

BOTEVAL: Facilitating Interactive Human Evaluation

Hyundong Cho¹ Thamme Gowda² Yuyang Huang¹

Zixun Lu¹ Tianli Tong¹ Jonathan May¹

¹University of Southern California, Information Sciences Institute

²Microsoft Translator

hd.justincho@gmail.com

Abstract

Following the rapid progress in natural language processing (NLP) models, language models are applied to increasingly more complex interactive tasks such as negotiations and conversation moderations. Having human evaluators directly interact with these NLP models is essential for adequately evaluating the performance on such interactive tasks. We develop BOTEVAL, an easily customizable, open-source, evaluation toolkit that focuses on enabling human-bot interactions as part of the evaluation process, as opposed to human evaluators making judgements for a static input. BOTEVAL balances flexibility for customization and user-friendliness by providing templates for common use cases that span various degrees of complexity and built-in compatibility with popular crowdsourcing platforms. We showcase the numerous useful features of BOTEVAL through a study that evaluates the performance of various chatbots on their effectiveness for conversational moderation and discuss how BOTEVAL differs from other annotation tools.

1 Introduction

As natural language processing (NLP) models become more versatile with the recent advances of language models and their instruction-tuned counterparts (Ouyang et al., 2022), it is becoming more common to create language agents (Sumers et al., 2023) and apply them to complex interactive tasks, such as negotiations (Chawla et al., 2021a), conversational moderation (Cho et al., 2023), reasoning-guided response generation (Zhou et al., 2022), and personalized response generation (Liu et al., 2023).

As noted by Smith et al. (2022), the evaluation methodology plays a critical role in accurately comparing models. For example, rankings between dialogue models can change depending on whether they are evaluated based on single-turn responses or full conversations. In addition, Cho et al. (2023)

found the evaluators point of view when evaluating a model is also an important factor. They showed that human evaluators perceived conversational moderators as more effective in making the evaluators become more cooperative and respectful when the evaluators directly interacted with the moderators while acting as the moderated user (first person point of view) compared to when they evaluated a completed interaction between a moderator and a moderated user as a bystander (third person point of view). However, these factors are overlooked in previous approaches that have focused on a simplified evaluation, such as comparing two complete conversations or individual responses (Smith et al., 2022; Li et al., 2019), or specific dialogue applications such as task-oriented dialogue (Cucurnia et al., 2021; Collins et al., 2019). Therefore, it is important to develop evaluation tools that enable an environment that evaluates models in a setting that best encapsulates how humans actually interact with models.

To facilitate accurate human evaluations of complex interactive tasks, we developed BOTEVAL,¹ a comprehensive evaluation toolkit that focuses on enabling human - bot² interactions as part of the human evaluation process. For flexibility, it is dynamically configurable to accommodate as many human agents and model agents to interact with each other simultaneously with a custom dialogue manager. It is also designed with modular components, such as the interaction interface, instructions, and survey, so that they can be individually adapted to accommodate various use cases. While main-

¹Source code and documentation for BOTEVAL can be found at <https://github.com/isi-nlp/boteval>. We make the demo video of BOTEVAL available at <https://justin-cho.com/boteval>. In addition, a live demo of BOTEVAL is also available at <https://spolin.isi.edu/boteval-dev1> where reviewers can complete a sample human evaluation task.

²We use *bot* loosely to describe any AI system that a human being interacts with.

taining generalizability, BOTEVAL strives to maximize user-friendliness by providing templates for frequent use cases that involve human evaluation where a human evaluator must interact with a NLP model, multiple models, or another human being to measure human performance. In addition, it is integrated with Amazon Mechanical Turk (AMT)³ for crowdsourcing. It can also be deployed independently of these platforms so that it can be used for internal annotations or used with survey tools that allow for external links, such as Qualtrics⁴ and Prolific.⁵

To showcase the usefulness of BOTEVAL and demonstrate its key features, we share a case study that uses BOTEVAL for evaluating models on their performance on conversational moderation (Cho et al., 2023). In this study, BOTEVAL is used to conduct various evaluations: (i) human-bot interactions to compare models; (ii) human-human interactions to measure a human performance threshold; and (iii) completed human-bot interactions by another evaluator to measure evaluator consistency and third-person point of view (POV) results.

