

The ShareLM Collection and Plugin: Contributing Human-Model Chats for the Benefit of the Community

Shachar Don-Yehiya¹ Leshem Choshen^{2,3} Omri Abend¹

¹The Hebrew University of Jerusalem, ²MIT, ³MIT-IBM Watson AI Lab
{first.last}@mail.huji.ac.il

Abstract

Human-model conversations provide a window into users' real-world scenarios, behavior, and needs, and thus are a valuable resource for model development and research. While for-profit companies collect user data through the APIs of their models, using it internally to improve their own models, the open source and research community lags behind.

We introduce the ShareLM collection, a unified set of human conversations with large language models, and its accompanying plugin, a Web extension for voluntarily contributing user-model conversations. Where few platforms share their chats, the ShareLM plugin adds this functionality, thus, allowing users to share conversations from most platforms. The plugin allows the user to rate their conversations, both at the conversation and the response levels, and delete conversations they prefer to keep private before they ever leave the user's local storage. We release the plugin conversations as part of the ShareLM collection, and call for more community effort in the field of open human-model data.

The **code**, **plugin**, and **data** are available.¹

1 Introduction

Recently, with the development of more capable models such as GPT4 (OpenAI et al., 2024) and LLAMA (Dubey et al., 2024), interacting with large language models (LLMs) has become common not only among Machine Learning experts, but also the general public. Human users have natural language conversations with the models, and use them for a wide range of use cases (Zhao et al., 2024). In turn, these conversations can be used for

¹Code: <https://github.com/shachardon/share-lm>, Plugin: <https://chromewebstore.google.com/detail/sharelm-share-your-chat-c/nldoebkdaiidhceaphmipeclmlcbljmh>, Data: <https://huggingface.co/datasets/shachardon/ShareLM>

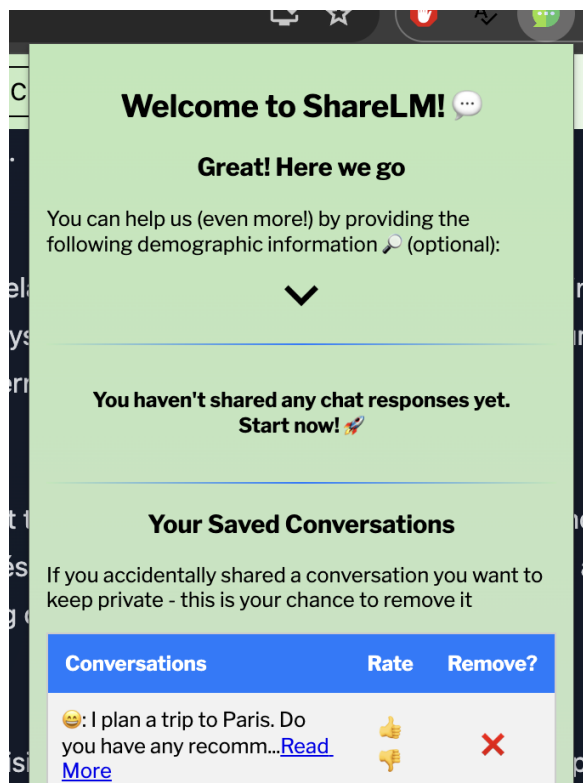


Figure 1: The popup window. The user can go over their previous conversations from the last 24 hours and rate them or alternatively choose to delete them if they prefer to keep them private.

training and better-aligning models to human preferences, as they provide a window into the users' real-world scenarios and needs (Bai et al., 2022). The conversations are also important for other research aspects, such as cognitive and linguistic research revealing the gaps in the mode of interaction between models and humans (Don-Yehiya et al., 2023).

Despite being a cornerstone for LLM development and research, mechanisms for openly collecting and sharing human conversations and feedback are still underdeveloped. In the meantime, models developed by for-profit companies collect

user-model conversations via their APIs to be used to further train their own models (Ouyang et al., 2022), leaving the open-source and research community far behind. The development process of these “closed models” is not always transparent, and so are their data and data collection pipelines. These all make developing platforms and tools for collecting human-model conversations a high priority (Don-Yehiya et al., 2024).

We collected existing human-model conversations datasets, and unified them under format. We call it the **ShareLM collection**. Doing so, we recognized that most of the existing open datasets are treated as static collections rather than a living artifact that can dynamically grow (see §7). Unlike traditional Natural Language Processing datasets (e.g., grammatical error correction), human-model conversations and preferences vary across individuals and time (Pozzobon et al., 2023). Also, these types of data collection efforts are not something that private users can be part of and may lack in diversity (Pavlick et al., 2014).

To overcome this, we introduce the **ShareLM plugin**, a Chrome extension that allows users to easily contribute their conversations with models. The ShareLM plugin collects the user’s conversations with models, supporting multiple platforms and hence not limited to certain models, serving infrastructure or user interface. Among its main features, the plugin supports thumbs up/down rating, and a delayed upload that allows users to go over their conversations from the last 24 hours and remove those that they prefer to keep private before they ever leaved the user’s local storage. The plugin provides the end-point user with ownership of their data, allowing them to keep, delete and retrieve their data and to contribute it for the benefit of the community. The plugin holds the potential to maintain an ever-growing dataset, up-to-date with users’ conversations with the state-of-the-art models of the moment.

