

Open Assistant

Conversational AI for everyone.



Interview with Andreas Köpf

Would you tell us shortly what the Open Assistant project is all about? What will be the key benefits and features of the model? What is it's history and maybe something about the organization and the people behind the project?

The mission of OpenAssistant is to develop an open-source AI chat-assistant that can follow natural language instructions, fulfill tasks, retrieve information and interact with third-party APIs.

At the heart of the project is the AI alignment of a large language model (LLM). This process aims to steer a language model to communicate in a positive, polite, honest, helpful and harmless manner with the user.

A crucial ingredient for training such a model is a dataset which captures the distribution of desired responses that the model after training will mimic. Building such a dataset is extremely labor and time intensive. We decided to start a large crowd-sourcing campaign and created a web site (open-assistant.io) to gather human demonstrations of prompter-assistant dialogues together with user ratings and reply-rankings to capture human preferences.

The dialogue & preference dataset is then used to fine-tune pre-trained LLMs in a process known as reinforcement learning from human feedback (RLHF). One of the main contributions of OA is to publish the dataset and make it available for everyone to train their own models.

The OA project is a joint effort of hundreds of volunteers from all over the world who collaborate via GitHub and Discord. Thousands of users contributed by completing tasks on our crowd-sourcing data collection website to the OA dataset between January and April this 2023.

OpenAssistant was founded mid December 2022 by Christoph Schuhmann, Huu Nguyen, Yannic Kilcher and me, Andreas Köpf. The project is a collaboration of the LAION and Yannic Kilcher discord communities. LAION is a German non-profit organization for open-source AI research projects of which Christoph Schuhmann is a co-founder. Yannic Kilcher is an influential ML/AI youtuber with a large follower base and an active Discord community. Huu Nguyen is co-founder of ontocord.ai, a former BigLaw partner, a recovering computer scientist and LAION volunteer. I am myself an active member of both the LAION & YK communities. I have led multiple teams before in challenging classic software development, machine learning and robotics projects but helping to organize OpenAssistant was an outstanding experience for me.

Has anything like this been created before? If so, how does this differ? Or is there something especially unique in the project that you would like to share with us?

Conversational AI has a long history, going back to early attempts like Weizenbaum's ELIZA program in 1966. Newer voice-assistants like Siri, Alexa or Google Assistant already successfully applied machine learning for speech-recognition and generation but remained largely rule based when it came to actual tasks. With the release of chatGPT and GPT-4 we have entered a new era of conversational assistant systems which use an aligned large language model at their core. The main difference of OpenAssistant to commercial for-profit products is the open data collection and development process which makes all results available under Apache 2.0 license.

How far along is the project? When can the model be tested?

We released the OpenAssistant Conversations Dataset (OASST1) with more than 160.000 messages in 35 different languages on April 15. Also on that day we published a first version of

our assistant model which can be accessed online at open-assistant.io/chat . Currently in testing are models tuned with reinforcement-learning and the first version of a plug-in system which will allow the assistant to interact with 3rd party web services.

Will Open Assistant be free?

Source code, dataset and models will always be free. Sustainable operation of a LLM online is definitely a major challenge on which we are working. For our launch we have extremely generous sponsors who donate parts of their compute clusters to us.

For the mid or long-term future we will either have to build a network of partners (like universities, businesses and private individuals) who run parts of our distributed inference system or we have to work together with commercial entities to build a refinancing system on top of OA. Another alternative might be a blockchain/token based solution. Of course any support from organizations with compute clusters or people who want to help us organize a network would be welcome.

What kind of contributions are you especially looking for at the moment? How can one contribute to the Open Assistant project?

We always look for front- and backend developers & ML engineers but also technical writers and prompt-engineers or plugin developers. People interested can visit our github repository or join the OA discord server. Besides that the human feedback data collection is still ongoing and we invite everyone to contribute in their language.

What are your goals in the future regarding the Open Assistant project?

Beside internet search we will likely integrate document retrieval systems like vector search databases to allow the assistant to work on larger amounts of data (e.g. read a whole manual). This is also related to extending the context size of our model from currently 2k tokens to 32k and beyond (you can think of a token as a word or fragment of a word).

New larger open-source foundation models will likely be released in the near future. Our current model size is not sufficient to replicate all the capabilities of the commercial-leaders in the field. We will try to work together with the whole AI/ML open-source and research community to improve capabilities and provide free access to this transformative technology to everyone.

Is there anything more that you would like to add regarding the project?

LLM conversational AI agents will have a profound impact on the way we work and what we can effectively automate. During the last months people showed me incredible use-cases. We must ensure that this technology is made accessible to everyone and not monopolized by a few tech giants. Regulators and decision makers face the challenging task to ensure that this technology is used for the benefit of mankind as a whole and does not lead to enormous concentration of power in the hands of a few.

More details about the Open Assistant project:

Data collection & chat: open-assistant.io/

Github: github.com/LAION-AI/Open-Assistant

OA discord: ykilcher.com/open-assistant-discord

Models & datasets: huggingface.co/OpenAssistant