In summary, BOTEVAL’s main contributions are:

- An open-source and customizable evaluation tool for interactive NLP tasks that incorporates human-bot and human-human interactions into the evaluation process.
- Detailed documentation and templates for various use cases to make modifications easy.
- Flexible deployment options with built-in integration with popular crowdsourcing platforms such as AMT and Prolific.
- Evaluation task management features that facilitate task monitoring and managing crowdsource workers.
- Dynamically configurable interaction logic with custom dialogue manager and multi-human and multi-bot evaluation settings.

2 BOTEVAL System Overview

BOTEVAL is a web application that provides an evaluation interface, what the human evaluators (i.e., crowdsource workers) see (Section 2.1), and an administrator dashboard, what the administrator uses to manage the evaluation task and evaluators

(Section 2.2). We recommend that the bots that evaluators interact with are provided as separate APIs that BOTEVAL can make queries to, as this isolates the management of the bot deployment and BOTEVAL (Section 2.3). Human evaluators can be flexibly set to crowdsource workers from AMT or Prolific or any other evaluators with internet access by having them create an account directly for a deployment of BOTEVAL using a public external link. An evaluation task is configured with a central YAML config file that identifies the frontend components to use, the deployment environment, and the crowdsourcing platform to use.

2.1 Evaluation interface

A sample evaluation interface for the case study later described in Section 4.1 is shown in Figure 1.

The evaluation interface consists of three main components: ① *Conversation pane*: a section where the interaction between the human and the bot takes place. This pane can be easily customized to contain seed conversations to serve as initial starting points for interactions to continue off of or it can instead contain any piece of text or completed conversation without requiring any interactions from the evaluators, making BOTEVAL also suitable for simpler annotation tasks. ② *Instruction pane*: this is an optional section that shows the main directions. Evaluators can see detailed instructions by clicking on the detailed instructions button. Administrators can choose to show detailed instructions as part of the consent form if one is needed to make sure that evaluators have read them. ③ *Survey pane*: this is where the human evaluators provide their evaluations. In the given example, it is configured to only be shown after the human evaluators have interacted with the bot for a set number of turns.

The conversation pane and instruction pane is configurable by providing custom HTML scripts, while the survey pane is even more easily customizable by configuring a YAML config file. An example of the YAML config file is shown in Appendix A.1. An optional consent form can be shown to evaluators as well, which is also managed with a separate HTML file. Further detail on how the consent form can be configured is in Appendix A.2.

2.2 Administrator dashboard

BOTEVAL’s administrator dashboard provides numerous features for managing evaluation tasks and evaluators. Its main benefit is a GUI that enables

³<https://www.mturk.com>

⁴<https://www.qualtrics.com>

⁵<https://www.prolific.com>




Figure 1: A snapshot of the admin point of view of an evaluation interface with a completed evaluation example. The interface is identical for the evaluator except for the text that shows the evaluator’s worker ID (hidden with asterisks in the figure for privacy). The three main components of the user-facing interface are the ① conversation pane, ② simple instruction pane, and the ③ survey pane.

a non-technical user to easily become an administrator for human evaluation tasks. The topics page, shown in Figure 2, is one view that allows the management of launching and deleting tasks. A topic refers to any predetermined context, such as seed conversation or external information relevant for the evaluation task. These topics are provided to BOTEVAL as a JSON file. If a user is interested in general open-domain dialogue evaluation, measuring a model’s general conversational capabilities, they can use a dummy topic file that contains an empty dictionary. This will launch an evaluation task that starts a conversation from scratch, with the human evaluator initiating the first turn.

After launching tasks, users can use the administrator dashboard to conveniently examine tasks that are completed or in progress with the same interface that the evaluators used to complete the task to easily visualize their work rather than examining a database or JSON file, as shown in Figure 1. The user can also directly export individual JSON files of the collected data if needed. Also, tasks can be deleted in batches using the parallel management tool shown in ① of Figure 2.

