Stories	

# OpenAI

Safety Approach

Security & Privacy

API Platform

Platform Overview

Pricing

API log in ↗

Documentation ↗

Developer Forum ↗



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

English United States















# OpenAI