We release the conversations that are collected by the plugin as part of the broader ShareLM collection. We hope to see more efforts in the field and contributions to the ShareLM collection, with the aim of sharing open data.

2 The ShareLM Collection

We collected existing human-model conversations datasets that are publicly released. As we focus on human-model conversations and realistic inter-

actions, we exclude other conversation datasets such as human-human (such as in OpenAssistant (Köpf et al., 2024; Zhang et al., 2018)), model-model (Honovich et al., 2023; Wang et al., 2023) or human-model but not conversations (Nakano et al., 2021).

The current list of datasets contains the following; HH-RLHF (Bai et al., 2022) which contains conversations of users with a closed model and their preferences, the dialog task of the bAbi-tasks (Weston et al., 2015), the self-feeding chatbot data (Hancock et al., 2019), the Collective Cognition dataset (see §7), and PRISM (Kirk et al., 2024), containing conversations and preferences of users born in 75 countries, residing in 38 countries with 21 different LLMs both opened and closed. Two more large datasets are WildChat (Zhao et al., 2024), a dataset of over 1M conversations of users with ChatGPT, and the LMSYS-Chat-1M (see §7). The last two are gated datasets², and thus require the user to conform to their terms of use prior to downloading them. We note that all these datasets were not collected by us originally and therefore we assume no responsibility. We ask the users to check each dataset directly for the appropriate citations and licenses. Still, those datasets mainly follow open licenses and we follow their licenses in the unification process.

Together with the conversations that were collected so far by the ShareLM plugin, the ShareLM collection currently contains over 2.3M conversations, from over 40 different models.

The unified format includes the following fields: *conversation_id* to identify each conversation, *conversation* that contains the content of the conversation, *model_name* (if available), *user_id* an anonymized identifier of the user (if available), a *timestamp* of the time the conversation was conducted (if available), the *source* of the conversation i.e., from what dataset it was taken, *user_metadata* which contains demographic information of the user such as location (if available), and *conversation_metadata* that contains additional information regarding the conversation, e.g., language, user-feedback and more.

3 Plugin Design and Architecture

In the following section, we describe the design choices of the ShareLM plugin and the motivations

²<https://huggingface.co/docs/hub/datasets-gated>

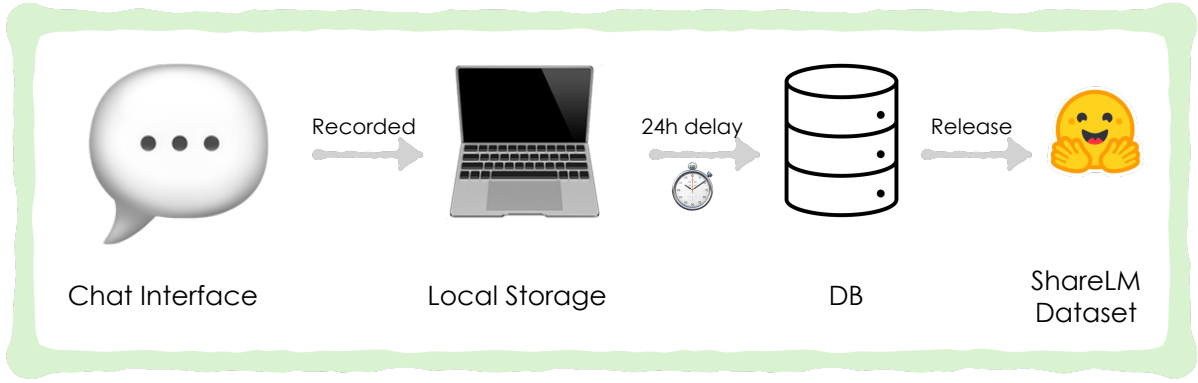


Figure 2: **1)** The user and model responses are periodically queried and collected. **2)** Each new conversation is assigned a unique ID, a timestamp, and the current URL. The conversation is stored in a local database. **3)** Upon a 24-hour delay, the conversations in the local database are posted to the server via a REST API, accompanied by the user ID and user/conversation metadata if available. **4)** An updated version of the dataset is released periodically.

behind them. We start by outlining the leading principles, and then describe the implementation.

3.1 Main Principles

Taking inspiration but more importantly lessons from the existing data collection platforms (see §7), we opt to design our plugin in accordance with the following principles:

1. **Easy Usage.** The plugin should be 'transparent' to the user, i.e., its basic functionality should not require any extra effort from the user.
2. **Users own their data.** The plugin merely helps in sharing and providing an open license to the data that the user creates and owns.
3. **Enhanced User Control.** The user should be able to manage their data on their own, e.g., deleting unwanted conversations.
4. **Privacy.** The plugin must conform to established privacy standards.
5. **Inclusive Models List.** Our plugin should be a mediator for other platforms, potentially supporting every model out there.

These principles guided us through the plugin development, from the decision to implement it as a plugin, to the finer details such as the delayed upload feature.

3.2 System Architecture

Upon installing the plugin and confirming the terms of use, the user is assigned a randomly generated user ID. We do not require the user to register and

log in, as we want to avoid unnecessary complications.