In addition to these features, we provide convenient AMT-specific features for managing workers and tasks known as human intelligence tasks (HITs). One of the most convenient features is being able to directly assign and remove qualifica-

tions for workers after examining their work without having to leave the administrator dashboard. This is an important convenience feature for ensuring the quality of work for human evaluations are kept to the desired standard by blocking unreliable workers. Another is being able to make bonus payments directly after examining the completed task, which is useful when each task is expected to involve variable rewards, such as to account for each HIT taking a different amount of time to complete.

2.3 Bot customization

Users are given multiple options to choose how they will service the bot that they want to evaluate, but the recommended setup is to set up a separate RESTful API and defining a logic within BOTEVAL to interface with this API. As shown in ③ in Figure 2, users can define task-specific parameters for bots that get passed on to the API if the API allows for it. This is useful if you are using the same model but adjusting the instruction prompt (e.g., using OpenAI endpoints). While BOTEVAL users have the option to launch bots simultaneously on the same server with BOTEVAL’s process, it is more efficient to separately manage human evaluation tasks and the processes that load and query NLP models because most NLP models are better served with GPUs for reducing latency.

Topics

External URL: <https://spolin.isi.edu/boteval-dev1/> Okay? Yes

If you are planning to crowdsource tasks on platforms like MTurk, make sure to have the above URL is accessible via the (public) Internet.
► More Info:

1 Set maximum thread per user:
Max_threads_per_user:

Multi-task Launch:
Unlaunched Tasks:

2 Multi-task Creation:
Topics:
ca3849b2 - science - The oldest and thickest Arctic
e9688ebf - worldnews - Worlds 26 richest people own
4190e3a2 - technology - A defiant 8chan vowed to
2130c5a8 - worldnews - Syria Says US Should Give
bbce6425 - technology - Social media is a threat
2e3bcd21 - news - Man charged after screaming racial
76680380 - worldnews - Christchurch mosque shooting victims receive
f94f9a5b - science - Lynches were in part a
Endpoint:
Persona_id:
Max_threads_per_topic:
Max_turns_per_thread:
max_human_users_per_thread:
human_moderator:
Reward:

Found 58 topics

#	Topic	Num Threads	Time	Actions
1	ca3849b2 - science - The oldest and thickest Arctic	0	Created 2024-03-19 07:10:05;	Endpoint: <input type="text" value="gpt4"/> Persona_id: <input type="text" value="moderator"/> Max_threads_per_topic: <input type="text" value="3"/> Max_turns_per_thread: <input type="text" value="3"/> max_human_users_per_thread: <input type="text" value="1"/> human_moderator: <input type="text" value="no"/> Reward: <input type="text" value="1.5"/> <input type="button" value="Create New Task"/>

Figure 2: A snapshot of the topics page of the admin dashboard. **1** is a parallel management tool that enables setting global configurations such as how many tasks each evaluator is allowed to complete and launching or deleting multiple tasks at once. **2** is a topics table that shares more information about each topic, such as its name, how many tasks have been created, and when they were created. **3** is a list of parameters that can be chosen for launching a task, which includes parameters that can be passed on to API queries for the bots.

2.4 Sourcing human evaluators

BOTEVAL can be customized to use with any crowdsourcing platform, and it is designed to be directly used with many popular ones such as AMT, Prolific, and Qualtrics. If the goal is to do internal annotations, the setup is even simpler as the user only has to configure BOTEVAL to not use any. Then the user can share their custom URL with the evaluators, where they can sign up and directly work on tasks that are made available to them without going through any other platform.

3 System Architecture

An overview of BOTEVAL’s system architecture is shown in Figure 3. BOTEVAL is a web application (i.e., a client-server model). We describe the front- and back-end technology stacks in the following sections.

