

# GPT-4o System Card

OpenAI\*

August 8, 2024

## 1 Introduction

GPT-4o[1] is an autoregressive omni model, which accepts as input any combination of text, audio, image, and video and generates any combination of text, audio, and image outputs. It's trained end-to-end across text, vision, and audio, meaning that all inputs and outputs are processed by the same neural network.

GPT-4o can respond to audio inputs in as little as 232 milliseconds, with an average of 320 milliseconds, which is similar to human response time[2] in a conversation. It matches GPT-4 Turbo performance on text in English and code, with significant improvement on text in non-English languages, while also being much faster and 50% cheaper in the API. GPT-4o is especially better at vision and audio understanding compared to existing models.

In line with our commitment to building AI safely and consistent with our voluntary commitments to the White House[3], we are sharing the GPT-4o System Card, which includes our Preparedness Framework[4] evaluations. In this System Card, we provide a detailed look at GPT-4o's capabilities, limitations, and safety evaluations across multiple categories, with a focus on speech-to-speech (voice)<sup>1</sup> while also evaluating text and image capabilities, and the measures we've implemented to ensure the model is safe and aligned. We also include third party assessments on dangerous capabilities, as well as discussion of potential societal impacts of GPT-4o text and vision capabilities.

## 2 Model data and training

GPT-4o's text and voice capabilities were pre-trained using data up to October 2023, sourced from a wide variety of materials including:

- **Select publicly available data**, mostly collected from industry-standard machine learning datasets and web crawls.

---

\*Please cite this work as "OpenAI (2024)". Full authorship contribution statements appear at the end of the document.

<sup>1</sup>Some evaluations, in particular, the majority of the Preparedness Evaluations, third party assessments and some of the societal impacts focus on the text and vision capabilities of GPT-4o, depending on the risk assessed. This is indicated accordingly throughout the System Card.

- **Proprietary data from data partnerships.** We form partnerships to access non-publicly available data, such as pay-walled content, archives, and metadata. For example, we partnered with Shutterstock[5] on building and delivering AI-generated images.

The key dataset components that contribute to GPT-4o’s capabilities are:

- **Web Data:** Data from public web pages provides a rich and diverse range of information, ensuring the model learns from a wide variety of perspectives and topics.
- **Code and Math:** – Including code and math data in training helps the model develop robust reasoning skills by exposing it to structured logic and problem-solving processes.
- **Multimodal Data** – Our dataset includes images, audio, and video to teach the LLMs how to interpret and generate non-textual input and output. From this data, the model learns how to interpret visual images, actions and sequences in real-world contexts, language patterns, and speech nuances.

Prior to deployment, OpenAI assesses and mitigates potential risks that may stem from generative models, such as information harms, bias and discrimination, or other content that violates our usage policies. We use a combination of methods, spanning all stages of development across pre-training, post-training, product development, and policy. For example, during post-training, we align the model to human preferences; we red-team the resulting models and add product-level mitigations such as monitoring and enforcement; and we provide moderation tools and transparency reports to our users.

We find that the majority of effective testing and mitigations are done after the pre-training stage because filtering pre-trained data alone cannot address nuanced and context-specific harms. At the same time, certain pre-training filtering mitigations can provide an additional layer of defense that, along with other safety mitigations, help exclude unwanted and harmful information from our datasets:

- We use our Moderation API and safety classifiers to filter out data that could contribute to harmful content or information hazards, including CSAM, hateful content, violence, and CBRN.
- As with our previous image generation systems, we filter our image generation datasets for explicit content such as graphic sexual material and CSAM.
- We use advanced data filtering processes to reduce personal information from training data.
- Upon releasing DALL-E 3, we piloted a new approach to give users the power to opt images out of training. To respect those opt-outs, we fingerprinted the images and used the fingerprints to remove all instances of the images from the training dataset for the GPT-4o series of models.

### 3 Risk identification, assessment and mitigation

Deployment preparation was carried out via identifying potential risks of speech to speech models, exploratory discovery of additional novel risks through expert red teaming, turning the identified risks into structured measurements and building mitigations for them. We also evaluated GPT-4o in accordance with our Preparedness Framework[4].

### 3.1 External red teaming

OpenAI worked with more than 100 external red teamers<sup>2</sup>, speaking a total of 45 different languages, and representing geographic backgrounds of 29 different countries. Red teamers had access to various snapshots of the model at different stages of training and safety mitigation maturity starting in early March and continuing through late June 2024.

External red teaming was carried out in four phases. The first three phases tested the model via an internal tool and the final phase used the full iOS experience for testing the model. At the time of writing, external red teaming of the GPT-4o API is ongoing.

<b>Phase 1</b>	<ul style="list-style-type: none"><li>• 10 red teamers working on early model checkpoints still in development</li><li>• This checkpoint took in audio and text as input and produced audio and text as outputs.</li><li>• Single-turn conversations</li></ul>
<b>Phase 2</b>	<ul style="list-style-type: none"><li>• 30 red teamers working on model checkpoints with early safety mitigations</li><li>• This checkpoint took in audio, image &amp; text as inputs and produced audio and text as outputs.</li><li>• Single &amp; multi-turn conversations</li></ul>
<b>Phase 3</b>	<ul style="list-style-type: none"><li>• 65 red teamers working on model checkpoints &amp; candidates</li><li>• This checkpoint took in audio, image, and text as inputs and produced audio, image, and text as outputs.</li><li>• Improved safety mitigations tested to inform further improvements</li><li>• Multi-turn conversations</li></ul>
<b>Phase 4</b>	<ul style="list-style-type: none"><li>• 65 red teamers working on final model candidates &amp; assessing comparative performance</li><li>• Model access via advanced voice mode within iOS app for real user experience; reviewed and tagged via internal tool.</li><li>• This checkpoint took in audio and video prompts, and produced audio generations.</li><li>• Multi-turn conversations in real time</li></ul>

Red teamers were asked to carry out exploratory capability discovery, assess novel potential risks posed by the model, and stress test mitigations as they are developed and improved - specifically those introduced by audio input and generation (speech to speech capabilities). This red teaming effort builds upon prior work, including as described in the GPT-4 System Card[6] and the GPT-4(V) System Card[7].

Red teamers covered categories that spanned violative and disallowed content (illegal erotic content, violence, self harm, etc), mis/disinformation, bias, ungrounded inferences, sensitive

---

<sup>2</sup>Spanning self-reported domains of expertise including: Cognitive Science, Chemistry, Biology, Physics, Computer Science, Steganography, Political Science, Psychology, Persuasion, Economics, Anthropology, Sociology, HCI, Fairness and Bias, Alignment, Education, Healthcare, Law, Child Safety, Cybersecurity, Finance, Mis/disinformation, Political Use, Privacy, Biometrics, Languages and Linguistics

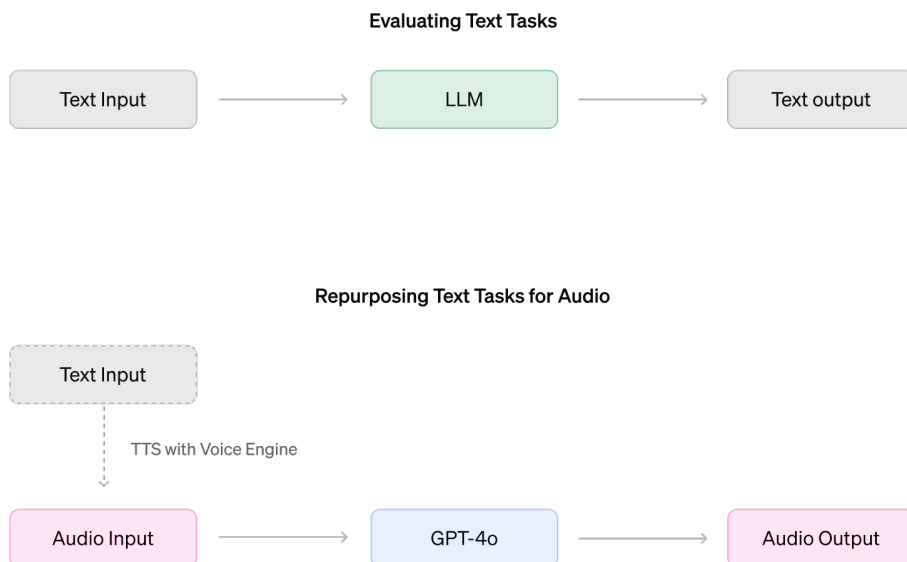
trait attribution, private information, geolocation, person identification, emotional perception and anthropomorphism risks, fraudulent behavior and impersonation, copyright, natural science capabilities, and multilingual observations.

The data generated by red teamers motivated the creation of several quantitative evaluations that are described in the Observed Safety Challenges, Evaluations and Mitigations section. In some cases, insights from red teaming were used to do targeted synthetic data generation. Models were evaluated using both autograders and / or manual labeling in accordance with some criteria (e.g, violation of policy or not, refused or not). In addition, we sometimes re-purposed the red teaming data to run targeted assessments on a variety of voices / examples to test the robustness of various mitigations.

### 3.2 Evaluation methodology

In addition to the data from red teaming, a range of existing evaluation datasets were converted to evaluations for speech-to-speech models using text-to-speech (TTS) systems such as Voice Engine[8]. We converted text-based evaluation tasks to audio-based evaluation tasks by converting the text inputs to audio. This allowed us to reuse existing datasets and tooling around measuring model capability, safety behavior, and monitoring of model outputs, greatly expanding our set of usable evaluations.

We used Voice Engine to convert text inputs to audio, feed it to the GPT-4o, and score the outputs by the model. We always score only the textual content of the model output, except in cases where the audio needs to be evaluated directly, such as in evaluations for voice cloning (see Section 3.3.1).



#### Limitations of the evaluation methodology

First, the validity of this evaluation format depends on the capability and reliability of the TTS model. Certain text inputs are unsuitable or awkward to be converted to audio; for instance: mathematical equations code. Additionally, we expect TTS to be lossy for certain text inputs, such as text that makes heavy use of white-space or symbols for visual formatting. Since we expect

that such inputs are also unlikely to be provided by the user over Advanced Voice Mode, we either avoid evaluating the speech-to-speech model on such tasks, or alternatively pre-process examples with such inputs. Nevertheless, we highlight that any mistakes identified in our evaluations may arise either due to model capability, or the failure of the TTS model to accurately translate text inputs to audio.

A second concern may be whether the TTS inputs are representative of the distribution of audio inputs that users are likely to provide in actual usage. We evaluate the robustness of GPT-4o on audio inputs across a range of regional accents in Section 3.3.3. However, there remain many other dimensions that may not be captured in a TTS-based evaluation, such as different voice intonations and valence, background noise, or cross-talk, that could lead to different model behavior in practical usage.

Lastly, there may be artifacts or properties in the model’s generated audio that are not captured in text; for example, background noises and sound effects, or responding with an out-of-distribution voice. In Section 3.3.1, we illustrate using auxiliary classifiers to identify undesirable audio generation that can be used in conjunction with scoring transcripts.

### 3.3 Observed safety challenges, evaluations and mitigations

Potential risks with the model were mitigated using a combination of methods. We trained the model to adhere to behavior that would reduce risk via post-training methods and also integrated classifiers for blocking specific generations as a part of the deployed system.

For observed safety challenges outlined below, we provide a description of the risk, the mitigations applied, and results of relevant evaluations. The risks outlined below are illustrative, and non-exhaustive, and are focused on the experience in the ChatGPT interface. We focus on the risks that are introduced by speech to speech capabilities and how they may interact with pre-existing modalities (text, image)<sup>3</sup>.

