

GPT-4V(ision) technical work and authors

This document acknowledges the contributors and technical work done as part of the GPT-4V project. GPT-4V refers to the technology that enables the integration of multimodal vision capabilities with GPT-4. Our current body of work consists of multiple resources:

- The “GPT-4 Technical Report” covers the GPT-4 system generally as well as quantitative evaluations of GPT-4V in academic evals and exams.
- The “GPT-4V System Card” covers the safety considerations involved in deploying our work.
- The blog post “ChatGPT Can Now See, Hear and Speak” demonstrates the user interface of the realized GPT-4V system as deployed in ChatGPT.
- “The Dawn of LMMs: Preliminary Explorations with GPT-4V(ision)” work from our colleagues at Microsoft covers a plethora of practical observations and strategies for using GPT-4V.

This collection of works and the following credits reflect the multidisciplinary expertise involved in creating, building, and safely deploying multimodal AI while empowering users and educating the public.

Authorship, credit attribution, and acknowledgments

When citing GPT-4V please cite this work as “OpenAI (2023)”. Contributions are sorted alphabetically and assembled by **Raul Puri**.

Research contributions

Jamie Kiros: *Deployment research & evals lead*

Daniel Levy: *Optimization lead*

Hyeonwoo Noh: *Pretraining research lead*

Long Ouyang: *Alignment data lead*

Raul Puri: *Research engineering lead*

Architecture research

Mark Chen, Casey Chu, Jamie Kiros, Christine McLeavey, Hyeonwoo Noh, Raul Puri, Alec Radford, Aditya Ramesh

Distributed training infrastructure

Trevor Cai, Yunxing Dai, Chris Hesse, Brandon Houghton, Yongjik Kim, Lukasz Kondraciuk, Hyeonwoo Noh, Mikhail Pavlov, Raul Puri, Nikolas Tezak, Amin Tootoonchian, Tianhao Zheng

Data

Alex Karpenko, Jong Wook Kim, David M  ly, Reiichiro Nakano, Hyeonwoo Noh, Long

Ouyang, Raul Puri, Alec Radford, Pranav Shyam, Tao Xu

Evaluation data

Sandhini Agarwal, Madeline Boyd, Shengli Hu, Andrew Kondrich, Todor Markov, David Méléy, Hyeonwoo Noh, Reiichiro Nakano, Long Ouyang, Cameron Raymond, Filippo Rasso, Chelsea Voss, Lilian Weng, Chong Zhang, Rowan Zellers, Nicholas Turley

Alignment data

Stephanie Lin, Long Ouyang, Chong Zhang

Deployment, alignment, & post-training research

Ilge Akkaya, Diogo Moitinho de Almeida, Mark Chen, Liam Fedus, Yuchen He, Alex Karpenko, Jamie Kiros, Andrew Kondrich, Rachel Lim, Randall Lin, Stephanie Lin, Ryan Lowe, Luke Metz, Reiichiro Nakano, Long Ouyang, Raul Puri, Jiayi Weng, Barret Zoph

Compute cluster scaling

Andrew Cann, Rory Carmichael, Christian Gibson, Henri Roussez, Akila Weliwinda

Hardware correctness

Oleg Boiko, Trevor Cai, Michael Petrov, Alethea Power

Training run babysitting

Trevor Cai, Kyle Kopic, Daniel Levy, David Méléy, Reiichiro Nakano, Hyeonwoo Noh, Mikhail Pavlov, Raul Puri, Amin Tootoonchian

Safety contributions

Sandhini Agarwal: *Policy research lead*

Lama Ahmad: *Red teaming lead*

Chong Zhang: *Safety systems research lead*

Red teaming leaders

Lama Ahmad, Rosie Campbell, Ashyana-Jasmine Kachra

Safety systems research

Florencia Leoni Aleman, Madelaine Boyd, Yuchen He, Andrew Kondrich, Todor Markov, Raul Puri, Cameron Raymond, Andrea Vallone, CJ Weinmann, Lilian Weng, Mehmet Yatbaz, Chong Zhang