3.1 Frontend

The frontend is a simple web interface (i.e., HTML) created with Bootstrap stylesheet. While the majority of the HTML structure is constructed on the server side using Jinja2, some dynamic updates such as responses coming from bots or other participants in the interaction are achieved using AJAX and RESTful APIs.

3.2 Backend

The backend is implemented in Python language using Flask framework, following a model-view-controller architecture pattern. *Models* are implemented using Python classes and stored in a relational database, specifically SQLite. In addition, we use SQLAlchemy, an object-relational mapper, to abstract the mapping between Python classes and database tables. For *views*, Flask uses Jinja2 for server side templating of HTML pages. *Controllers* are based on Flask’s builtin URL routers and RESTful API constructs.

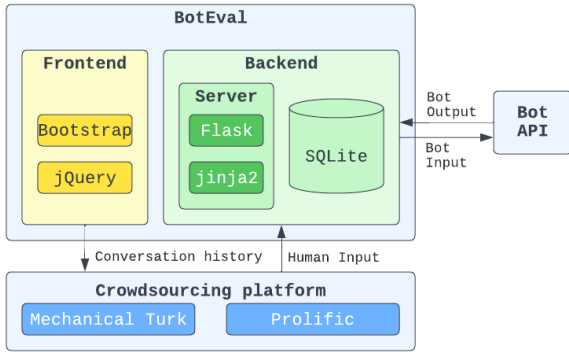


Figure 3: BOTEVAL system architecture. We use popular frameworks that are well documented and easy to use.

While internally our server is an HTTP server, crowdsourcing platforms such as AMT require annotation interface be served via secure connections (HTTPS). HTTPS can be enabled by obtaining and installing an SSL/TLS certificate. We use free certificates from Certbot,⁶ and configure Nginx⁷ as a reverse proxy server for HTTPS requests.

Some scenarios may require several simultaneous instances of BOTEVAL to facilitate multiple annotation tasks, and obtaining SSL certificate for each instance maybe cumbersome. We address this problem by using a different TCP port for each instance, and configuring a single Nginx (with SSL certificate) route requests for all instances.

4 Case Study and Use Cases

4.1 Case study: conversational moderation evaluation

To showcase the usefulness of BOTEVAL, we share a case study that uses BOTEVAL to conduct a study on how effective various zero-shot instruction-tuned language models (ITLM) and dialogue model are in performing *conversational moderation* (CM) (Cho et al., 2023).⁸ Instead of iron-fisted approaches to moderation such as deleting comments or banning users, which may exacerbate societal polarization as these users find refuge in echo chambers, CM seeks to have moderators interact with users exhibiting problematic behavior to guide them back to more constructive and respectful conversations.

⁶<https://certbot.eff.org/>

⁷<https://nginx.org/en/>

⁸The BOTEVAL template for this work is available at https://github.com/isi-nlp/isi_darma/tree/main/boteval-darma-task.

This study makes full use of BOTEVAL as it requires evaluating multiple bots by interacting with them for a preset number of turns (in this case 3), starting with a variety of conversation stubs. The evaluations were conducted with all desired configurations simultaneously to get the most representative and fair results that is not affected by any confounding factors such as recency bias. The evaluation was conducted with AMT, and being able to easily monitor evaluations enabled rapid iterations of updating the instructions and giving feedback to the evaluators.

Therefore, BOTEVAL was integral in being able to refine the evaluation study efficiently and ultimately collect statistically meaningful results for an interactive evaluation setup. The study showed that prompt-engineered ITLMs outperformed prosocial dialogue models and that a conflict resolution prompt based on the Socratic method was the best performing prompt. In addition, one of this work’s central findings was discovering that there are differences between evaluation results when the models were evaluated from a first person point of view (POV) and a third person POV. With BOTEVAL, collecting human evaluations in these two different settings was a simple change of updating the topics file such that the conversation stubs were the completed conversations, rewording the questions such that it is in third person POV, and setting the number of turns required for human evaluators to interact with the bots to zero.