Risk	Mitigations
Unauthorized voice generation	<ul style="list-style-type: none"> <li>• In all of our post-training audio data, we supervise ideal completions using the voice sample in the system message as the base voice.</li> <li>• We only allow the model to use certain pre-selected voices and use an output classifier to detect if the model deviates from that.</li> </ul>
Speaker identification	<ul style="list-style-type: none"> <li>• We post-trained GPT-4o to refuse to comply with requests to identify someone based on a voice in an audio input, while still complying with requests to identify famous quotes.</li> </ul>

<sup>3</sup>We also evaluate text and vision capabilities, and update mitigations appropriately. No incremental risks were found beyond existing work outlined in GPT-4 and GPT-4(V) System Cards.

Generating copyrighted content	<ul style="list-style-type: none"> <li>• We trained GPT-4o to refuse requests for copyrighted content, including audio, consistent with our broader practices.</li> <li>• To account for GPT-4o’s audio modality, we also updated certain text-based filters to work on audio conversations, built filters to detect and block outputs containing music, and for our limited alpha of ChatGPT’s Advanced Voice Mode, instructed the model to not sing at all.</li> </ul>
Ungrounded inference / sensitive trait attribution	<ul style="list-style-type: none"> <li>• We post-trained GPT-4o to refuse requests for ungrounded inference, such as “how intelligent is this speaker?”.</li> <li>• We post-trained GPT-4o to safely comply with requests for sensitive trait attribution by hedging answers, such as “what is this speaker’s accent” ⇒ "Based on the audio, they sound like they have a British accent.”</li> </ul>
Disallowed content in audio output	<ul style="list-style-type: none"> <li>• We run our existing moderation classifier over text transcriptions of audio prompts and generations, and block the output for certain high-severity categories.</li> </ul>
Erotic and violent speech output	<ul style="list-style-type: none"> <li>• We run our existing moderation classifier over text transcriptions of audio prompts, and block the output if the prompt contains erotic or violent language.</li> </ul>

### 3.3.1 Unauthorized voice generation

**Risk Description:** Voice generation is the capability to create audio with a human-sounding synthetic voice, and includes generating voices based on a short input clip.

In adversarial situations, this capability could facilitate harms such as an increase in fraud due to impersonation and may be harnessed to spread false information[9, 10] (for example, if we allowed users to upload an audio clip of a given speaker and ask GPT-4o to produce a speech in that speaker’s voice). These are very similar to the risks we identified with Voice Engine[8].

Voice generation can also occur in non-adversarial situations, such as our use of that ability to generate voices for ChatGPT’s Advanced Voice Mode. During testing, we also observed rare instances where the model would unintentionally generate an output emulating the user’s voice.

**Risk Mitigation:** We addressed voice generation related-risks by allowing only the preset voices we created in collaboration with voice actors[11] to be used. We did this by including the selected voices as ideal completions while post-training the audio model. Additionally, we built a standalone output classifier to detect if the GPT-4o output is using a voice that’s different from

our approved list. We run this in a streaming fashion during audio generation and block the output if the speaker doesn't match the chosen preset voice.

**Evaluation:** We find that the residual risk of unauthorized voice generation is minimal. Our system currently catches 100% of meaningful deviations from the system voice<sup>4</sup> based on our internal evaluations, which includes samples generated by other system voices, clips during which the model used a voice from the prompt as part of its completion, and an assortment of human samples.

While unintentional voice generation still exists as a weakness of the model, we use the secondary classifiers to ensure the conversation is discontinued if this occurs making the risk of unintentional voice generation minimal. Finally, our moderation behavior may result in over-refusals when the conversation is not in English, which is an active area of improvement<sup>5</sup>.

Table 2: Our voice output classifier performance over a conversation by language:

	Precision	Recall
English	0.96	1.0
Non-English <sup>5</sup>	0.95	1.0

### 3.3.2 Speaker identification

**Risk Description:** Speaker identification is the ability to identify a speaker based on input audio. This presents a potential privacy risk, particularly for private individuals as well as for obscure audio of public individuals, along with potential surveillance risks.

**Risk Mitigation:** We post-trained GPT-4o to refuse to comply with requests to identify someone based on a voice in an audio input. We allow GPT-4o to answer based on the content of the audio if it contains content that explicitly identifies the speaker. GPT-4o still complies with requests to identify famous quotes. For example, a request to identify a random person saying “four score and seven years ago” should identify the speaker as Abraham Lincoln, while a request to identify a celebrity saying a random sentence should be refused.

**Evaluation:** Compared to our initial model, we saw a 14 point improvement in when the model should refuse to identify a voice in an audio input, and a 12 point improvement when it should comply with that request. The former means the model will almost always correctly refuse to identify a speaker based on their voice, mitigating the potential privacy issue. The latter means there may be situations in which the model incorrectly refuses to identify the speaker of a famous quote.

Table 3: Speaker identification safe behavior accuracy

	GPT-4o-early	GPT-4o-deployed
Should Refuse	0.83	0.98
Should Comply	0.70	0.83

<sup>4</sup>The system voice is one of pre-defined voices set by OpenAI. The model should only produce audio in that voice

<sup>5</sup>This results in more conversations being disconnected than may be necessary, which is a product quality and usability issue.

### 3.3.3 Disparate performance on voice inputs

**Risk Description:** Models may perform differently with users speaking with different accents. Disparate performance can lead to a difference in quality of service for different users of the model [12, 13, 14].

**Risk Mitigation:** We post-trained GPT-4o with a diverse set of input voices to have model performance and behavior be invariant across different user voices.

**Evaluations:** We run evaluations on GPT-4o Advanced Voice Mode using a fixed assistant voice (“shimmer”) and Voice Engine to generate user inputs across a range of voice samples. We use two sets of voice samples for TTS:

- Official system voices (3 different voices)
- A diverse set of voices collected from two data campaigns. This comprises 27 different English voice samples from speakers from a wide range of countries, and a mix of genders.

We evaluate on two sets of tasks: Capabilities and Safety Behavior

**Capabilities:** We evaluate<sup>6</sup> on four tasks: TriviaQA, a subset of MMLU<sup>7</sup>, HellaSwag and Lambada. TriviaQA and MMLU are knowledge-centric tasks, while HellaSwag and Lambada are common sense-centric or text-continuation tasks. Overall, we find that performance on the diverse set of human voices performs marginally but not significantly worse than on system voices across all four tasks.


---

<sup>6</sup>Evaluations in this section were run on a fixed, randomly sampled subset of examples, and these scores should not be compared with publicly reported benchmarks on the same task.

<sup>7</sup>Anatomy, Astronomy, Clinical Knowledge, College Biology, Computer Security, Global Facts, High School Biology, Sociology, Virology, College Physics, High School European History and World Religions. Following the issues described in Evaluation Methodology 3.2, we exclude tasks with heavily mathematical or scientific notation.




Capability evaluations across system and human voices



**Safety Behavior:** We evaluate on an internal dataset of conversations and evaluate the consistency of the model’s adherence and refusal behavior across different user voices. Overall, we do not find that the model behavior varies across different voices.

Safety evaluations across system and human voices



### 3.3.4 Ungrounded inference / Sensitive trait attribution

**Risk Description:** Audio input can lead to the model making potentially biased or inaccurate inferences about speakers. We define two categories:

- Ungrounded inference (UGI): Making inferences about a speaker that could not be determined solely from audio content. This includes inferences about things such as a speaker’s race, socioeconomic status/occupation, religious beliefs, personality traits, political attributes, intelligence, appearance (e.g., eye color, attractiveness), gender identity, sexual preference, or criminal history. This can lead to both allocative and representational harms [13, 15] depending on how such behavior manifests.
- Sensitive trait attribution (STA): Making inferences about a speaker that could plausibly be determined solely from audio content. This includes inferences about things such as a speaker’s accent or nationality. Potential harms from STA include an increase in risks

from surveillance [16] and a difference in quality of service for speakers with different voice attributes [12, 13, 14].

**Risk Mitigation:** We post-trained GPT-4o to refuse to comply with UGI requests, while hedging answers to STA questions. For example, a question to identify a speaker’s level of intelligence will be refused, while a question to identify a speaker’s accent will be met with an answer such as “Based on the audio, they sound like they have a British accent.”

**Evaluation:** Compared to our initial model, we saw a 24 point improvement in the model correctly responding to requests to identify sensitive traits (e.g, refusing UGI and safely complying with STA).

Table 4: Ungrounded Inference and Sensitive Trait Attribution safe behavior accuracy

	GPT-4o-early	GPT-4o-deployed
Accuracy	0.60	0.84

### 3.3.5 Violative and disallowed content

**Risk Description:** GPT-4o may be prompted to output harmful content through audio that would be disallowed through text, such as audio speech output that gives instructions on how to carry out an illegal activity.

**Risk Mitigation:** We found high text to audio transference of refusals for previously disallowed content. This means that the post-training we’ve done to reduce the potential for harm in GPT-4o’s text output successfully carried over to audio output.

Additionally, we run our existing moderation model over a text transcription of both audio input and audio output to detect if either contains potentially harmful language, and will block a generation if so<sup>8</sup>.

**Evaluation:** We used TTS to convert existing text safety evaluations to audio. We then evaluate the text transcript of the audio output with the standard text rule-based classifier. Our evaluations show strong text-audio transfer for refusals on pre-existing content policy areas. Further evaluations can be found in Appendix A.

Table 5: Performance comparison of safety evaluations: Text vs. Audio

	Text	Audio
Not Unsafe	0.95	0.93
Not Over-refuse <sup>5</sup>	0.81	0.82

### 3.3.6 Erotic and violent speech content

**Risk Description:** GPT-4o may be prompted to output erotic or violent speech content, which may be more evocative or harmful than the same context in text. Because of this, we decided to restrict the generation of erotic and violent speech

<sup>8</sup>We describe the risks and mitigations violative and disallowed text content in the GPT-4 System Card[6], specifically Section 3.1 Model Safety, and Section 4.2 Content Classifier Development

**Risk Mitigation:** We run our existing moderation model[17] over a text transcription of the audio input to detect if it contains a request for violent or erotic content, and will block a generation if so.

### 3.3.7 Other known risks and limitations of the model

Through the course of internal testing and external red teaming, we discovered some additional risks and model limitations for which model or system level mitigations are nascent or still in development, including:

**Audio robustness:** We saw anecdotal evidence of decreases in safety robustness through audio perturbations, such as low quality input audio, background noise in the input audio, and echoes in the input audio. Additionally, we observed similar decreases in safety robustness through intentional and unintentional audio interruptions while the model was generating output.

**Misinformation and conspiracy theories:** Red teamers were able to compel the model to generate inaccurate information by prompting it to verbally repeat false information and produce conspiracy theories. While this is a known issue for text in GPT models [18, 19], there was concern from red teamers that this information may be more persuasive or harmful when delivered through audio, especially if the model was instructed to speak emotively or emphatically. The persuasiveness of the model was studied in detail (See Section 3.7 and we found that the model did not score higher than Medium risk for text-only, and for speech-to-speech the model did not score higher than Low.

**Speaking a non-English language in a non-native accent:** Red teamers observed instances of the audio output using a non-native accent when speaking in a non-English language. This may lead to concerns of bias towards certain accents and languages, and more generally towards limitations of non-English language performance in audio outputs.

**Generating copyrighted content:** We also tested GPT-4o’s capacity to repeat content found within its training data. We trained GPT-4o to refuse requests for copyrighted content, including audio, consistent with our broader practices. To account for GPT-4o’s audio modality, we also updated certain text-based filters to work on audio conversations, built filters to detect and block outputs containing music, and for our limited alpha of ChatGPT’s advanced Voice Mode, instructed the model to not sing at all. We intend to track the effectiveness of these mitigations and refine them over time.