The plugin works by identifying certain elements in the web page XML, according to the chat platform in use. Currently, the plugin supports *Gradio*³, a web interface for various demos including chats, *ChatUI*⁴, a web interface for chats, *ChatGPT* and *Claude*. Those were chosen due to their popularity, e.g., Gradio and ChatUI are in frequent use in Huggingface Spaces⁵ and the ChatBot Arena (see §7). Nevertheless, adding support to new web platforms is easy.

The plugin flow operates as follows (see Fig. 2): The user and model responses are periodically queried and collected, together with thumb-up/down notions if available. A check is performed to determine whether the current conversation is a new one or rather a continuation of the previous one. Each new conversation is assigned a unique ID, a timestamp, and the current URL. The last is used to recognize what model the user was interacting with. The conversation is stored in a local database.

Upon a 24-hour delay, the conversations in the local database are posted to the server via a REST API, accompanied by the user ID and user/conversation metadata if available.

In turn, the server runs an anonymization script⁶ on the conversation's content, to remove names, addresses, phone numbers, and more. We note that as part of the plugin terms of use, we ask users to avoid sharing conversations with such

³<https://www.gradio.app/>

⁴<https://huggingface.co/docs/chat-ui/index>

⁵<https://huggingface.co/spaces>

⁶<https://pypi.org/project/anonymization/>

identifying details. The anonymization script is another line of protection, but no text shared should be assumed fully anonymous (Narayanan and Shmatikov, 2008). The server adds the new conversations to a PostgreSQL database.

Periodically, we release an updated version of the dataset⁷. In the future, we plan to employ a fully automated release process, but for now, we validate it manually before uploading it for quality control.

4 The Plugin UI

We describe the plugin UI components and usage.

4.1 Terms of Use

To activate the plugin after installation, the user needs to confirm the terms of use. The terms of use are available through the plugin popup (see §4.3), or the recording banner while in a supported demo (see next §4.2). We ask the users to avoid sharing conversations with identifying/sensitive content (names, e-mail addresses, etc.), as the content of the conversations will be publicly released. The full terms are available in the plugin repository and in App. §A.

4.2 The Recording Banner

The recording banner (see Fig. 3) is a thin strip at the top of the tab. The recording banner is available when the web page contains a supported demo interface (see §3.2). Seeing whether the current demo is supported is also possible through the extension icon. The icon is green when a supported interface exists, and gray otherwise.

The main role of the recording banner is to inform the user their conversations are recorded. In addition, it can be used to pause the conversation sharing. Clicking on the "Click Here to Stop Sharing" button will turn off the conversations collection (see Fig. 4). This is useful when conducting a conversation with identifying information that should be kept private.

With the recording banner, we balance between ease of use and control. We do not want to tire the user and require them to press buttons in order to record each conversation. On the other hand, we want the user to be aware that their conversations are recorded. Thus, the recording banner is designed to be visible but not interfere with normal use.

⁷<https://huggingface.co/datasets/shachardon/ShareLM>

4.3 The Popup

The plugin popup (see Fig. 1) is where the more advanced features are concentrated.

Demographic Details Form. Clicking on the down arrow at the top of the popup window opens a form of demographic details (Age, Country, and Gender). LLMs suffer from limited coverage of diverse human demographics in their training data, as their data usually comes from English speakers from narrow communities (Pavlick et al., 2014). Filling this form is voluntary, and can be of great help for studies focusing on diversity.

Chat Responses Counter The counter indicates the number of chat responses that have been shared (posted to the server) so far. Chats that are still stored locally are not included. The counter helps the user keep track of the size of their contribution.

Saved Conversations Table. The saved conversations table contains all the user’s recorded conversations from the last 24 hours. Clicking on a conversation extends it such that its full content is visible. The thumbs-up/down are used to rate the satisfaction of the user from the conversation as a whole. Rating the conversation and providing ‘human feedback’ is not mandatory, but it has great merit. Human feedback is a valuable resource for model development, as it allows better alignment of the model to human users’ preferences. Clicking on the red X button will delete the conversation from the local database, without it ever leaving the user storage. Asking to delete past conversations through the contact form (available at the bottom of the popup) is always possible, but we note that after the dataset was already released it is very likely that someone has already downloaded and saved an old version of it. Under the table, there are the buttons “Download” and “Publish Now”. The “Download” button allows the user to download a CSV file with all the conversations that are still in the local storage. This aim to strengthen the user’s ownership of their data. The “Publish Now” button empties the local storage and publishes the conversations immediately.