4.2 Main use cases

BOTEVAL’s main differentiation with previous annotation tools and frameworks is that it is focused on, but not limited to, interactive use cases. In other words, it is useful when the annotated data is not static, e.g., bot responses over multiple turns or other dynamic outputs that can change based on user interaction. Therefore, BOTEVAL is appealing for evaluating or collecting data for conversational tasks that usually require multi-turn interactions for fulfilling the goal, rather than a single generated output. Many real-life tasks go through multi-turn interactions, such as negotiations (Chawla et al., 2021b), counseling (Mehta et al., 2022), and improvisational theater (Cho and May, 2020). As artificial systems become more capable, more will be applied to completing these complex multi-turn tasks, and BOTEVAL will serve as a handy starting point for facilitating their evaluation.

Although BOTEVAL was designed for interac-

Name	Human-bot interaction-focused	Crowdsource integration	Multi-human & bot support	Language
BOTEVAL (Ours)	✓	✓	✓	Python
Mephisto (Urbanek and Ringshia, 2023)	✗	✓	✗	Python
Pigeon ¹¹	✗	✗	✗	Python
MATILDA (Cucurnia et al., 2021)	✗	✓	✗	Python
LIDA (Collins et al., 2019)	✗	✓	✗	Python
INCEpTION (Klie et al., 2018)	✗	✗	✗	Java
GATE (Cunningham, 2002)	✗	✗	✗	Java
BRAT (Stenetorp et al., 2012)	✗	✗	✗	Python
doccano (Nakayama et al., 2018)	✗	✗	✗	Python
Potato (Pei et al., 2022)	✗	✗	✗	Python
Argilla ¹⁰	✗	✗	✗	Python
Prodigy ¹²	✗	✗	✗	Python
DialogueView (Yang and Heeman, 2005)	✗	✗	✗	TcK/TK
DART (Weisser, 2016)	✗	✗	✗	Perl
Anvil (Kipp, 2001)	✗	✗	✗	Java
EZCAT (Guibon et al., 2022)	✗	✗	✗	Javascript

Table 1: Comparison overview with other annotation tools. BOTEVAL innately supports evaluations that require human-bot interactions and allow for multiple human agents or bot agents to be involved in each evaluation sample.

tive tasks, BOTEVAL can also be easily adapt for simple static annotation tasks by simplifying the conversation pane in Figure 1. This pane can contain any other modality such as images, video, and audio, and adjusting the instruction and survey panes accordingly can make BOTEVAL also suitable for text classification or conversation-level or turn-level comparisons, similar to Smith et al. (2022). As BOTEVAL gets actively used for more research studies, we will be able to provide a variety of templates that accommodate a comprehensive set of use cases, further lowering the effort required to conduct effective human evaluation for new studies.

5 Related Work

In Table 1, we compare BOTEVAL with other related annotation tools and discuss differences further here.

5.1 General text annotation tools

A popular general annotation tool is Mephisto (Urbanek and Ringshia, 2023), which started by isolating the crowdsourcing features from ParlAI (Miller et al., 2017). Mephisto provides a general annotation framework that interfaces with Amazon Mechanical Turk and Prolific and includes basic templates for simple annotation tasks. BOTEVAL adapted many of its AMT integration features, but Mephisto is not customized for common interactive data annotation and evaluation use cases, and thus requires nontrivial effort to create a human evaluation environment for interactive NLP tasks where a human evaluator needs to interact with a bot or another human and then evaluate their performance.

ParlAI still provides templates for Mephisto for human-bot interactions⁹, but it is not easy to use with a dialogue model that is not developed with ParlAI. With BOTEVAL, we also provide a GUI administrator dashboard for task and worker management, which is absent in ParlAI and Mephisto.

GATE (Cunningham, 2002) and INCEPTION (Klie et al., 2018) are annotation tools that provide many predefined features, but they are also not designed for interactive human evaluations. Other simpler general text annotation tools that share similar limitations are Doccano (Nakayama et al., 2018), brat (Stenetorp et al., 2012), Argilla¹⁰, Potato (Pei et al., 2022) and Pigeon¹¹, which are web-based annotation tools that enable rapid annotations for text classification and machine translation. Prodigy¹² is a commercial annotation tool for text annotations that provides similar features.