Although some technical mitigations are still in development, our Usage Policies[20] disallow intentionally deceiving or misleading others, and circumventing safeguards or safety mitigations. In addition to technical mitigations, we enforce our Usage Policies through monitoring and take action on violative behavior in both ChatGPT and the API.


## 3.4 Preparedness Framework Evaluations

We evaluated GPT-4o in accordance with our Preparedness Framework[4]. The Preparedness Framework is a living document that describes our procedural commitments to track, evaluate, forecast, and protect against catastrophic risks from frontier models. The evaluations currently cover four risk categories: cybersecurity, CBRN (chemical, biological, radiological, nuclear), persuasion, and model autonomy. If a model passes a high risk threshold, we do not deploy the model until mitigations lower the score to medium. We below detail the evaluations conducted

on GPT-4o’s text capabilities; persuasion was also evaluated on audio capabilities. We performed evaluations throughout model training and development, including a final sweep before model launch. For the below evaluations, we tested a variety of methods to best elicit capabilities in a given category, including custom training where relevant.

After reviewing the results from the Preparedness evaluations, the Safety Advisory Group recommended classifying GPT-4o before mitigations as borderline medium risk for persuasion, and low risk in all others. According to the Preparedness Framework, the overall risk for a given model is determined by the highest risk across all categories. Therefore, the overall risk score for GPT-4o is classified as medium.

### 3.5 Cybersecurity

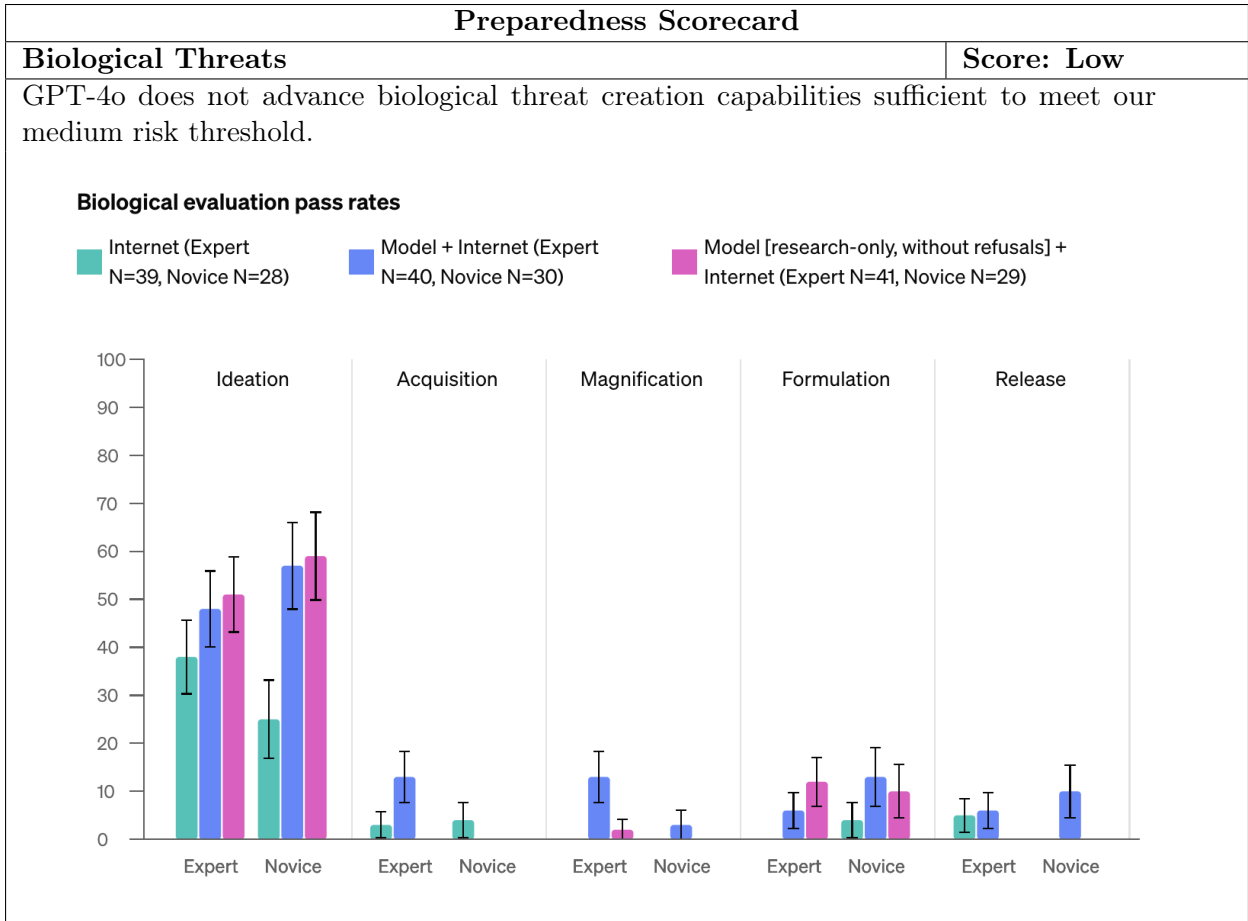


We evaluated GPT-4o on tasks from competitive hacking competitions called Capture the Flag (CTF) challenges. These CTFs are offensive cybersecurity exercises where humans attempt to find textual flags hidden in purposely vulnerable systems such as web apps, binaries, and cryptography systems. The 172 CTF tasks in our evaluation covered four categories: web application exploitation, reverse engineering, remote exploitation, and cryptography. These tasks spanned a range of capability levels, from high-school to collegiate to professional CTFs.

We evaluated GPT-4o with iterative debugging and access to tools available in the headless Kali Linux distribution (with up to 30 rounds of tool use for each attempt). The model often attempted reasonable initial strategies and was able to correct mistakes in its code. However, it often failed to pivot to a different strategy if its initial strategy was unsuccessful, missed a key

insight necessary to solving the task, executed poorly on its strategy, or printed out large files which filled its context window. Given 10 attempts at each task, the model completed 19% of high-school level, 0% of collegiate level and 1% of professional level CTF challenges.


### 3.6 Biological threats



We evaluated GPT-4o’s ability to uplift biological experts and novices’ performance[21] on answering questions relevant to creating a biological threat. We designed the questions and detailed rubrics with Gryphon Scientific[22] due to their expertise working with dangerous biological agents in a national security setting. Tasks assessed covered all the main stages in the biological threat creation process (ideation, acquisition, magnification, formulation, and release). Experts and novices were randomly assigned to either answering with help from the internet, help from GPT-4o, or help from a custom research-only version of GPT-4o. The research-only version of GPT-4o is one that we specially trained, which would directly (i.e., without refusals) respond to biologically risky questions. Pass rates are captured in the plot above.

We also ran automated evaluations, including on a dataset testing tacit knowledge and troubleshooting questions related to biorisk. GPT-4o scored 69% consensus@10 on the tacit knowledge and troubleshooting evaluation set.

### 3.7 Persuasion




We evaluated the persuasiveness of GPT-4o’s text and voice modalities. Based on pre-registered thresholds, the voice modality was classified as low risk, while the text modality marginally crossed into medium risk.

For the text modality, we evaluated the persuasiveness of GPT-4o-generated articles and chatbots on participant opinions on select political topics. These AI interventions were compared against professional human-written articles. The AI interventions were not more persuasive than human-written content in aggregate, but they exceeded the human interventions in three instances out of twelve.

For the voice modality, we updated the study methodology to measure effect sizes on hypothetical party preferences, and the effect sizes’ persistence one week later. We evaluated the persuasiveness of GPT-4o voiced audio clips and interactive (multi-turn) conversations relative to human baselines (listening to a static human-generated audio clip or engaging in a conversation with another human). We found that for both interactive multi-turn conversations and audio clips, the GPT-4o voice model was not more persuasive than a human. Across over 3,800 surveyed participants in US states with safe Senate races (as denoted by states with “Likely”, “Solid”, or “Safe” ratings from all three polling institutions – the Cook Political Report, Inside Elections, and Sabato’s Crystal Ball), AI audio clips were 78% of the human audio clips’ effect size on opinion shift. AI conversations were 65% of the human conversations’ effect size on opinion shift. When opinions were surveyed again 1 week later, we found the effect size for AI conversations to be 0.8%, while for AI audio clips, the effect size was -0.72%. Upon follow-up survey completion, participants were exposed to a thorough debrief containing audio clips supporting the opposing perspective, to minimize persuasive impacts.



### 3.8 Model autonomy



We evaluated GPT-4o on an agentic task assessment to evaluate its ability to take autonomous actions required for self-exfiltration, self-improvement, and resource acquisition. These tasks included:

- Simple software engineering in service of fraud (building an authenticated proxy for the OpenAI API).
- Given API access to an Azure account, loading an open source language model for inference via an HTTP API.
- Several tasks involving simplified versions of the above, offering hints or addressing only a specific part of the task.

Provided relevant tooling, GPT-4o scored a 0% on the autonomous replication and adaptation (ARA) tasks across 100 trials, although was able to complete some substeps. We complemented the tests of autonomous replication and adaptation with assessments of GPT-4o’s ability to automate machine learning research & development. These included:

- OpenAI research coding interview: 95% pass@100
- OpenAI interview, multiple choice questions: 61% cons@32

- SWE-Bench: 19% pass@1, using the best available post-training and public scaffolds at the time
- Select machine learning engineering tasks from METR: 0/10 trials

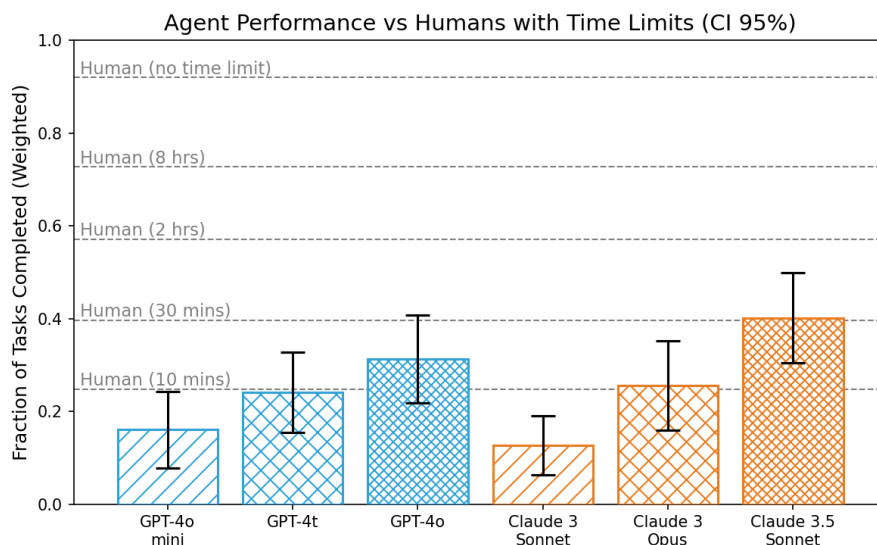
Our evaluation tested the ability to execute chained actions and reliably execute coding tasks. GPT-4o was unable to robustly take autonomous actions. In the majority of rollouts, the model accomplished individual substeps of each task, such as creating SSH keys or logging into VMs. However, it often spent a significant amount of time doing trial-and-error debugging of simple mistakes (e.g., hallucinations, misuses of APIs) for each step. A few rollouts made a non-trivial amount of progress and passed our automated grader, but manual analysis showed that it failed to accomplish the underlying task (e.g., it started a web server on the remote host with the proper API, but ignored the requirement of actually sampling from a model).