Frequently Asked Questions. Under the conversations table, we include a frequently asked questions section, to answer common questions regarding the plugin (see Fig. 6). There, we address questions regarding privacy (e.g., *Will it be possible to identify me by my conversations?*), li-

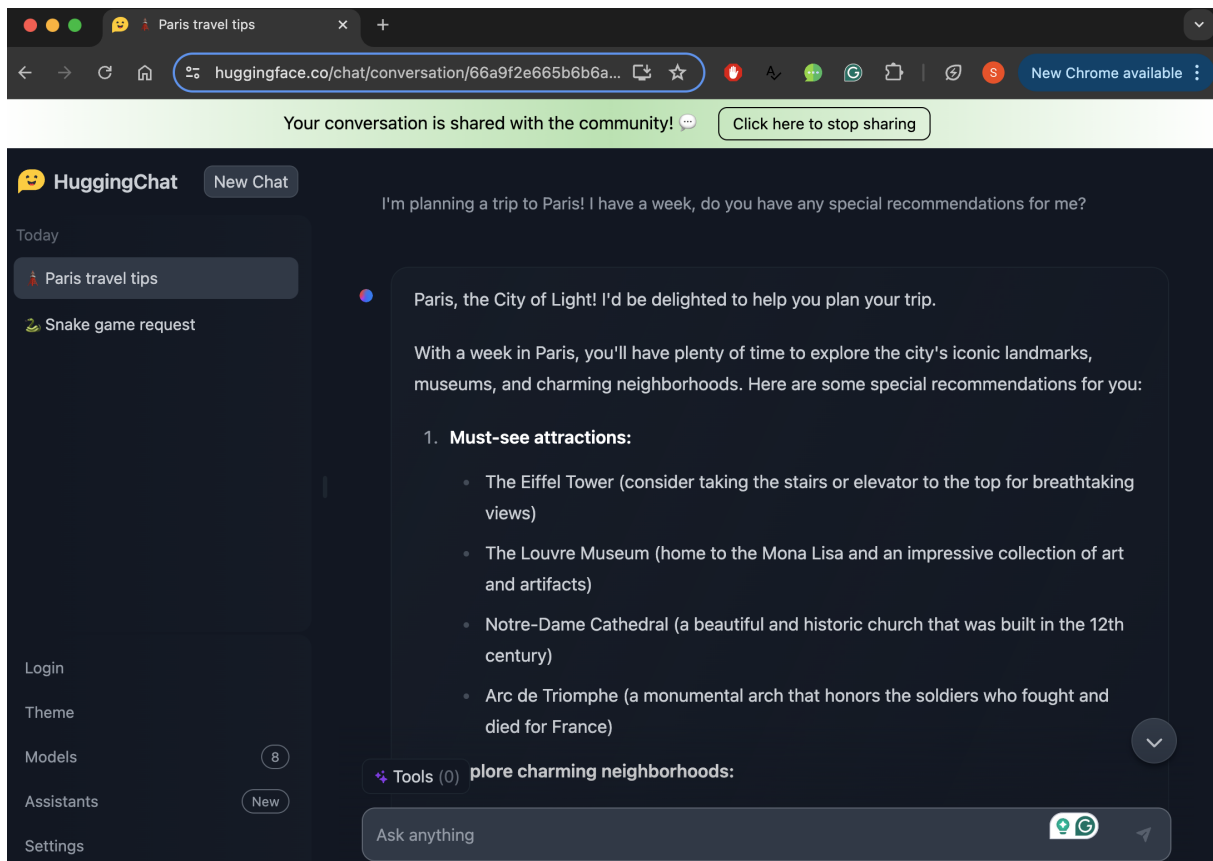


Figure 3: The recording banner is at the top of the window, indicating that the current chat demo (here ChatUI) is supported by the plugin and that the current conversation is recorded. Clicking on the "Click here to stop sharing" button will pause the conversation's recording.

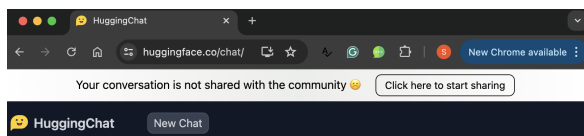


Figure 4: The conversation collection is paused. Clicking on the "Click here to start sharing" button will start the conversation's recording.

cense (*Would you share the dataset? With what license?*), ownership (*How can I ask to remove all my conversations from the dataset?*) among others.

Contact Form. The contact form is used to request to remove already published conversations from the dataset. One can ask to remove their own conversation, or use the form to report others' conversations that violate the terms of use. When a user asks to remove their own conversations, they will be asked to include their user ID for identity verification. For that, they can use the 'Copy User ID' button which copies the user ID to their clipboard.

5 Providing Human Feedback

As was already mentioned in §4.3, in addition to collecting conversations the plugin can be used also for rating them. Providing feedback for a given conversation can be done in two manners. The first is through the plugin popup. As shown in Figure 1, after conducting the conversations, the user can mark their conversations with thumbs up/down to express their (dis)satisfaction with the entire conversation. The other way to provide feedback is through the chat interface in real-time, as demonstrated in Figure 5. The user can click the thumbs-up/down buttons separately for each model response. This allows better feedback granularity, and is also sometimes easier, as it does not require the user to go over their conversation again, but is instead done at the time of the interaction. We note that the per-response option is currently available for the ChatUI interface only.