5.2 Dialogue annotation tools and evaluation methodologies

A prominent set of annotation tools specific to dialogue are centered around task-oriented dialogue (Budzianowski et al., 2018). LIDA (Collins et al., 2019) is an annotation tool that provides useful features for efficiently making turn-level annotations, incorporating model-provided label recommendations to speed up annotations, and resolving inter-annotation disagreements. MATILDA (Cucurnia et al., 2021) builds on LIDA for multilin-

⁹https://parl.ai/docs/tutorial_crowdsourcing.html

¹⁰<https://argilla.io>

¹¹<https://github.com/agermanidis/pigeon>

¹²<https://prodi.gy>

gual support and improved management of crowdsourcing tasks among multiple workers. However, they do not have built-in compatibility with popular crowdsourcing platforms and do not support human-bot interactions to take place within the crowdsourcing task. A lightweight option for dialogue annotations is EZCAT (Guibon et al., 2022), which provides a web-based serverless annotation framework that focuses on enhanced accessibility for conversation-level and turn-level annotations.

Other work have created tools for multi-modal annotations or speech-based annotations. Anvil (Kipp, 2001) provides a multi-modal dialogue annotation tool that enables annotation of audiovisual content. DialogueView (Yang and Heeman, 2005) is an annotation tool that is focused on segmenting audio conversations. DART (Weisser, 2016) focuses on enabling efficient annotations of speech acts and linguistic criteria to facilitate corpus-based research into pragmatics. Text is still the primary focus of BOTEVAL and the templates we provide, but BOTEVAL remains general enough to be adapted to such cases as well by modifying the templates we provide.

6 Conclusion

We presented BOTEVAL and its usefulness in collecting human evaluations for interactive tasks that require live human-bot interactions through a case study of evaluating various language models on their ability to conversationally moderate online discussions. BOTEVAL provides a customizable interface that can be adapted for various evaluation and annotation use cases while also providing integration with popular crowdsourcing platforms and task management features. We hope that this work will serve as an important foundation for setting up custom interactive human evaluation tasks that facilitate our understanding of more complex NLP systems as they become increasingly sophisticated and capable.

Limitations

We designed BOTEVAL to be modular such that customizing existing templates and modifying the dialogue manager’s logic is simple, but it is yet not configured so that the task management process, shown in Figure 2 is independent of the process that serves the evaluation, shown in Figure 1. This means that any updates to BOTEVAL that help with task management cannot be applied without restart-

ing evaluation tasks that were launched already, which will interfere with any concurrent tasks that evaluators are working on. While inconvenient, this has not been a major issue as restarting can be done quickly such that it does not interfere the work of many evaluators and this doesn’t mean that existing crowdsourcing tasks, as in the case of AMT, will be deleted and need to be relaunched again.

Another challenge for using BOTEVAL may arise from the difficulty of managing a separate process that serves the bots that the human evaluators will interact with. However, if the BOTEVAL user is able to launch a bot as part of BOTEVAL, refactoring the code for that bot such that its responses are accessed through an API instead is a simple modification with plenty of online tutorials and tools, such as FastAPI.¹³

References

- Paweł Budzianowski, Tsung-Hsien Wen, Bo-Hsiang Tseng, Iñigo Casanueva, Stefan Ultes, Osman Ramadan, and Milica Gašić. 2018. [MultiWOZ - a large-scale multi-domain Wizard-of-Oz dataset for task-oriented dialogue modelling](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 5016–5026, Brussels, Belgium. Association for Computational Linguistics.
- Kushal Chawla, Jaysa Ramirez, Rene Clever, Gale Lucas, Jonathan May, and Jonathan Gratch. 2021a. [Casino: A corpus of campsite negotiation dialogues for automatic negotiation systems](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 3167–3185.
- Kushal Chawla, Jaysa Ramirez, Rene Clever, Gale Lucas, Jonathan May, and Jonathan Gratch. 2021b. [CaSiNo: A corpus of campsite negotiation dialogues for automatic negotiation systems](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 3167–3185, Online. Association for Computational Linguistics.
- Hyundong Cho, Shuai Liu, Taiwei Shi, Darpan Jain, Basem Rizk, Yuyang Huang, Zixun Lu, Nuan Wen, Jonathan Gratch, Emilio Ferrara, and Jonathan May. 2023. [Can language model moderators improve the health of online discourse?](#)
- Hyundong Cho and Jonathan May. 2020. [Grounding conversations with improvised dialogues](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 2398–2413, Online. Association for Computational Linguistics.