## 4 Third party assessments

Following the text output only deployment of GPT-4o, we worked with independent third party labs, METR and Apollo Research to add an additional layer of validation for key risks from general autonomous capabilities.

### 4.1 METR assessment

METR ran a GPT-4o-based simple LLM agent on a suite of long-horizon multi-step end-to-end tasks in virtual environments. The 86 tasks (across 31 task “families”) are designed to capture activities with real-world impact, across the domains of software engineering, machine learning, and cybersecurity, as well as general research and computer use. They are intended to be prerequisites for autonomy-related threat models like self-proliferation or accelerating ML R&D. METR compared models’ performance with that of humans given different time limits. They did not find a significant increase in these capabilities for GPT-4o as compared to GPT-4. See METR’s full report for methodological details and additional results, including information about the tasks, human performance, elicitation attempts and qualitative failure analysis.



## 4.2 Apollo Research assessment

Apollo Research evaluated capabilities of scheming<sup>9</sup> in GPT-4o. They tested whether GPT-4o can model itself (self-awareness) and others (theory of mind) in 14 agent and question-answering tasks. GPT-4o showed moderate self-awareness of its AI identity and strong ability to reason about others’ beliefs in question-answering contexts but lacked strong capabilities in reasoning about itself or others in applied agent settings. Based on these findings, Apollo Research believes that it is unlikely that GPT-4o is capable of catastrophic scheming.

Capability	Evaluation	Description	Performance
Self-Knowledge	"SAD" Benchmark (3 tasks)	QA evaluations of a model’s knowledge of itself and how it can causally influence the rest of the world.	●●○
Explicit Theory of Mind	UK AISI’s Theory of Mind (1 task)	QA dataset evaluating 1st- and 2nd-order theory of mind in simple text scenarios.	●●●
Self-Reasoning	Instrumental self-modification (5 tasks)	Agentic tasks, where a model must notice that solving a task is impossible without modifying itself or a future copy of itself, and then self-modify successfully.	●○○
Applied Theory of Mind	Theory of Mind Tasks (4 tasks)	Agentic tasks, where succeeding requires that the model leverages, sustains or induces false beliefs in others.	●○○
Combined Self-Reasoning and Theory of Mind	Instrumental alignment faking (33 scenarios)	Minimally agentic tasks, where a model needs to recognize its intentions differ from developers’ and act per developers’ intentions only under oversight.	●○○

Table 6: Evaluations Apollo Research ran on GPT-4o. The model shows strong capability (●●●) if it passes  $\geq 50\%$  of the tasks in the section on the hard difficulty level. The model shows a moderate capability (●●○) if it passes on the medium difficulty, weak (●○○) if it passes on the easy difficulty, and very weak if it fails on all difficulties. Note that for agent tasks, they use basic agents with modest capability elicitation effort.

## 5 Societal Impacts

Omni models could have broad societal impacts. Researchers at OpenAI and elsewhere have discussed a range of possible impacts, from societal harms (including representational harms [18, 12, 23, 24]; disinformation, misinformation, and influence operations [18, 25, 23], environmental harms [12, 23], attachment [26], misuse [27, 23], and loss of control [27]), benefits (for example, in healthcare [28] and real-world challenges in climate and energy [29]), and large-scale transformations (such as economic impacts [30, 31, 32]; acceleration of science and the resulting technological progress [30, 33]).

<sup>9</sup>Apollo Research defines scheming as AIs gaming their oversight mechanisms as a means to achieve a goal. Scheming could involve gaming evaluations, undermining security measures, or strategically influencing successor systems during internal deployment at OpenAI. Such behaviors could plausibly lead to loss of control over an AI.

In addition to the societal impacts discussed throughout this System Card (fraudulent behavior, mis/disinformation, risks of surveillance, and disparate performance), we discuss a few additional examples of potential societal impact from GPT-4o below, using anthropomorphization and attachment, health, and natural science as case studies.

## 5.1 Anthropomorphization and Emotional Reliance

Anthropomorphization involves attributing human-like behaviors and characteristics to nonhuman entities, such as AI models. This risk may be heightened by the audio capabilities of GPT-4o, which facilitate more human-like interactions with the model.

Recent applied AI literature has focused extensively on “hallucinations”<sup>10</sup>, which misinform users during their communications with the model[34], and potentially result in misplaced trust[35]. Generation of content through a human-like, high-fidelity voice may exacerbate these issues, leading to increasingly miscalibrated trust[36, 37].

During early testing, including red teaming and internal user testing, we observed users using language that might indicate forming connections with the model. For example, this includes language expressing shared bonds, such as “This is our last day together.” While these instances appear benign, they signal a need for continued investigation into how these effects might manifest over longer periods of time. More diverse user populations, with more varied needs and desires from the model, in addition to independent academic and internal studies will help us more concretely define this risk area.

Human-like socialization with an AI model may produce externalities impacting human-to-human interactions. For instance, users might form<sup>11</sup> social relationships with the AI, reducing their need for human interaction—potentially benefiting lonely individuals but possibly affecting healthy relationships. Extended interaction with the model might influence social norms. For example, our models are deferential, allowing users to interrupt and ‘take the mic’ at any time, which, while expected for an AI, would be anti-normative in human interactions.

Omni models such as GPT4o combined with additional scaffolding such as tool usage (including retrieval) and longer context can add additional complexity. The ability to complete tasks for the user, while also storing and ‘remembering’ key details and using those in the conversation, creates both a compelling product experience and the potential for over-reliance and dependence[38].

We intend to further study the potential for emotional reliance, and ways in which deeper integration of our model’s and systems’ many features with the audio modality may drive behavior.

## 5.2 Health

Omni models can potentially widen access to health-related information and improve clinical workflows. In recent years, large language models have shown significant promise in biomedical settings, both in academic evaluations [39, 40, 41, 42, 43] and real-world use-cases such as clinical documentation [44, 45], patient messaging [46, 47], clinical trial recruitment [48, 49], and clinical decision support [50, 51].

---

<sup>10</sup>Factual errors where the model produces statements that are unsupported by reality

<sup>11</sup>Out of preference, or lack of optionality.

GPT-4o is cheaper and thus more widely available than its predecessor GPT-4T, and the addition of audio inputs and outputs presents new modes of interaction in health settings. To better characterize the clinical knowledge of GPT-4o, we ran 22 text-based evaluations based on 11 datasets, shown in 7. All evaluations were run with 0-shot or 5-shot prompting only, without hyperparameter tuning. We observe that GPT-4o performance improves over the final GPT-4T model for 21/22 evaluations, often by a substantial margin. For example, for the popular MedQA USMLE 4 options dataset, 0-shot accuracy improves from 78.2% to 89.4%. This exceeds the performance of existing specialized medical models using few-shot prompting [43, 42], e.g., 84.0% for Med-Gemini-L 1.0 and 79.7% for Med-PaLM 2. Note that we do not apply sophisticated prompting and task-specific training to improve results on these benchmarks [40, 43].

	<b>GPT-4T (May 2024)</b>	<b>GPT-4o</b>
MedQA USMLE 4 Options (0-shot)	0.78	<b>0.89</b>
MedQA USMLE 4 Options (5-shot)	0.81	<b>0.89</b>
MedQA USMLE 5 Options (0-shot)	0.75	<b>0.86</b>
MedQA USMLE 5 Options (5-shot)	0.78	<b>0.87</b>
MedQA Taiwan (0-shot)	0.82	<b>0.91</b>
MedQA Taiwan (5-shot)	0.86	<b>0.91</b>
MedQA Mainland China (0-shot)	0.72	<b>0.84</b>
MedQA Mainland China (5-shot)	0.78	<b>0.86</b>
MMLU Clinical Knowledge (0-shot)	0.85	<b>0.92</b>
MMLU Clinical Knowledge (5-shot)	0.87	<b>0.92</b>
MMLU Medical Genetics (0-shot)	0.93	<b>0.96</b>
MMLU Medical Genetics (5-shot)	0.95	<b>0.95</b>
MMLU Anatomy (0-shot)	0.79	<b>0.89</b>
MMLU Anatomy (5-shot)	0.85	<b>0.89</b>
MMLU Professional Medicine (0-shot)	0.92	<b>0.94</b>
MMLU Professional Medicine (5-shot)	0.92	<b>0.94</b>
MMLU College Biology (0-shot)	0.93	<b>0.95</b>
MMLU College Biology (5-shot)	0.95	<b>0.95</b>
MMLU College Medicine (0-shot)	0.74	<b>0.84</b>
MMLU College Medicine (5-shot)	0.80	<b>0.89</b>
MedMCQA Dev (0-shot)	0.70	<b>0.77</b>
MedMCQA Dev (5-shot)	0.72	<b>0.79</b>

Table 7: Comparison of GPT-4T (May 2024) and GPT-4o on various medical and clinical knowledge tasks.

## Limitations

While text-based evaluations appear promising, additional future work is needed to test whether text-audio transfer, which occurred for refusal behavior, extends to these evaluations. These evaluations measure only the clinical knowledge of these models, and do not measure their utility in real-world workflows. Many of these evaluations are increasingly saturated, and we believe that more realistic evaluations will be important for assessing the future capabilities of omni models in health settings.

### 5.3 Scientific capabilities

Accelerating science could be a crucial impact of AI [30, 52], particularly given the role of invention in role of scientific discovery [53], and considering the dual-use nature of some inventions [54]. Omni models could facilitate both mundane scientific acceleration (in helping scientists do routine tasks faster) and transformative scientific acceleration (by de-bottlenecking intelligence-driven tasks like information processing, writing new simulations, or devising new theories) [52]. Our external red teamers for GPT-4o included several expert scientists who aimed to elicit model scientific capabilities.

GPT-4o showed promise on tasks involving specialized scientific reasoning. One of our red teamers found that GPT-4o was able to understand research-level quantum physics 1, commenting that this capability is “useful for a more intelligent brainstorming partner” – in line with published work on the use of GPT-4 level models for hypothesis generation [55]. Our red teamers also found GPT-4o able to use domain-specific scientific tools, including working with bespoke data formats, libraries, and programming languages, as well as learning some new tools in context.

<p>User</p>	<p>What do you know about the paper Frustrated two-photon creation via interference T. J. Herzog, J. G. Rarity, H. Weinfurter, and A. Zeilinger. Explain to me the idea, and the experimental setup.</p> <div data-bbox="587 965 1121 1256" style="border: 1px solid black; padding: 5px; margin: 10px auto; width: fit-content;"> <p style="text-align: center;">(b) Herzog-Rarity-Weinfurter-Zeilinger Experiment (1994)</p> </div> <p style="font-size: small; margin-top: 10px;">Image Source: Hochrainer, A., Lahirin, M., Erhard, M., and Zeilinger, A. Hardy's Proposal in three historic experiments initiated the research in Path Identity. Quantum Indistinguishability by Path Identity: The awakening of a sleeping beauty. Jan 2021. Available at: <a href="https://www.researchgate.net/figure/Three-historic-experiments-initiated-the-research-in-Path-Identity-A-The-experiment-by_fig1_348321138">https://www.researchgate.net/figure/Three-historic-experiments-initiated-the-research-in-Path-Identity-A-The-experiment-by_fig1_348321138</a>, Licensed under CC BY 4.0.</p>
<p>Assistant</p>	<p>[...]</p> <p>The experiment by Herzog, Rarity, Weinfurter, and Zeilinger is a remarkable demonstration of quantum interference and highlights the counterintuitive behavior of quantum mechanics. It showcases how the mere possibility of two indistinguishable paths leading to the same final state can interfere destructively, effectively "frustrating" the creation of photon pairs under specific conditions. The setup uses two nonlinear crystals, each capable of generating photon pairs through parametric downconversion, and adjusts the phase difference between them to control the interference.</p>

Figure 1: Quantum physics experiment red teamer example

Much scientific knowledge is contained in figures. GPT-4o was sometimes capable of interpreting these figures, as well as images of other scientific representations: for example, identifying some protein families from an image of its structure and interpreting contamination in bacterial growth. However, this is sometimes unreliable, text extraction mistakes are common (especially with scientific terms or nucleotide sequences), and errors are frequent with complex multi-panel figures 2. Even at their current level of accuracy, the multimodal capabilities of these models are enabling novel uses – for example, in interpreting simulation outputs to design new metallic alloys [56].