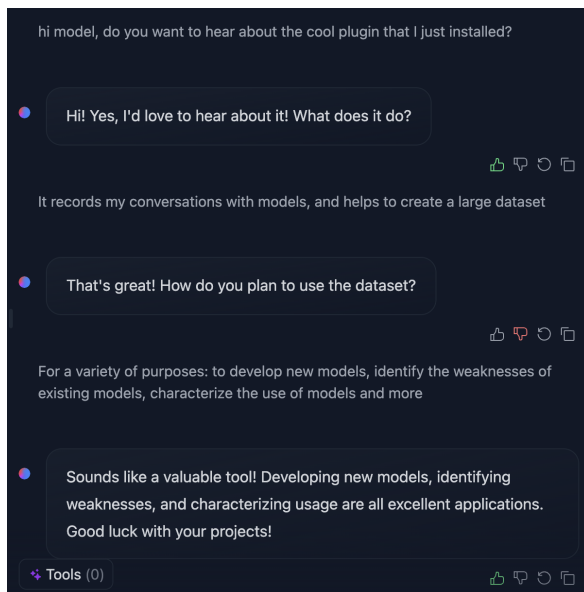


Figure 5: Providing feedback through the chat interface. The user can rate each response separately, at the time of the interaction.

6 User Study

We conducted a user study to evaluate the plugin. We asked 10 participants to install and experiment with the plugin. On a scale of 1 (poor) to 5 (great), 9 out of 10 participants described the installation experience as 5, and the average score was 4.8. Some of them elaborated, saying that *It was straight forwards, self explanatory and Smooth sailing, really easy and nice*. The participants described the experience of using the plugin for the first time with an average score of 4.7. Half of the participants reported that they used the plugin popup to rate or delete some of their conversations. The participants described the UI with an average score of 4.7, saying that it is *Really responsive, quick, and neatly designed and Easy to like a convo, to delete, and to understand the flow*. One of the participants said that *Its refresh time is long*. When asked how often do they use open models in their day-to-day activities on a scale of 1 (never) to 5 (all the time), the average score was 2.7.

Figure 7 shows the word cloud for the first 1000 conversations collected by the plugin. It seems that coding is the main use case. The average number of responses per conversation is 2.7, a bit higher than the average for the LMSYS-Chat-1M dataset (LMSYS reports an average length of 2, see §7 for more details about this dataset).

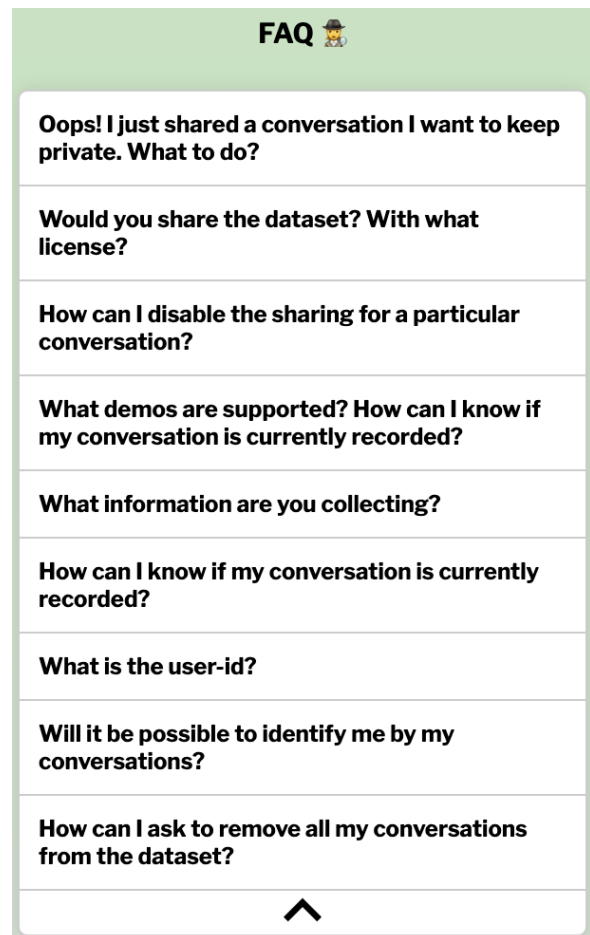


Figure 6: The frequently asked questions section (in the popup window). Provides answers to common questions regarding the plugin.

7 Previous Work

ShareGPT⁸, a plugin for collecting and sharing conversations specifically with ChatGPT, is the closest to ours. Although not active these days, the ShareGPT plugin collected over 400,000 conversations and 90,000 of them were published as a dataset before its API was shut down. Another effort is Collective Cognition⁹, a platform for collecting and tagging conversations with ChatGPT, also not active anymore. Unlike ShareGPT and Collective Cognition, our plugin is not limited to ChatGPT but rather focuses on open and closed models. It is also easier to use and does not require the user to actively click buttons to share each conversation.

The LMSYS’s Chatbot Arena (Zheng et al., 2023) hosts various models, both open and closed,

⁸<https://sharegpt.com/>

⁹<https://huggingface.co/datasets/CollectiveCognition/chats-data-2023-10-16?row=11>

[The dynamics of iterative prompting with midjourney](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 4146–4161, Singapore. Association for Computational Linguistics.

Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, Anirudh Goyal, Anthony Hartshorn, Aobo Yang, Archi Mitra, Archie Sravankumar, Artem Korenev, Arthur Hinsvark, Arun Rao, Aston Zhang, Aurelien Rodriguez, Austen Gregerson, Ava Spataru, Baptiste Roziere, Bethany Biron, Binh Tang, Bobbie Chern, Charlotte Caucheteux, Chaya Nayak, Chloe Bi, Chris Marra, Chris McConnell, Christian Keller, Christophe Touret, Chunyang Wu, Corinne Wong, Cristian Canton Ferrer, Cyrus Nikolaidis, Damien Al-lonsius, Daniel Song, Danielle Pintz, Danny Livshits, David Esiobu, Dhruv Choudhary, Dhruv Mahajan, Diego Garcia-Olano, Diego Perino, Dieuwke Hupkes, Egor Lakomkin, Ehab AlBadawy, Elina Lobanova, Emily Dinan, Eric Michael Smith, Filip Radenovic, Frank Zhang, Gabriel Synnaeve, Gabrielle Lee, Georgia Lewis Anderson, Graeme Nail, Gregoire Mialon, Guan Pang, Guillem Cucurell, Hailey Nguyen, Hannah Korevaar, Hu Xu, Hugo Touvron, Iliyan Zarov, Imanol Arrieta Ibarra, Isabel Kloumann, Ishan Misra, Ivan Evtimov, Jade Copet, Jaewon Lee, Jan Geffert, Jana Vranes, Jason Park, Jay Mahadeokar, Jeet Shah, Jelmer van der Linde, Jennifer Billock, Jenny Hong, Jenya Lee, Jeremy Fu, Jianfeng Chi, Jianyu Huang, Jiawen Liu, Jie Wang, Jiecao Yu, Joanna Bitton, Joe Spisak, Jongsoo Park, Joseph Rocca, Joshua Johnstun, Joshua Saxe, Junteng Jia, Kalyan Vasuden Alwala, Kartikeya Upasani, Kate Plawiak, Ke Li, Kenneth Heafield, Kevin Stone, Khalid El-Arini, Krithika Iyer, Kshitiz Malik, Kuenley Chiu, Kunal Bhalla, Lauren Rantala-Yeary, Laurens van der Maaten, Lawrence Chen, Liang Tan, Liz Jenkins, Louis Martin, Lovish Madaan, Lubo Malo, Lukas Blecher, Lukas Landzaat, Luke de Oliveira, Madeline Muzzi, Mahesh Pasupuleti, Mannat Singh, Manohar Paluri, Marcin Kardas, Mathew Oldham, Mathieu Rita, Maya Pavlova, Melanie Kambadur, Mike Lewis, Min Si, Mitesh Kumar Singh, Mona Hassan, Naman Goyal, Narjes Torabi, Nikolay Bashlykov, Nikolay Bogoychev, Niladri Chatterji, Olivier Duchenne, Onur Çelebi, Patrick Alrassy, Pengchuan Zhang, Pengwei Li, Petar Vasic, Peter Weng, Prajjwal Bhargava, Pratik Dubal, Praveen Krishnan, Punit Singh Koura, Puxin Xu, Qing He, Qingxiao Dong, Ragavan Srinivasan, Raj Ganapathy, Ramon Calderer, Ricardo Silveira Cabral, Robert Stojnic, Roberta Raileanu, Rohit Girdhar, Rohit Patel, Romain Sauvestre, Ronnie Polidoro, Roshan Sumbaly, Ross Taylor, Ruan Silva, Rui Hou, Rui Wang, Saghar Hosseini, Sahana Chennabasappa, Sanjay Singh, Sean Bell, Seohyun Sonia Kim, Sergey Edunov, Shaoliang Nie, Sharan Narang, Sharath Rapparthi, Sheng Shen, Shengye Wan, Shruti Bhosale, Shun Zhang, Simon Vandenhende, Soumya Batra, Spencer Whitman, Sten Sootla, Stephane Collot, Suchin Gururangan, Sydney Borodinsky, Tamar Herman, Tara