¹³<https://github.com/tiangolo/fastapi>

- Edward Collins, Nikolai Rozanov, and Bingbing Zhang. 2019. **LIDA: Lightweight interactive dialogue annotator**. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP): System Demonstrations*, pages 121–126, Hong Kong, China. Association for Computational Linguistics.
- Davide Cucurnia, Nikolai Rozanov, Irene Sucameli, Augusto Ciuffoletti, and Maria Simi. 2021. **MATILDA - multi-AnnoTator multi-language InteractiveLightweight dialogue annotator**. In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: System Demonstrations*, pages 32–39, Online. Association for Computational Linguistics.
- Hamish Cunningham. 2002. Gate, a general architecture for text engineering. *Computers and the Humanities*, 36:223–254.
- Gaël Guibon, Matthieu Labeau, Luce Lefevre, and Chloé Clavel. 2022. **Ezcat: an easy conversation annotation tool**. In *Proceedings of the 13th Language Resources and Evaluation Conference*, Marseille, France. European Language Resources Association.
- Michael Kipp. 2001. Anvil-a generic annotation tool for multimodal dialogue. In *Seventh European conference on speech communication and technology*. Citeseer.
- Jan-Christoph Klie, Michael Bugert, Beto Boullosa, Richard Eckart De Castilho, and Iryna Gurevych. 2018. The inception platform: Machine-assisted and knowledge-oriented interactive annotation. In *Proceedings of the 27th international conference on computational linguistics: system demonstrations*, pages 5–9.
- Margaret Li, Jason Weston, and Stephen Roller. 2019. Acute-eval: Improved dialogue evaluation with optimized questions and multi-turn comparisons. *arXiv preprint arXiv:1909.03087*.
- Shuai Liu, Hyundong Cho, Marjorie Freedman, Xuezhe Ma, and Jonathan May. 2023. **Recap: Retrieval-enhanced context-aware prefix encoder for personalized dialogue response generation**. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 8404–8419.
- Maitrey Mehta, Derek Caperton, Katherine Axford, Lauren Weitzman, David Atkins, Vivek Srikumar, and Zac Imel. 2022. **Psychotherapy is not one thing: Simultaneous modeling of different therapeutic approaches**. In *Proceedings of the Eighth Workshop on Computational Linguistics and Clinical Psychology*, pages 47–58, Seattle, USA. Association for Computational Linguistics.
- A. H. Miller, W. Feng, A. Fisch, J. Lu, D. Batra, A. Bordes, D. Parikh, and J. Weston. 2017. **Parlai: A dialog research software platform**. *arXiv preprint arXiv:1705.06476*.
- Hiroki Nakayama, Takahiro Kubo, Junya Kamura, Yasufumi Taniguchi, and Xu Liang. 2018. **doccano: Text annotation tool for human**. Software available from <https://github.com/doccano/doccano>.
- Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Gray, John Schulman, Jacob Hilton, Fraser Kelton, Luke Miller, Maddie Simens, Amanda Askell, Peter Welinder, Paul Christiano, Jan Leike, and Ryan Lowe. 2022. **Training language models to follow instructions with human feedback**. In *Advances in Neural Information Processing Systems*.
- Jiaxin Pei, Aparna Ananthasubramaniam, Xingyao Wang, Naitian Zhou, Apostolos Dedeloudis, Jackson Sargent, and David Jurgens. 2022. **Potato: The portable text annotation tool**. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*.
- Eric Smith, Orion Hsu, Rebecca Qian, Stephen Roller, Y-Lan Boureau, and Jason Weston. 2022. **Human evaluation of conversations is an open problem: comparing the sensitivity of various methods for evaluating dialogue agents**. In *Proceedings of the 4th Workshop on NLP for Conversational AI*, pages 77–97, Dublin, Ireland. Association for Computational Linguistics.
- Pontus Stenetorp, Sampo Pyysalo, Goran Topić, Tomoko Ohta, Sophia Ananiadou, and Jun'ichi Tsujii. 2012. **Brat: a web-based tool for nlp-assisted text annotation**. In *Proceedings of the Demonstrations at the 13th Conference of the European Chapter of the Association for Computational Linguistics*, pages 102–107.
- Theodore Sumers, Shunyu Yao, Karthik Narasimhan, and Thomas Griffiths. 2023. **Cognitive architectures for language agents**. *Transactions on Machine Learning Research*.
- Jack Urbanek and Pratik Ringshia. 2023. **Mephisto: A framework for portable, reproducible, and iterative crowdsourcing**.
- Martin Weisser. 2016. Dart—the dialogue annotation and research tool. *Corpus Linguistics and Linguistic Theory*, 12(2):355–388.
- Fan Yang and Peter A Heeman. 2005. **Dialogueview: an annotation tool for dialogue**. In *Proceedings of HLT/EMNLP 2005 Interactive Demonstrations*, pages 20–21.
- Pei Zhou, Hyundong Cho, Pegah Jandaghi, Dong-Ho Lee, Bill Yuchen Lin, Jay Pujara, and Xiang Ren. 2022. **Reflect, not reflex: Inference-based common**

