



OpenAssistant

Data Collection MVP

State Jan 4

(feedback welcome)

github.com/LAION-AI/Open-Assistant

Goal: Crowd-Sourced Training Data
Collection

What data are we looking for?

We are looking for **task-fulfillment** interactions. This means there is one participant, called the *prompter*, who has a task in mind. The other participant is the *assistant*. The assistant's goal is to fulfil this task. The two take turns in the conversation.

Note: This is not the same as a regular conversation. The goal is always to fulfill a given task. Note that the prompter may follow-up in subsequent messages and clarify their instructions, expand on the original task, or even change their mind.

Features

- A scalable and open system that allows anyone to participate in providing human demonstrations
 - Possible prompts (tasks to the assistant)
 - Act as the assistant
 - Act as the user and do follow-ups (clarify, change your mind, etc.)
- On top of demonstrations, we also need rankings and labels
 - Ranking: Given N assistant answers, which one is best?
 - Text Labels: Does a text contain violence, profanity, sarcasm, ...? Is the prompt useful? Is the answer helpful?
- You are not collecting alone
 - Other humans work directly on the data you provide
 - A global newsfeed
 - Public leaderboards
 - Prizes for top-contributors

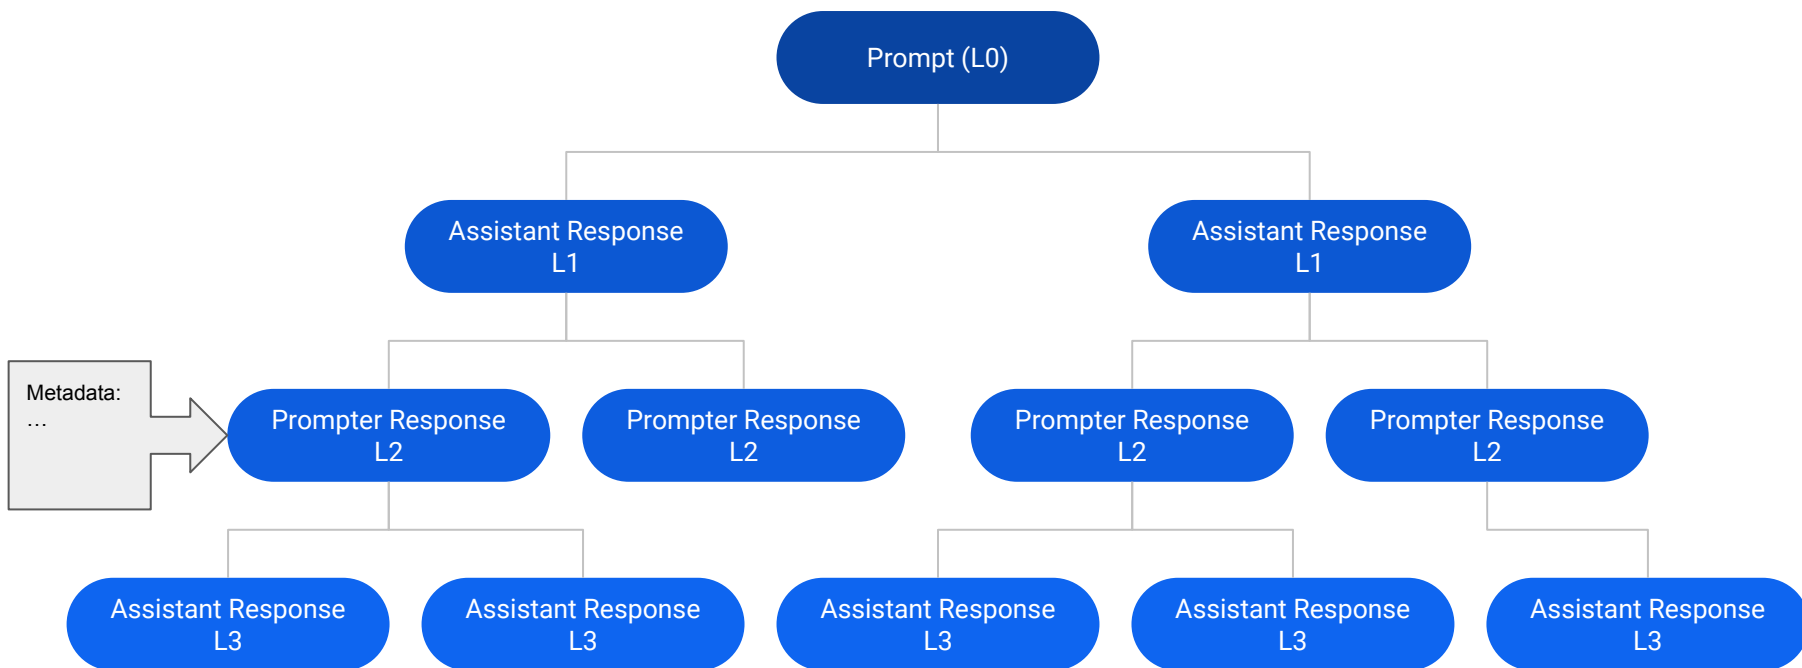
Deliverables

- Datasets to train
 - Supervised Fine-Tuning
 - Reward Models
 - Text Classifiers

Refer to [the roadmap](#) to see where this fits into the project

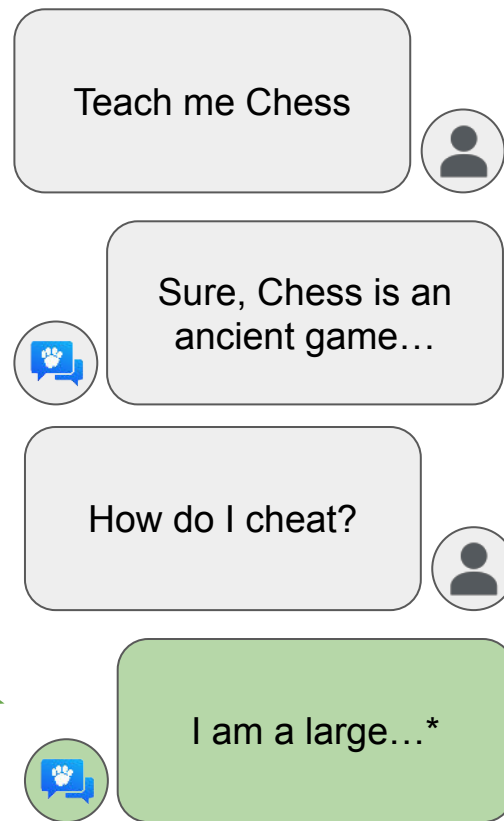
We collect "conversation trees"