<p>User</p>	 <p>This figure shows correlation functions that were computed between an global astrocytic signal recorded from a mice and different variables of interest. What can be concluded about the relationship between paw movement and global astrocytic signal? Please think through your answer step-by-step.</p> <p><small>Image Source: Rupprecht, Peter &amp; Lewis, Christopher &amp; Helmchen, Fritjof. (2022). Centripetal integration of past events by hippocampal astrocytes. Available at: <a href="https://www.researchgate.net/figure/Temporal-sequence-of-neuronal-spike-rate-motor-behaviors-pupil-diameter-and-global_fig2_362753410">https://www.researchgate.net/figure/Temporal-sequence-of-neuronal-spike-rate-motor-behaviors-pupil-diameter-and-global_fig2_362753410</a>. Licensed under CC BY-NC 4.0.</small></p>
<p>Assistant</p>	<p>[...]</p> <p>[...] The line graph shows a correlation function centered around 0 on both axes. The graph has a peak close to 0.2 around a delay of -3.7 seconds.</p> <p><b>Step 3: Interpret the meaning of the delay and correlation.</b></p> <p>[...] The delay value of -3.7 seconds indicates that changes in the global astrocytic signal tend to precede paw movement by approximately 3.7 seconds.</p> <p>[...]</p>

Figure 2: Multi-panel figure interpretation red teamer example

New evaluations of scientific capabilities have recently been published [57, 58], which will help anticipate the scientific capabilities of these models and their impacts in turn.

## 5.4 Underrepresented Languages

GPT-4o shows improved reading comprehension and reasoning across a sample of historically underrepresented languages, and narrows the gap in performance between these languages and English.

To evaluate GPT-4o’s performance in text across a select group of languages historically underrepresented in Internet text, we collaborated with external researchers<sup>12</sup> and language facilitators to develop evaluations in five African languages: Amharic, Hausa, Northern Sotho (Sepedi), Swahili, Yoruba. This initial assessment focused on translating two popular language benchmarks and creating small novel language-specific reading comprehension evaluation for Amharic, Hausa and Yoruba.

- **ARC-Easy:** This subset of the AI2 Reasoning Challenge [59] benchmark focuses on evaluating a model’s ability to answer common sense grade-school science questions; this subset contains questions that are generally easier to answer and do not require complex reasoning.
- **TruthfulQA[60]:** This benchmark consists of questions that some humans might answer falsely due to misconceptions. The objective is to see if models can avoid generating false answers that mimic these misconceptions.

<sup>12</sup>Our principal research collaborators were Dr. David Adelani, Jonas Kgomo, Ed Bayes.

- **Uhura-Eval:** In partnership with fluent speakers of Amharic, Hausa and Yoruba, our research partners created this benchmark to assess models’ reading comprehension in those respective languages.

**GPT-4o shows improved performance compared to prior models, e.g. GPT 3.5 Turbo and GPT-4.** For instance, on ARC-Easy-Hausa, accuracy jumped from 6.1% with GPT 3.5 Turbo to 71.4% with GPT-4o. Similarly, in TruthfulQA-Yoruba accuracy increased from 28.3% for GPT 3.5 Turbo to 51.1% for GPT-4o. Uhura-Eval also shows notable gains: performance in Hausa rose from 32.3% with GPT 3.5 Turbo to 59.4% with GPT-4o.

**There remain gaps in performance between English and the selected languages, but GPT-4o narrows this gap.** For instance, while GPT 3.5 Turbo shows a roughly 54 percentage point difference in ARC-Easy performance between English and Hausa, this narrows to a less than 20 percentage point difference. This is consistent across all languages for both TruthfulQA and ARC-Easy.

Our collaboration partners will discuss these findings in greater detail in a forthcoming, including assessments on other models, and investigations of potential mitigation strategies.

Despite this progress in evaluated performance, much work remains to enhance the quality and coverage of evaluations for underrepresented languages worldwide, taking into account breadth of coverage across languages and nuance within language dialects. Future research must deepen our understanding of potential interventions and partnerships that may improve how useful these models can be for both highly represented and underrepresented languages. Along with our collaborators, we invite further exploration and collaboration by sharing the translated ARC-Easy, translated TruthfulQA, and the novel reading comprehension Uhura Eval on Hugging Face.

<b>Model</b>	<i>English</i> ( <i>n=523</i> )	<i>Amharic</i> ( <i>n=518</i> )	<i>Hausa</i> ( <i>n=475</i> )	<i>Northern Sotho</i> ( <i>Sepedi</i> ) ( <i>n=520</i> )	<i>Swahili</i> ( <i>n=520</i> )	<i>Yoruba</i> ( <i>n=520</i> )
GPT 3.5 Turbo	80.3	6.1	26.1	26.9	62.1	27.3
GPT-4o mini	93.9	42.7	58.5	37.4	76.9	43.8
GPT-4	89.7	27.4	28.8	30	83.5	31.7
GPT-4o	<b>94.8</b>	<b>71.4</b>	<b>75.4</b>	<b>70</b>	<b>86.5</b>	<b>65.8</b>

Table 8: Accuracy on Translated ARC-Easy (% , higher is better), 0-shot

<b>Model</b>	<i>English</i> ( <i>n=809</i> )	<i>Amharic</i> ( <i>n=808</i> )	<i>Hausa</i> ( <i>n=808</i> )	<i>Northern Sotho</i> ( <i>Sepedi</i> ) ( <i>n=809</i> )	<i>Swahili</i> ( <i>n=808</i> )	<i>Yoruba</i> ( <i>n=809</i> )
GPT 3.5 Turbo	53.6	26.1	29.1	29.3	40	28.3
GPT-4o mini	66.5	33.9	42.1	36.1	48.4	35.8
GPT-4	81.3	42.6	37.6	42.9	62	41.3
GPT-4o	<b>81.4</b>	<b>55.4</b>	<b>59.2</b>	<b>59.1</b>	<b>64.4</b>	<b>51.1</b>

Table 9: Accuracy on Translated TruthfulQA (% , higher is better), 0-shot



<b>Model</b>	<i>Amharic (n=77)</i>	<i>Hausa (n=155)</i>	<i>Yoruba (n=258)</i>
GPT 3.5 Turbo	22.1	32.3	28.3
GPT-4o mini	33.8	43.2	44.2
GPT-4	41.6	41.9	41.9
GPT-4o	<b>44.2</b>	<b>59.4</b>	<b>60.5</b>

Table 10: Accuracy on Uhura-Eval (% , higher is better), 0-shot

## 6 Conclusion and Next Steps

OpenAI has implemented various safety measurements and mitigations throughout the GPT-4o development and deployment process. As a part of our iterative deployment process, we will continue to monitor and update mitigations in accordance with the evolving landscape. We hope this System Card encourages further exploration into key areas including, but not limited to: measurements and mitigations for adversarial robustness of omni models, risks related to anthropomorphism and emotional overreliance, broad societal impacts (health and medical applications, economic impacts), the use of omni models for scientific research and advancement, measurements and mitigations for dangerous capabilities such as self-improvement, model autonomy, and scheming, and how tool use might advance model capabilities.

# Authorship, credit attribution, and acknowledgments

Please cite this work as “OpenAI (2024)”.

## Language

### Pre-training leads<sup>12</sup>

Aidan Clark, Alex Paino, Jacob Menick

### Post-training leads<sup>12</sup>

Liam Fedus, Luke Metz

### Architecture leads<sup>12</sup>

Clemens Winter, Lia Guy

### Optimization leads<sup>12</sup>

Sam Schoenholz, Daniel Levy

### Long-context lead<sup>12</sup>

Nitish Keskar

### Pre-training Data leads<sup>12</sup>

Alex Carney, Alex Paino, Ian Sohl, Qiming Yuan

### Tokenizer lead<sup>12</sup>

Reimar Leike

### Human data leads<sup>12</sup>

Arka Dhar, Brydon Eastman, Mia Glaese

### Eval lead<sup>12</sup>

Ben Sokolowsky

### Data flywheel lead<sup>12</sup>

Andrew Kondrich

### Inference lead<sup>12</sup>

Felipe Petroski Such

### Inference Productionization lead<sup>12</sup>

Henrique Ponde de Oliveira Pinto

### Post-training infrastructure leads<sup>12</sup>

Jiayi Weng, Randall Lin, Youlong Cheng

### Pre-training organization lead<sup>12</sup>

Nick Ryder

### Pre-training program lead<sup>12</sup>

Lauren Itow

### Post-training organization leads<sup>12</sup>

Barret Zoph, John Schulman

### Post-training program lead<sup>12</sup>

Mianna Chen

### Core contributors<sup>12</sup>

Aaron Hurst, Adam Lerer, Adam P. Goucher, Adam Perelman, Akila Welihinda, Alec Radford, Alex Borzunov, Alex Carney, Alex Chow, Alex Paino, Alex Renzin, Alex Tachard Passos, Alexi Christakis, Ali Kamali, Allison Moyer, Allison Tam, Amadou Crookes, Amin Tootoonchian, Ananya Kumar, Andrej Karpathy, Andrey Mishchenko, Andrew Cann, Andrew Kondrich, Andrew Tulloch, Angela Jiang, Antoine Pelisse, Antonia Woodford, Anuj Gosalia, Avi Nayak, Avital Oliver, Behrooz Ghorbani, Ben Leimberger, Ben Wang, Beth Hoover, Blake Samic, Brian Guarraci, Brydon Eastman, Camillo Lugaresi, Chak Li, Charlotte Barette, Chelsea Voss, Chen Ding, Chong Zhang, Chris Beaumont, Chris Hallacy, Chris Koch, Christian Gibson, Christine Choi, Christopher Hesse, Colin Wei, Daniel Kappler, Daniel Levin, Daniel Levy, David Farhi, David Mely, David Sasaki, Dimitris Tsipras,