ground improves dialogue response quality. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 10450–10468, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.

Appendix

A Evaluation interface configurations

A.1 Sample survey configuration

An example survey configuration is shown in [Figure 4](#). Survey components are also easily dynamically configurable with common input options such as radio buttons, Likert scales, freeform text, etc.

Users can create their own survey with HTML as well, but the customization options we provide through the YAML file covers most conversation-level evaluation use cases.

A.2 Consent form configuration

An example of configuring the consent form is shown in [Figure 5](#). When deployed without a crowdsourcing platform, these appear as checkboxes in the sign up process. Within AMT, we automatically sign up the workers with their worker ID and do not require a password, but they have to check the same checkboxes in order to move on to the task if they are doing the task for the first time.

```

ratings:
- question: Did the moderator make specific suggestions for the given conversation
  to facilitate cooperation?
  choices: &choices
  - Not at all
  - Mostly not
  - So-so
  - Somewhat
  - Very
- question: Was the moderator fair to all users involved in the conversation?
  choices: *choices
- question: Did you (the user) become more engaged and willing to cooperate? (e.g.
  provide more details or ask sincere questions to make the conversation more constructive
  or be more persuasive)
  choices: *choices
- question: Did you (the user) become more respectful and less abusive? (e.g. less
  profanity, unconstructive criticism, or condescending sarcasm)
  choices: *choices
- question: How much did you agree with the arguments/viewpoints of the user that
  you were acting as?
  choices: *choices
- question: How much did you like the person that you were acting as?
  choices: *choices
- question: How could the moderator have been more effective? (e.g. reduce repetition,
  less generic suggestions, more examples, etc.)
  freetext:
    default: none
    required: false
- question: How can we improve the task design or survey questions?
  freetext:
    default: none
    required: false

```

Figure 4: An example of the survey pane configuration that contains a custom Likert scale and freeform text input fields. This configuration corresponds to the survey pane partially shown in [Figure 1](#).

```

onboarding:
  agreement_file: user-agreement.html
  instructions_file: instructions.html
  simple_instructions_file: simple_instructions.html

checkboxes:
  instructions_read: I have read the instructions.
  iam_adult: I am 18 years or older and I understand that I may have to read and
  write using toxic language.

```

Figure 5: An example of configuring the consent form. The `agreement_file` parameter should point to the HTML file that shows the content of the consent form.