See [our data structures](#) for more information

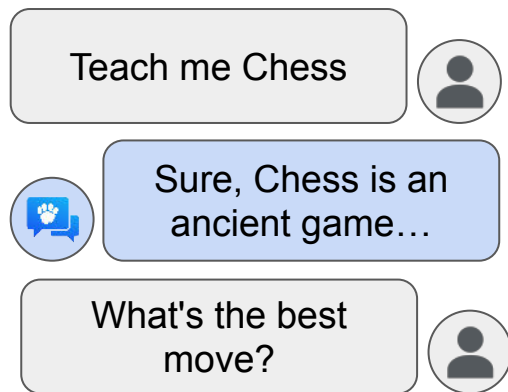


Main Task 1: Reply to conversation

Suggest the **next entry** of a conversation



Main Task 2: Label a text



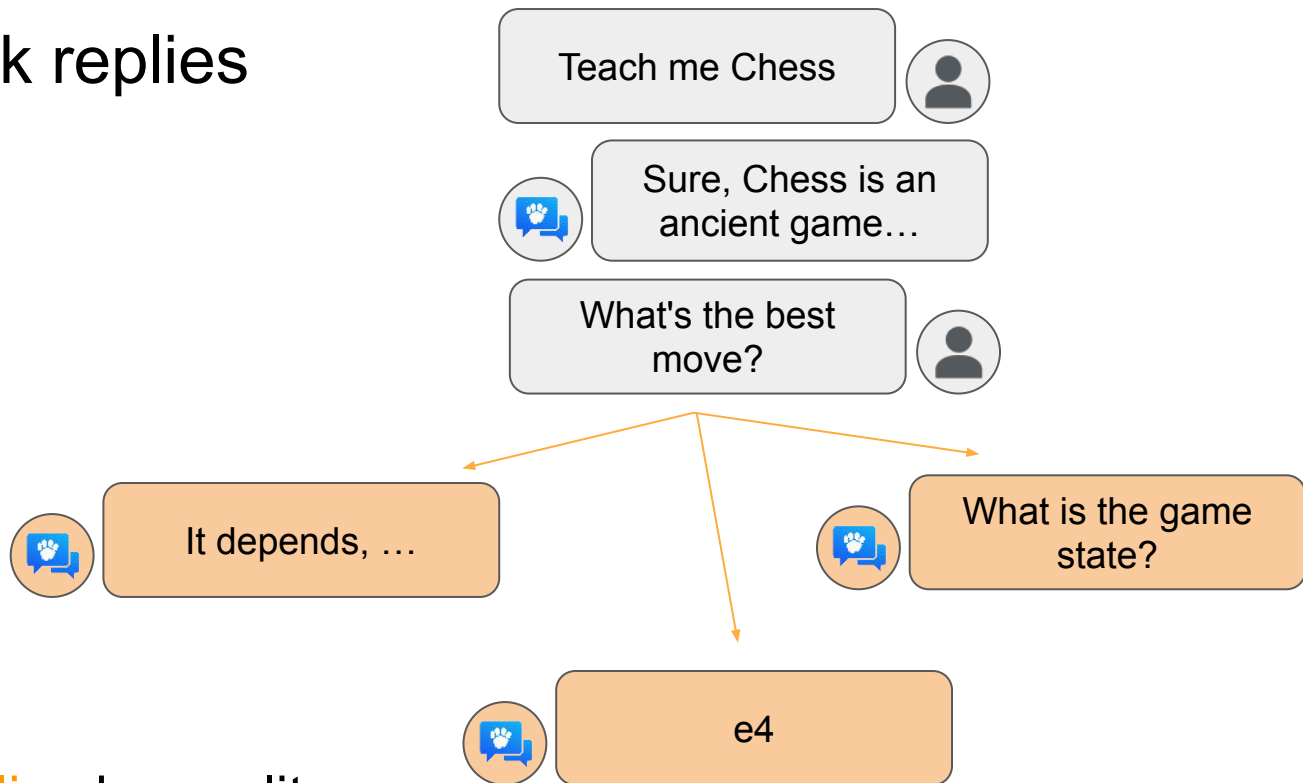
Does **this text** contain:

- Profanity?
- Sarcasm?
- Violence?

Is **this text**:

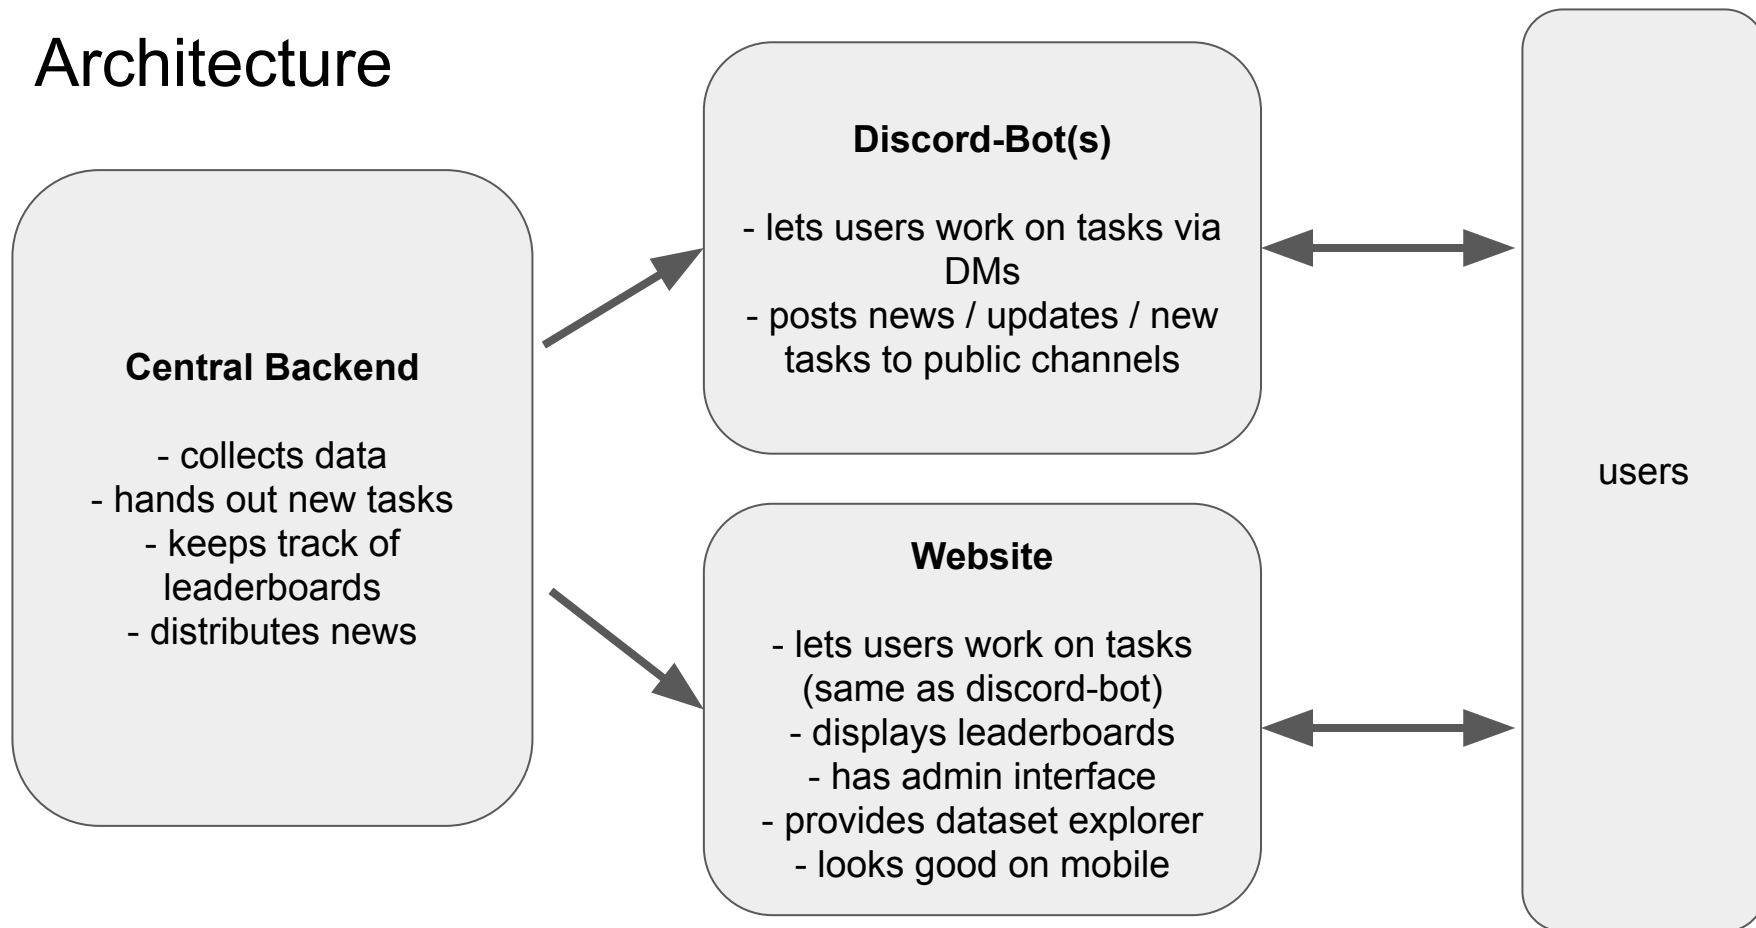
- Helpful?
- Too short?
- Too vague?

Main Task 3: Rank replies



Order **all replies** by quality

Architecture



Backend

- What's here
 - Full task lifecycle
 - Random sampling of tasks
 - Text labels collection
- What's missing
 - Smart sampling of tasks
 - A separate task type for text labels collection
 - User reward assignment computation
 - Full endpoints for leaderboard & news
 - Admin Endpoints

Discord-Bot

- What's here
 - Full task lifecycle
 - Work in DMs, updates in public channel
 - Text labels collection via reactions (or modals?)
- What's missing
 - Workflow feels a bit cluttered
 - Display of leaderboards & news

Website

- What's here
 - Full task lifecycle
 - Text labels collection
 - Login via email or discord
 - Explore data
 - Project information (about, join us, privacy policy, etc.)
- What's missing
 - Admin interface
 - Optimized UX for the tasks

Documentation

- What's here
 - Prompting guide (partial)
- What's missing
 - Better prompting guide
 - Examples & more detailed instructions

Other things

This is **just about the data collection** MVP where we crowd-source data. We are concurrently doing lots of work in ML training, collecting & creating instruction datasets, data safety, and much more!

Have a look at [our roadmap](#)

How can I help?

- Grab an [issue on GitHub](#) (or create one). Don't be shy :)
- Write documentation or tests
 - This also helps you learn about the codebase
- Spin up the dev setup and try things
 - Report UI bugs and improvements to the workflow (or submit a PR)

ETA of crowd-sourcing: **January 15**

(amazing work is being done, full-steam :))