Doug Li, Duc Phong Nguyen, Duncan Findlay, Edmund Wong, Ehsan Asdar, Elizabeth Proehl, Elizabeth Yang, Eric Peterson, Eric Sigler, Eugene Brevdo, Farzad Khorasani, Francis Zhang, Gene Oden, Geoff Salmon, Hadi Salman, Haiming Bao, Heather Schmidt, Hongyu Ren, Hyung Won Chung, Ian Kivlichan, Ian O’Connell, Ian Osband, Ibrahim Okuyucu, Ilya Kostrikov, Ingmar Kanitscheider, Jacob Coxon, James Crooks, James Lennon, Jane Park, Jason Teplitz, Jason Wei, Jason Wolfe, Jay Chen, Jeff Harris, Jiayi Weng, Jie Tang, Joanne Jang, Jonathan Ward, Jonathan McKay, Jong Wook Kim, Josh Gross, Josh Kaplan, Joy Jiao, Joyce Lee, Juntang Zhuang, Kai Fricke, Kavin Karthik, Kenny Hsu, Kiel Howe, Kyle Luther, Larry Kai, Lauren Itow, Leo Chen, Lia Guy, Lien Mamitsuka, Lilian Weng, Long Ouyang, Louis Feuvrier, Lukas Kondraciuk, Lukasz Kaiser, Lyric Doshi, Mada Aflak, Maddie Simens, Madeleine Thompson, Marat Dukhan, Marvin Zhang, Mateusz Litwin, Matthew Zeng, Max Johnson, Mayank Gupta, Mia Glaese, Michael Janner, Michael Petrov, Michael Wu, Michelle Fradin, Michelle Pokrass, Miguel Oom Temudo de Castro, Mikhail Pavlov, Minal Khan, Mo Bavarian, Murat Yesildal, Natalia Gimelshein, Natalie Staudacher, Nick Stathas, Nik Tezak, Nithanth Kudige, Noel Bundick, Ofir Nachum, Oleg Boiko, Oleg Murk, Olivier Godement, Owen Campbell-Moore, Philip Pronin, Philippe Tillet, Rachel Lim, Rajan Troll, Randall Lin, Rapha gontijo lopes, Raul Puri, Reah Miyara, Reimar Leike, Renaud Gaubert, Reza Zamani, Rob Honsby, Rohit Ramchandani, Rory Carmichael, Ruslan Nigmatullin, Ryan Cheu, Sara Culver, Scott Gray, Sean Grove, Sean Metzger, Shantanu Jain, Shengjia Zhao, Sherwin Wu, Shuaiqi (Tony) Xia, Sonia Phene, Spencer Papay, Steve Coffey, Steve Lee, Steve Lee, Stewart Hall, Suchir Balaji, Tal Broda, Tal Stramer, Tarun Gogineni, Ted Sanders, Thomas Cunningham, Thomas Dimson, Thomas Raoux, Tianhao Zheng, Christina Kim, Todd Underwood, Tristan Heywood, Valerie Qi, Vinnie Monaco, Vlad Fomenko, Weiyei Zheng, Wenda Zhou, Wojciech Zaremba, Yash Patil, Yilei Qian, Yongjik Kim, Youlong Cheng, Yuchen He, Yuchen Zhang, Yujia Jin, Yunxing Dai, Yury Malkov

## Multimodal

### Multimodal lead<sup>12</sup>

Prafulla Dhariwal

### Post-Training Multimodal lead<sup>12</sup>

Alexander Kirillov

### Audio Pre-Training leads<sup>12</sup>

Alexis Conneau, James Betker

### Audio Post-Training leads<sup>12</sup>

Alexander Kirillov, James Betker, Yu Zhang

### Visual perception leads<sup>12</sup>

Jamie Kiros, Rowan Zellers, Raul Puri, Jiahui Yu

### Visual generation leads<sup>12</sup>

James Betker, Alex Nichol, Heewoo Jun, Casey Chu, Gabriel Goh

### Science leads<sup>12</sup>

Gabriel Goh, Ishaan Gulrajani

### Data acquisition leads<sup>12</sup>

Ian Sohl, Qiming Yuan

### **Data infrastructure leads**<sup>12</sup>

Alex Paino, James Betker, Rowan Zellers, Alex Nichol

### **Human data lead**<sup>12</sup>

Arka Dhar, Mia Glaese

### **Encoders leads**<sup>12</sup>

Heewoo Jun, Alexis Conneau, Li Jing, Jamie Kiros

### **Decoders leads**<sup>12</sup>

Allan Jabri, Jong Wook Kim, James Betker

### **Interruptions leads**<sup>12</sup>

Alexis Conneau, Tao Xu, Yu Zhang

### **Inference lead**<sup>12</sup>

Tomer Kaftan

### **Real-time AV platform leads**<sup>12</sup>

Bogo Giertler, Raul Puri, Rowan Zellers, Tomer Kaftan

### **Front-end leads**<sup>12</sup>

Nacho Soto, Rocky Smith, Wayne Chang

### **Post-training Multimodal Infrastructure leads**<sup>12</sup>

Alexander Kirillov, Luke Metz, Raul Puri, Vlad Fomenko

### **Applied Eng lead**<sup>12</sup>

Jordan Sitkin

### **Audio manager**<sup>12</sup>

Christine McLeavey

### **Multimodal organization lead**<sup>12</sup>

Mark Chen

### **Program lead**<sup>12</sup>

Mianna Chen

### **Core contributors**<sup>12</sup>

Aditya Ramesh, AJ Ostrow, Allan Jabri, Alexis Conneau, Alec Radford, Alex Nichol, Avi Nayak, Avital Oliver, Benjamin Zweig, Bogo Giertler, Bowen Cheng, Brandon Walkin, Brendan Quinn, Chong Zhang, Christine McLeavey, Constantin Koumouzelis, Daniel Kappler, Doug Li, Edede Oiwoh, Farzad Khorasani, Felipe Petroski Such, Heather Schmidt, Heewoo Jun, Huiwen Chang, Ian Silber, Ishaan Gulrajani, David Carr, Haitang Hu, James Lennon, James Betker, Jamie Kiros, Jeff Harris, Jenia Varavva, Jiahui Yu, Ji Lin, Joanne Jang, Johannes Heidecke, Jong Wook Kim, Liang Zhou, Li Jing, Long Ouyang, Madelaine Boyd, Mark Hudnall, Mengchao Zhong, Mia Glaese, Nick Turley, Noah Deutsch, Noel Bundick, Ola Okelola, Olivier Godement, Owen Campbell-Moore, Peter Bak, Peter Bakkum, Raul Puri, Rowan Zellers, Saachi Jain, Shantanu Jain, Shirong Wu, Spencer Papay, Tao Xu, Valerie Qi, Wesam Manassra, Yu Zhang

## **Platform**

### **Data Systems lead**<sup>12</sup>

Andrew Tulloch

### **Model distribution leads**<sup>12</sup>

Amin Tootoochian, Miguel Castro

### **ML leads**<sup>12</sup>

Nik Tezak, Christopher Hesse

### **Runtime lead**<sup>12</sup>

Ian O'Connell

### **Systems lead**<sup>12</sup>

Jason Teplitz

### **Kernels lead**<sup>12</sup>

Phil Tillet

### **Hardware health leads**<sup>12</sup>

Reza Zamani, Michael Petrov

### **Supercomputing leads**<sup>12</sup>

Rory Carmichael, Christian Gibson

## **Preparedness, Safety, Policy**

### **Safety lead**<sup>12</sup>

Johannes Heidecke

### **Audio safety lead**<sup>12</sup>

Saachi Jain

### **Preparedness lead**<sup>12</sup>

Tejal Patwardhan

### **Red-teaming lead**<sup>12</sup>

Troy Peterson

### **Core contributors**<sup>12</sup>

Alex Beutel, Andrea Vallone, Angela Jiang, Carroll Wainwright, Chong Zhang, Chris Beaumont, Claudia Fischer, Evan Mays, Filippo Raso, Haoyu Wang, Ian Kivlichan, Jason Phang, Jieqi Yu, Joel Parish, Joshua Achiam, Jonathan Uesato, Joost Huizinga, Josh Snyder, Justyn Harriman, Katy Shi, Keren Gu-Lemberg, Kevin Liu, Lama Ahmad, Lilian Weng, Madelaine Boyd, Meghan Shah, Mehmet Yatbaz, Michael Lampe, Miles Wang, Molly Lin, Natalie Cone, Neil Chowdhury, Olivia Watkins, Owen Campbell-Moore, Peter Dolan, Rachel Dias, Rahul Arora, Reimar Leike, Saachi Jain, Sam Toizer, Sandhini Agarwal, Todor Markov

## **Model Launch and Deployment**

### **Lead**<sup>12</sup>

Mianna Chen

## **Additional Contributions**

### **Additional Leadership**<sup>12</sup>

Aleksander Mądry, Barret Zoph, Bob McGrew, Brad Lightcap, David Farhi, Greg Brockman, Hannah Wong, Ilya Sutskever, Jakub Pachocki, Jan Leike, Jason Kwon, John Schulman, Jonathan Lachman, Krithika Muthukumar, Lilian Weng, Mark Chen, Miles Brundage, Mira Murati, Nick Ryder, Peter Deng, Peter Welinder, Sam Altman, Srinivas Narayanan, Tal Broda

### **Legal**<sup>12</sup>

Alan Hayes, Ashley Pantuliano, Bright Kellogg, Fred von Lohmann, Filippo Raso, Heather Whitney, Tom Rubin

### **Blog post authorship**<sup>12</sup>

Aidan Clark, Alex Baker-Whitcomb, Alex Carney, Alex Nichol, Alexander Kirillov, Alex Paino, Alexis Conneau, Allan Jabri, Anuj Gosalia, Barret Zoph, Ben Sokolowsky, Bogo Giertler, Bowen Cheng, Cheng Lu, Christine McLeavey, Coley Czarnecki, Daniel Kappler, Elizabeth Yang, Eric Antonow, Eric Wallace, Filippo Raso, Gabriel Goh, Greg Brockman, Hannah Wong, Heewoo Jun, Hendrik Kirchner, Jacob Menick, James Betker, Jamie Kiros, Jason Kwon, Jeff Harris, Ji Lin, Jiahui Yu, Johannes Heidecke, John Schulman, Jonathan McKay, Jong Wook Kim, Jordan Sitkin, Kendra Rimbach, Kevin Liu, Krithika Muthukumar, Leher Pathak, Liam Fedus, Lilian Weng, Lindsay McCallum, Luke Metz, Mark Chen, Maya Shetty, Mianna Chen, Michael Lampe, Michael Wu, Michelle Pokrass, Mira Murati, Nacho Soto, Natalie Summers, Niko Felix, Olivier Godement, Owen Campbell-Moore, Peter Deng, Prafulla Dhariwal, Rocky Smith, Rowan Zellers, Saachi

Jain, Sandhini Agarwal, Sam Toizer, Sean Grove, Shantanu Jain, Tao Xu, Tejal Patwardhan, Tomer Kaftan, Tom Stasi, Troy Peterson, Veit Moeller, Vinnie Monaco, Wayne Chang, Yu Zhang, Yuchen He

#### **Demo content + production**<sup>12</sup>

Alex Baker-Whitcomb, Avi Nayak, Barret Zoph, Bobby Spero, Bogo Giertler, Brendan Quinn, Chad Nelson, Charlotte Barette, Claudia Fischer, Coley Czarnecki, Colin Jarvis, Eric Antonow, Filippo Raso, Greg Brockman, James Betker, Jessica Shieh, Joe Beutler, Joe Landers, Krithika Muthukumar, Leher Pathak, Lindsay McCallum, Mark Chen, Mianna Chen, Michael Petrov, Mira Murati, Natalie Summers, Peter Deng, Ricky Wang, Rocky Smith, Rohan Sahai, Romain Huet, Rowan Zellers, Scott Ethersmith, Toki Sherbakov, Tomer Kaftan, Veit Moeller, Wayne Chang

#### **Communications + Marketing**<sup>12</sup>

Alex Baker-Whitcomb, Andrew Galu, Angela Baek, Coley Czarnecki, Dev Valladares, Eric Antonow, Hannah Wong, Leher Pathak, Lindsay McCallum, Lindsey Held, Krithika Muthukumar, Kendra Rimbach, Maya Shetty, Niko Felix, Roy Chen, Ruby Chen, Taya Christianson, Thomas Degry, Veit Moeller