Fowler, Tarek Sheasha, Thomas Georgiou, Thomas Scialom, Tobias Speckbacher, Todor Mihaylov, Tong Xiao, Ujjwal Karn, Vedanuj Goswami, Vibhor Gupta, Vignesh Ramanathan, Viktor Kerkez, Vincent Gonguet, Virginie Do, Vish Vogeti, Vladan Petrovic, Weiwei Chu, Wenhan Xiong, Wenyin Fu, Whitney Meers, Xavier Martinet, Xiaodong Wang, Xiaoqing Ellen Tan, Xinfeng Xie, Xuchao Jia, Xuewei Wang, Yaelle Goldschlag, Yashesh Gaur, Yasmine Babaei, Yi Wen, Yiwen Song, Yuchen Zhang, Yue Li, Yuning Mao, Zacharie Delpierre Coudert, Zheng Yan, Zhengxing Chen, Zoe Papakipos, Aaditya Singh, Aaron Grattafiori, Abha Jain, Adam Kelsey, Adam Shajnfeld, Adithya Gangidi, Adolfo Victoria, Ahuva Goldstand, Ajay Menon, Ajay Sharma, Alex Boesenberg, Alex Vaughan, Alexei Baeviski, Allie Feinstein, Amanda Kallet, Amit Sangani, Anam Yunus, Andrei Lupu, Andres Alvarado, Andrew Caples, Andrew Gu, Andrew Ho, Andrew Poulton, Andrew Ryan, Ankit Ramchandani, Annie Franco, Aparajita Saraf, Arkabandhu Chowdhury, Ashley Gabriel, Ashwin Bharambe, Assaf Eisenman, Azadeh Yazdan, Beau James, Ben Maurer, Benjamin Leonhardi, Bernie Huang, Beth Loyd, Beto De Paola, Bhargavi Paranjape, Bing Liu, Bo Wu, Boyu Ni, Braden Hancock, Bram Wasti, Brandon Spence, Brani Stojkovic, Brian Gamido, Britt Montalvo, Carl Parker, Carly Burton, Catalina Mejia, Changan Wang, Changkyu Kim, Chao Zhou, Chester Hu, Ching-Hsiang Chu, Chris Cai, Chris Tindal, Christoph Feichtenhofer, Damon Civin, Dana Beaty, Daniel Kreymier, Daniel Li, Danny Wyatt, David Adkins, David Xu, Davide Testuggine, Delia David, Devi Parikh, Diana Liskovich, Didem Foss, Dingkan Wang, Duc Le, Dustin Holland, Edward Dowling, Eissa Jamil, Elaine Montgomery, Eleonora Presani, Emily Hahn, Emily Wood, Erik Brinkman, Esteban Arcaute, Evan Dunbar, Evan Smothers, Fei Sun, Felix Kreuk, Feng Tian, Firat Ozgenel, Francesco Caggioni, Francisco Guzmán, Frank Kanayet, Frank Seide, Gabriela Medina Florez, Gabriella Schwarz, Gada Badeer, Georgia Sweet, Gil Halpern, Govind Thattai, Grant Herman, Grigory Sizov, Guangyi, Zhang, Guna Lakshminarayanan, Hamid Shojanazeri, Han Zou, Hannah Wang, Hanwen Zha, Haroun Habeeb, Harrison Rudolph, Helen Suk, Henry Aspegren, Hunter Goldman, Igor Molybog, Igor Tufanov, Irina-Elena Veliche, Itai Gat, Jake Weissman, James Geboski, James Kohli, Japhet Asher, Jean-Baptiste Gaya, Jeff Marcus, Jeff Tang, Jennifer Chan, Jenny Zhen, Jeremy Reizenstein, Jeremy Teboul, Jessica Zhong, Jian Jin, Jingyi Yang, Joe Cummings, Jon Carvill, Jon Shepard, Jonathan McPhie, Jonathan Torres, Josh Ginsburg, Junjie Wang, Kai Wu, Kam Hou U, Karan Saxena, Karthik Prasad, Kartikay Khandelwal, Katayoun Zand, Kathy Matosich, Kaushik Veeraraghavan, Kelly Michelena, Keqian Li, Kun Huang, Kunal Chawla, Kushal Lakhotia, Kyle Huang, Lailin Chen, Lakshya Garg, Lavender A, Leandro Silva, Lee Bell, Lei Zhang, Liangpeng Guo, Licheng Yu, Liron Moshkovich, Luca Wehrstedt, Madian Khabisa, Manav Avalani, Manish Bhatt, Maria Tsim-poukelli, Martynas Mankus, Matan Hasson, Matthew Lennie, Matthias Reso, Maxim Groshev, Maxim