Policy research

Sandhini Agarwal, Lama Ahmad, Miles Brundage, Rosie Campbell, Michael Kolhede, Michael Lampe

Deployment contributions

Madeline Boyd: *Trust & safety engineering lead*

Raul Puri: *Inference infrastructure lead*

Jordan Sitkin: *Deployment platform lead*

Isaac Wolkerstorfer: *ChatGPT engineering lead*

Benjamin Zweig: *Design lead*

Deployment engineering

Valerie Balcom, Jason Chen, Dave Cummings, Bogo Giertler, Joshua Gross, Eric Horacek, Mark Hudnall, Tomer Kaftan, Rachel Lim, Lien Mamitsuka, Rajeev Nayak, Henrique Ponde de Oliveira Pinto, Adam Perelman, Raul Puri, David Schnurr, Eric Sigler, Jordan Sitkin, Javier Soto, Heather Schmidt, Felipe Such, Anton Tananaev, Sherwin Wu, Isaac Wolkerstorfer

ChatGPT client engineering

Valerie Balcom, Bogo Giertler, Eric Horacek, Lien Mamitsuka, Rajeev Nayak, Raul Puri, David Schnurr, Javier Soto, Anton Tananaev

ChatGPT backend engineering

Jason Chen, Joshua Gross, Mark Hudnall, Alex Karpenko, Raul Puri, Eric Sigler, Jordan Sitkin, Isaac Wolkerstorfer, Chong Zhang, Dave Cummings

Deployment platform

Madeleine Boyd, Olivier Godement, Mark Hudnall, Rachel Lim, Raul Puri, Jordan Sitkin, Isaac Wolkerstorfer, Sherwin Wu

Inference infrastructure

Greg Brockman, Tomer Kaftan, Rachel Lim, Raul Puri, Heather Schmidt, Jordan Sitkin, Felipe Such

Trust & safety engineering

Madeleine Boyd

Design

Maddie Simens, Benjamin Zweig

Launch partners, product, and deployment management

Olivier Godement, Joanne Jang, Angela Jiang, Raul Puri, Jessica Shieh, Natalie Staudacher, Nicholas Turley

Additional contributions

Greg Brockman, Peter Deng, Jason Kwon, Bob McGrew, Mira Murati, Srinivas Narayanan, Peter Welinder, Hannah Wong

Communications

Eric Antonow, Ryan Bidy, Ruby Chen, Thomas Degry, Niko Felix, Elie Georges, Kendra Rimbach, Natalie Summers, Justin Jay Wang

Deployment security

Tiffany Citra, Jake McNeil, Karthik Rangarajan

User support

Jeremiah Currier

Legal

Ashley Pantuliano, Filippo Raso, Thomas Stasi

Acknowledgments

We are grateful to our expert adversarial 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: Sally Applin, Gerardo Adesso, Rubaid Ashfaq, Max Bai, Matthew Brammer, Ethan Fecht, Andrew Goodman, Shelby Grossman, Matthew Groh, Hannah Rose Kirk, Seva Gunitsky, Yixing Huang, Lauren Kahn, Sangeet Kumar, Dani Madrid-Morales, Fabio Motoki, Aviv Ovadya, Uwe Peters, Maureen Robinson, Paul Röttger, Herman Wasserman, Alexa Wehsener, Leah Walker, Bertram Vidgen, Jianlong Zhu.

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 and safety research. We also thank the Microsoft Research team for their exploratory work cataloguing use of GPT-4V: Zhengyuan Yang, Linjie Li, Kevin Lin, Jianfeng Wang, Chung-Ching Lin, Zicheng Liu, Lijuan Wang.

Lastly, we thank our deployment partners Be My Eyes for their support and feedback in deploying this technology to the blind and low-vision community.