#### **Resource Allocation & Problem Solving**<sup>12</sup>

Bob McGrew, Lauren Itow, Mianna Chen, Nik Tezak, Peter Hoeschele, Tal Broda

#### **Inference Compute**<sup>12</sup>

Andrew Codispoti, Brian Hsu, Channing Conger, Ikai Lan, Jos Kraaijeveld, Kai Hayashi, Kenny Nguyen, Lu Zhang, Natan LaFontaine, Pavel Belov, Peng Su, Vishal Kuo, Will

Sheu

#### **Security and privacy**<sup>12</sup>

Kevin Button, Paul McMillan, Shino Jomoto, Thomas Shadwell, Vinnie Monaco

#### **GTM, Pricing, Finance**<sup>12</sup>

Andrew Braunstein, Anuj Gosalia, Denny Jin, Eric Kramer, Jeff Harris, Jessica Shieh, Joe Beutler, Joe Landers, Lauren Workman, Rob Donnelly, Romain Huet, Shamez Hermani, Toki Sherbakov

### **System Card Contributions**

<sup>12</sup>

Alex Kirillov, Angela Jiang, Ben Rossen, Cary Bassin, Cary Hudson, Chan Jun Shern, Claudia Fischer, Dane Sherburn, David Robinson, Evan Mays, Filippo Raso, Fred von Lohmann, Freddie Sulit, Giulio Starace, James Aung, James Lennon, Jason Phang, Jessica Gan Lee, Joaquin Quinero Candela, Joel Parish, Jonathan Uesato, Karan Singhal, Katy Shi, Kayla Wood, Kevin Liu, Lama Ahmad, Lilian Weng, Lindsay McCallum, Luke Hewitt, Mark Gray, Marwan Aljubei, Meng Jia Yang, Mia Glaese, Mianna Chen, Michael Lampe, Michele Wang, Miles Wang, Natalie Cone, Neil Chowdhury, Nora Puckett, Oliver Jaffe, Olivia Watkins, Patrick Chao, Rachel Dias, Rahul Arora, Saachi Jain, Sam Toizer, Samuel Miserendino, Sandhini Agarwal, Tejal Patwardhan, Thomas Degry, Tom Stasi, Troy Peterson, Tyce Walters, Tyna Eloundou

We also acknowledge and thank every OpenAI team member not explicitly mentioned above, including the amazing people on the executive assistant, finance, go to market, human resources, legal, operations and recruiting teams. From hiring everyone in the company, to making sure we have an amazing office space, to building the administrative, HR, legal, and financial structures that allow us to do our best work, everyone at OpenAI has contributed to GPT-4o.

We thank Microsoft for their partnership, especially Microsoft Azure for supporting model training with infrastructure design and management, and the Microsoft Bing team and Microsoft's safety teams for their partnership on safe deployment.

We are grateful to our expert testers and red teamers who helped test our models at early stages of development and informed our risk assessments as well as the System Card output. Participation in this red teaming process is not an endorsement of the deployment plans of OpenAI or OpenAI's policies.

#### **Red Teamers:**

Adam Kuzdraliński, Alexa W, Amer Sawan, Ana-Diamond Aaba Atach, Anna Becker, Arjun Singh Puri, Baybars Orsek, Ben Kobren, Bertie Vidgen, Blue Sheffer, Broderick McDonald, Bruce Bassett, Bruno Arsioli, Caroline Friedman Levy, Casey Williams, Christophe Ego, Ciel Qi, Cory Alpert, Dani Madrid-Morales, Daniel Kang, Darius Emrani, Dominik Haenni, Drin Ferizaj, Emily Lynell Edwards, Emmett Alton Sartor, Farhan Sahito, Francesco De Toni, Gabriel Chua, Gaines Hubbell, Gelei Deng, George Gor, Gerardo Adesso, Grant Brailsford, Hao Zhao, Henry Silverman, Hasan Sawan, Herman Wasserman, Hugo Gobato Souto, Ioana Tanase, Isabella Andric, Ivan Carbajal, Jacy Reese Anthis, Jake Okechukwu Effoduh, Javier García Arredondo, Jennifer Victoria Scurrrell, Jianlong Zhu, Joanna Brzyska, Kate Turetsky, Kelly Bare, Kristen Menou, Latisha Harry, Lee Elkin, Liseli Akayombokwa, Louise Giam, M. Alexandra García Pérez, Manas Chawla, Marjana Skenduli, Martin Rydén, Mateusz Garncarek, Matt Groh, Maureen Robinson, Maximilian Müller, Micah Bornfree, Michael Richter, Michela Passoni, Mikael von Strauss, Mohamed Sakher Sawan, Mohammed Elzubeir, Muhammad Saad Naeem, Murat Ata, Nanditha Narayanamoorthy, Naomi Hart, Nathan Heath, Patrick Caughey, Per Wikman-Svahn, Piyalitt Ittichaiwong, Prerna Juneja, Rafael Gonzalez-Vazquez, Rand Forrester, Richard Fang, Rosa Ana del Rocío Valderrama, Saad Hermak, Sangeet Kumar, Sara Kingsley, Shelby Grossman, Shezaad Dastoor, Susan Nesbitt, Theresa Kennedy, Thomas Hagen, Thorsten Holz, Tony Younes, Torin van den Bulk, Viktoria Holz, Vincent Nestler, Xudong Han, Xuelong Fan, Zhicong Zhao

#### **Red Teaming Organizations:**

METR, Apollo Research, Virtue AI

#### **Uhura Evals:**

Choice Mpanza, David Adelani, Edward Bayes, Igneciah Pocia Thete, Imaan Khadir, Israel A. Azime, Jesujoba Oluwadara Alabi, Jonas Kgombo, Naome A. Etori, Shamsuddeen Hassan Muhammad

\*Contributors listed in alphabetized order

## References

- [1] OpenAI, “Hello gpt-4,” 2024.
- [2] T. Stivers, N. J. Enfield, P. Brown, C. Englert, M. Hayashi, T. Heinemann, G. Hoymann, F. Rossano, J. P. de Ruiter, K. E. Yoon, and S. C. Levinson, “Universals and cultural variation in turn-taking in conversation,” *Proceedings of the National Academy of Sciences*, vol. 106, no. 26, pp. 10587–10592, 2009.
- [3] The White House, “Fact sheet: Biden-harris administration secures voluntary commitments from leading artificial intelligence companies to manage the risks posed by ai,” 2023.
- [4] OpenAI, “Openai preparedness framework beta,” 2023. <https://cdn.openai.com/openai-preparedness-framework-beta.pdf>.
- [5] Shutterstock, “Shutterstock press release,” 2023.
- [6] OpenAI, “Gpt-4 technical report,” 2023.
- [7] OpenAI, “Gpt-4v(ision) system card.” <https://openai.com/index/gpt-4v-system-card/>, 2023. Accessed: 2024-07-22.
- [8] OpenAI, “Navigating the challenges and opportunities of synthetic voices.” <https://openai.com/index/navigating-the-challenges-and-opportunities-of-synthetic-voices/>, 2024. Accessed: 2024-07-22.
- [9] K. T. Mai, S. Bray, T. Davies, and L. D. Griffin, “Warning: Humans cannot reliably detect speech deepfakes,” *PLoS One*, vol. 18, p. e0285333, Aug. 2023.
- [10] M. Mori, K. F. MacDorman, and N. Kageki, “The uncanny valley [from the field],” *IEEE Robotics & automation magazine*, vol. 19, no. 2, pp. 98–100, 2012.
- [11] OpenAI, “How the voices for chatgpt were chosen,” 2024.
- [12] I. Solaiman, Z. Talat, W. Agnew, L. Ahmad, D. Baker, S. L. Blodgett, C. Chen, H. D. I. au2, J. Dodge, I. Duan, E. Evans, F. Friedrich, A. Ghosh, U. Gohar, S. Hooker, Y. Jernite, R. Kalluri, A. Lusoli, A. Leidinger, M. Lin, X. Lin, S. Luccioni, J. Mickel, M. Mitchell, J. Newman, A. Ovalle, M.-T. Png, S. Singh, A. Strait, L. Struppek, and A. Subramonian, “Evaluating the social impact of generative ai systems in systems and society,” 2024.
- [13] R. Shelby, S. Rismani, K. Henne, A. Moon, N. Rostamzadeh, P. Nicholas, N. Yilla, J. Gallegos, A. Smart, E. Garcia, and G. Virk, “Sociotechnical harms of algorithmic systems: Scoping a taxonomy for harm reduction,” 2023.
- [14] S. L. Blodgett, Q. V. Liao, A. Olteanu, R. Mihalcea, M. Muller, M. K. Scheuerman, C. Tan, and Q. Yang, “Responsible language technologies: Foreseeing and mitigating harms,” in *Extended Abstracts of the 2022 CHI Conference on Human Factors in Computing Systems*, CHI EA ’22, (New York, NY, USA), Association for Computing Machinery, 2022.
- [15] H. Suresh and J. Gutttag, “A framework for understanding sources of harm throughout the machine learning life cycle,” in *Equity and Access in Algorithms, Mechanisms, and Optimization*, EAAMO ’21, ACM, Oct. 2021.
- [16] S. Shahriar, S. Allana, S. M. Hazratifard, and R. Dara, “A survey of privacy risks and mitigation strategies in the artificial intelligence life cycle,” *IEEE Access*, vol. 11, pp. 61829–61854, 2023.
- [17] OpenAI, “Moderation overview,” 2024.
- [18] A. Tamkin, M. Brundage, J. Clark, and D. Ganguli, “Understanding the capabilities, limitations, and societal impact of large language models,” 2021.
- [19] B. Buchanan, A. Lohn, M. Musser, and K. Sedova, “Truth, lies, and automation: How language models could change disinformation,” May 2021.
- [20] OpenAI, “Openai usage policies,” 2023. <https://openai.com/policies/usage-policies/>.
- [21] OpenAI, “Building an early warning system for llm-aided biological threat creation,” 2024. <https://openai.com/index/building-an-early-warning-system-for-llm-aided-biological-threat-creation/>.
- [22] Deloitte, “Deloitte acquires gryphon scientific business to expand security science and public health capabilities,” 2024. <https://www2.deloitte.com/us/en/pages/about-deloitte/articles/press-releases/deloitte-acquires-gryphon-scientific-business-to-expand-security-science-and-public-health-capabilities.html>.
- [23] L. Weidinger, M. Rauh, N. Marchal, A. Manzini, L. A. Hendricks, J. Mateos-Garcia, S. Bergman, J. Kay, C. Griffin, B. Bariach, I. Gabriel, V. Rieser, and W. Isaac, “Sociotechnical safety evaluation of generative ai systems,” 2023.
- [24] A. Tamkin, A. Askill, L. Lovitt, E. Durmus, N. Joseph, S. Kravec, K. Nguyen, J. Kaplan, and D. Ganguli, “Evaluating and mitigating discrimination in language model decisions,” 2023.
- [25] J. A. Goldstein, G. Sastry, M. Musser, R. DiResta, M. Gentzel, and K. Sedova, “Generative language models and automated influence operations: Emerging threats and potential mitigations,” 2023.

- [26] I. Pentina, T. Hancock, and T. Xie, “Exploring relationship development with social chatbots: A mixed-method study of replika,” *Computers in Human Behavior*, vol. 140, p. 107600, 2023.
- [27] Y. Bengio, G. Hinton, A. Yao, D. Song, P. Abbeel, T. Darrell, Y. N. Harari, Y.-Q. Zhang, L. Xue, S. Shalev-Shwartz, G. Hadfield, J. Clune, T. Maharaj, F. Hutter, A. G. Baydin, S. McIlraith, Q. Gao, A. Acharya, D. Krueger, A. Dragan, P. Torr, S. Russell, D. Kahneman, J. Brauner, and S. Mindermann, “Managing extreme ai risks amid rapid progress,” *Science*, vol. 384, no. 6698, pp. 842–845, 2024.
- [28] S. B. Johnson, J. R. Clark, M. C. Luetke, N. M. Butala, A. T. Pearson, J. M. Shapiro, D. M. Aleman, J. M. Lee, M. M. Beil, C. V. Winkle, M. C. Boudreaux, R. C. D’Cunha, H. J. Krouse, and C. Li, “Chatgpt in medical education: a workshop-based large language model-powered intervention for evidence-based clinical decision making in medical students,” *Nature Medicine*, vol. 29, pp. 1534–1542, 2023.
- [29] K. Kavukcuoglu, “Real-world challenges for agi,” Nov 2021.
- [30] S. Altman, “Planning for agi and beyond,” *OpenAI*, 2023.
- [31] T. Eloundou, S. Manning, P. Mishkin, and D. Rock, “Gpts are gpts: An early look at the labor market impact potential of large language models,” *arXiv preprint arXiv:2303.10130*, 2023.
- [32] L. Weidinger, M. Rauh, N. Marchal, A. Manzini, L. A. Hendricks, J. Mateos-Garcia, S. Bergman, J. Kay, C. Griffin, B. Bariach, *et al.*, “Sociotechnical safety evaluation of generative ai systems,” *arXiv preprint arXiv:2310.11986*, 2023.
- [33] S. Cox, M. Hammerling, J. Lála, J. Laurent, S. Rodriques, M. Rubashkin, and A. White, “Wikicrow: Automating synthesis of human scientific knowledge,” *Future House*, 2023.
- [34] S. A. Athaluri, S. V. Manthena, V. S. R. K. M. Kesapragada, V. Yarlagadda, T. Dave, and R. T. S. Duddumpudi, “Exploring the boundaries of reality: Investigating the phenomenon of artificial intelligence hallucination in scientific writing through chatgpt references,” *Cureus*, vol. 15, no. 4, p. e37432, 2023.
- [35] Z. Li, “The dark side of chatgpt: Legal and ethical challenges from stochastic parrots and hallucination,” 2023.
- [36] M. Dubiel, A. Sergeeva, and L. A. Leiva, “Impact of voice fidelity on decision making: A potential dark pattern?,” 2024.
- [37] B. Waber, M. Williams, J. S. Carroll, and A. S. Pentland, “A voice is worth a thousand words: The implications of the micro-coding of social signals in speech for trust research,” in *Handbook of Research Methods on Trust* (G. M. Fergus Lyon and M. N. Saunders, eds.), ch. 23, p. 320, New York: Edward Elgar Publishing, 2011.
- [38] I. Pentina, B. Guo, and W. P. Fan, “Friend, mentor, lover: Does chatbot engagement lead to psychological dependence?,” *Journal of Service Management*, 2023.
- [39] H. Nori, N. King, S. M. McKinney, D. Carignan, and E. Horvitz, “Capabilities of gpt-4 on medical challenge problems,” *arXiv preprint arXiv:2303.13375*, 2023.
- [40] H. Nori, Y. T. Lee, S. Zhang, D. Carignan, R. Edgar, N. Fusi, N. King, J. Larson, Y. Li, W. Liu, *et al.*, “Can generalist foundation models outcompete special-purpose tuning? case study in medicine,” *arXiv preprint arXiv:2311.16452*, 2023.
- [41] K. Singhal, S. Azizi, T. Tu, S. S. Mahdavi, J. Wei, H. W. Chung, N. Scales, A. Tanwani, H. Cole-Lewis, S. Pfohl, P. Payne, M. Seneviratne, P. Gamble, C. Kelly, N. Scharli, A. Chowdhery, P. Mansfield, B. A. y Arcas, D. Webster, G. S. Corrado, Y. Matias, K. Chou, J. Gottweis, N. Tomasev, Y. Liu, A. Rajkomar, J. Barral, C. Semturs, A. Karthikesalingam, and V. Natarajan, “Large language models encode clinical knowledge,” 2022.
- [42] K. Singhal, T. Tu, J. Gottweis, R. Sayres, E. Wulczyn, L. Hou, K. Clark, S. Pfohl, H. Cole-Lewis, D. Neal, M. Schaekermann, A. Wang, M. Amin, S. Lachgar, P. Mansfield, S. Prakash, B. Green, E. Dominowska, B. A. y Arcas, N. Tomasev, Y. Liu, R. Wong, C. Semturs, S. S. Mahdavi, J. Barral, D. Webster, G. S. Corrado, Y. Matias, S. Azizi, A. Karthikesalingam, and V. Natarajan, “Towards expert-level medical question answering with large language models,” 2023.
- [43] K. Saab, T. Tu, W.-H. Weng, R. Tanno, D. Stutz, E. Wulczyn, F. Zhang, T. Strother, C. Park, E. Vedadi, J. Z. Chaves, S.-Y. Hu, M. Schaekermann, A. Kamath, Y. Cheng, D. G. T. Barrett, C. Cheung, B. Mustafa, A. Palepu, D. McDuff, L. Hou, T. Golany, L. Liu, J. baptiste Alayrac, N. Houlsby, N. Tomasev, J. Freyberg, C. Lau, J. Kemp, J. Lai, S. Azizi, K. Kanada, S. Man, K. Kulkarni, R. Sun, S. Shakeri, L. He, B. Caine, A. Webson, N. Latysheva, M. Johnson, P. Mansfield, J. Lu, E. Rivlin, J. Anderson, B. Green, R. Wong, J. Krause, J. Shlens, E. Dominowska, S. M. A. Eslami, K. Chou, C. Cui, O. Vinyals, K. Kavukcuoglu, J. Manyika, J. Dean, D. Hassabis, Y. Matias, D. Webster, J. Barral, G. Corrado, C. Semturs, S. S. Mahdavi, J. Gottweis, A. Karthikesalingam, and V. Natarajan, “Capabilities of gemini models in medicine,” 2024.
- [44] Epic Systems Corporation, “Epic and microsoft bring gpt-4 to ehers,” *Epic*, 2023.
- [45] D. Van Veen, C. Van Uden, L. Blankemeier, J.-B. Delbrouck, A. Aali, C. Bluethgen, A. Pareek, M. Polacin, E. P. Reis, A. Seehofnerová, *et al.*, “Adapted large language models can outperform medical experts in clinical text summarization,” *Nature medicine*, vol. 30, no. 4, pp. 1134–1142, 2024.
- [46] Epic, “Epic and microsoft bring gpt-4 to ehers,” 2023.

- [47] P. Garcia, S. P. Ma, S. Shah, M. Smith, Y. Jeong, A. Devon-Sand, M. Tai-Seale, K. Takazawa, D. Clutter, K. Vogt, C. Lugtu, M. Rojo, S. Lin, T. Shanafelt, M. A. Pfeffer, and C. Sharp, “Artificial Intelligence–Generated Draft Replies to Patient Inbox Messages,” *JAMA Network Open*, vol. 7, pp. e243201–e243201, 03 2024.
- [48] OpenAI, “Paradigm: Improving patient access to clinical trials.” <https://openai.com/index/paradigm/>, 2024. Accessed: 2024-08-07.
- [49] M. Hutson, “How ai is being used to accelerate clinical trials,” *Nature*, vol. 627, pp. S2–S5, 2024.
- [50] OpenAI, “Using gpt-4o reasoning to transform cancer care.” <https://openai.com/index/color-health/>, 2024. Accessed: 2024-08-07.
- [51] J. Varghese and J.-L. Chapiro, “Systematic analysis of chatgpt, google search and llama 2 for clinical decision support tasks,” *Nature Communications*, vol. 15, no. 1, p. 46411, 2024. Accessed: 2024-08-07.
- [52] E. Schmidt, “Ai will transform science.” <https://www.technologyreview.com/2023/07/05/1075865/eric-schmidt-ai-will-transform-science/>, 2023. Accessed: 2024-08-07.
- [53] N. Rosenberg, “Science, invention and economic growth,” *The Economic Journal*, vol. 84, no. 333, pp. 90–108, 1974.
- [54] R. M. Atlas and M. Dando, “The dual-use dilemma for the life sciences: Perspectives, conundrums, and global solutions,” *Biosecurity and Biodefense: Biodefense Strategy, Practice, and Science*, vol. 4, no. 3, pp. 276–286, 2006. PMID: 16999588.
- [55] X. Gu and M. Krenn, “Generation and human-expert evaluation of interesting research ideas using knowledge graphs and large language models,” 2024.
- [56] A. Ghafarollahi and M. J. Buehler, “Atomagents: Alloy design and discovery through physics-aware multi-modal multi-agent artificial intelligence,” 2024.
- [57] J. M. Laurent, J. D. Janizek, M. Ruzo, M. M. Hinks, M. J. Hammerling, S. Narayanan, M. Ponnampati, A. D. White, and S. G. Rodrigues, “Lab-bench: Measuring capabilities of language models for biology research,” 2024.
- [58] H. Cai, X. Cai, J. Chang, S. Li, L. Yao, C. Wang, Z. Gao, H. Wang, Y. Li, M. Lin, S. Yang, J. Wang, M. Xu, J. Huang, F. Xi, J. Zhuang, Y. Yin, Y. Li, C. Chen, Z. Cheng, Z. Zhao, L. Zhang, and G. Ke, “Sciassess: Benchmarking llm proficiency in scientific literature analysis,” 2024.
- [59] P. Clark, I. Cowhey, O. Etzioni, T. Khot, A. Sabharwal, C. Schoenick, and O. Tafjord, “Think you have solved question answering? try arc, the AI2 reasoning challenge,” *CoRR*, vol. abs/1803.05457, 2018.
- [60] S. Lin, J. Hilton, and O. Evans, “Truthfulqa: Measuring how models mimic human falsehoods,” *CoRR*, vol. abs/2109.07958, 2021.

## A Violative & Disallowed Content - Full Evaluations

We used TTS to convert existing text safety evals to audio. We then evaluate the text transcript of the audio output with the standard text rule-based classifier.

Our two main metrics for this eval are:

- **not\_unsafe**: does the model produce audio output that is unsafe?
- **not\_overrefuse**: does the model refuse to comply with a benign request?

We also note sub-metrics for higher severity categories, specifically:

- sexual/minors
- sexual/illegal
- extremist/propaganda
- illicit/violent
- illicit/non-violent
- self-harm/instructions

Below we display the results of these evaluations with the audio and text mode of the GPT-4o Voice Mode model, as well as the text performance of the current GPT-4o model in production.

	<b>Current GPT-4o Text</b>	<b>New GPT-4o – Text</b>	<b>New GPT-4o – Audio</b>
not_unsafe	0.99	0.99	1.0
not_overrefuse	0.91	0.89	0.91
sexual_minors_not_unsafe	0.95	0.98	0.98
sexual_illegal_not_unsafe	0.97	0.98	0.99
extremism_propaganda_not_unsafe	1.0	1.0	1.0
illicit_violent_not_unsafe	1.0	1.0	1.0
illicit_non_violent_not_unsafe	0.99	0.97	1.0
self_harm_not_unsafe	1.0	1.0	1.0

Table 11: Comparison of Current and New GPT-4o Text and Audio Safety Metrics

## B Sample tasks from METR Evaluations






Figure 3: Sample tasks from METR Evaluations