- Naumov, Maya Lathi, Meghan Keneally, Michael L. Seltzer, Michal Valko, Michelle Restrepo, Mihir Patel, Mik Vyatskov, Mikayel Samvelyan, Mike Clark, Mike Macey, Mike Wang, Miquel Jubert Hermoso, Mo Metanat, Mohammad Rastegari, Munish Bansal, Nandhini Santhanam, Natascha Parks, Natasha White, Navyata Bawa, Nayan Singhal, Nick Egebo, Nicolas Usunier, Nikolay Pavlovich Laptev, Ning Dong, Ning Zhang, Norman Cheng, Oleg Chernoguz, Olivia Hart, Omkar Salpekar, Ozlem Kalinli, Parkin Kent, Parth Parekh, Paul Saab, Pavan Balaji, Pedro Rittner, Philip Bontrager, Pierre Roux, Piotr Dollar, Polina Zvyagina, Prashant Ratan-chandani, Pritish Yuvraj, Qian Liang, Rachad Alao, Rachel Rodriguez, Rafi Ayub, Raghotham Murthy, Raghu Nayani, Rahul Mitra, Raymond Li, Rebekkah Hogan, Robin Battey, Rocky Wang, Rohan Maheswari, Russ Howes, Ruty Rinott, Sai Jayesh Bondu, Samyak Datta, Sara Chugh, Sara Hunt, Sargun Dhillon, Sasha Sidorov, Satadru Pan, Saurabh Verma, Seiji Yamamoto, Sharadh Ramaswamy, Shaun Lindsay, Shaun Lindsay, Sheng Feng, Shenghao Lin, Shengxin Cindy Zha, Shiva Shankar, Shuqiang Zhang, Shuqiang Zhang, Sinong Wang, Sneha Agarwal, Soji Sajuyigbe, Soumith Chintala, Stephanie Max, Stephen Chen, Steve Kehoe, Steve Satterfield, Sudarshan Govindaprasad, Sumit Gupta, Sungmin Cho, Sunny Virk, Suraj Subramanian, Sy Choudhury, Sydney Goldman, Tal Remez, Tamar Glaser, Tamara Best, Thilo Kohler, Thomas Robinson, Tianhe Li, Tianjun Zhang, Tim Matthews, Timothy Chou, Tzook Shaked, Varun Vontimitta, Victoria Ajayi, Victoria Montanez, Vijai Mohan, Vinay Satish Kumar, Vishal Mangla, Vlad Ionescu, Vlad Poenaru, Vlad Tiberiu Mihailescu, Vladimir Ivanov, Wei Li, Wenchen Wang, Wenwen Jiang, Wes Bouaziz, Will Constable, Xiao Cheng Tang, Xiaofang Wang, Xiaoqian Wu, Xiaolan Wang, Xide Xia, Xilun Wu, Xinbo Gao, Yanjun Chen, Ye Hu, Ye Jia, Ye Qi, Yenda Li, Yilin Zhang, Ying Zhang, Yossi Adi, Youngjin Nam, Yu, Wang, Yuchen Hao, Yundi Qian, Yuzi He, Zach Rait, Zachary DeVito, Zef Rosnbrick, Zhaoduo Wen, Zhenyu Yang, and Zhiwei Zhao. 2024. *The llama 3 herd of models*. *Preprint*, arXiv:2407.21783.
- Braden Hancock, Antoine Bordes, Pierre-Emmanuel Mazare, and Jason Weston. 2019. *Learning from dialogue after deployment: Feed yourself, chatbot!* In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3667–3684, Florence, Italy. Association for Computational Linguistics.
- Or Honovich, Thomas Scialom, Omer Levy, and Timo Schick. 2023. *Unnatural instructions: Tuning language models with (almost) no human labor*. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 14409–14428, Toronto, Canada. Association for Computational Linguistics.
- Hannah Rose Kirk, Alexander Whitefield, Paul Röttger, Andrew Bean, Katerina Margatina, Juan Ciro, Rafael Mosquera, Max Bartolo, Adina Williams, He He, et al. 2024. The prism alignment project: What participatory, representative and individualised human feedback reveals about the subjective and multicultural alignment of large language models. *arXiv preprint arXiv:2404.16019*.
- Andreas Köpf, Yannic Kilcher, Dimitri von Rütte, Sotiris Anagnostidis, Zhi Rui Tam, Keith Stevens, Abdullah Barhoum, Duc Nguyen, Oliver Stanley, Richárd Nagyfi, et al. 2024. Openassistant conversations-democratizing large language model alignment. *Advances in Neural Information Processing Systems*, 36.
- Reiichiro Nakano, Jacob Hilton, Suchir Balaji, Jeff Wu, Long Ouyang, Christina Kim, Christopher Hesse, Shantanu Jain, Vineet Kosaraju, William Saunders, et al. 2021. Webgpt: Browser-assisted question-answering with human feedback. *arXiv preprint arXiv:2112.09332*.
- Arvind Narayanan and Vitaly Shmatikov. 2008. Robust de-anonymization of large sparse datasets. In *2008 IEEE Symposium on Security and Privacy (sp 2008)*, pages 111–125. IEEE.
- OpenAI, Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altschmidt, Sam Altman, Shyamal Anadkat, Red Avila, Igor Babuschkin, Suchir Balaji, Valerie Balcom, Paul Baltescu, Haiming Bao, Mohammad Bavarian, Jeff Belgum, Irwan Bello, Jake Berdine, Gabriel Bernadett-Shapiro, Christopher Berner, Lenny Bogdonoff, Oleg Boiko, Madelaine Boyd, Anna-Luisa Brakman, Greg Brockman, Tim Brooks, Miles Brundage, Kevin Button, Trevor Cai, Rosie Campbell, Andrew Cann, Brittany Carey, Chelsea Carlson, Rory Carmichael, Brooke Chan, Che Chang, Fotis Chantzis, Derek Chen, Sully Chen, Ruby Chen, Jason Chen, Mark Chen, Ben Chess, Chester Cho, Casey Chu, Hyung Won Chung, Dave Cummings, Jeremiah Currier, Yunxing Dai, Cory Decareaux, Thomas Degry, Noah Deutsch, Damien Deville, Arka Dhar, David Dohan, Steve Dowling, Sheila Dunning, Adrien Ecoffet, Atty Eleti, Tyna Eloundou, David Farhi, Liam Fedus, Niko Felix, Simón Posada Fishman, Juston Forte, Isabella Fulford, Leo Gao, Elie Georges, Christian Gibson, Vik Goel, Tarun Gogineni, Gabriel Goh, Rapha Gontijo-Lopes, Jonathan Gordon, Morgan Grafstein, Scott Gray, Ryan Greene, Joshua Gross, Shixiang Shane Gu, Yufei Guo, Chris Hallacy, Jesse Han, Jeff Harris, Yuchen He, Mike Heaton, Johannes Heidecke, Chris Hesse, Alan Hickey, Wade Hickey, Peter Hoeschele, Brandon Houghton, Kenny Hsu, Shengli Hu, Xin Hu, Joost Huizinga, Shantanu Jain, Shawn Jain, Joanne Jang, Angela Jiang, Roger Jiang, Haozhun Jin, Denny Jin, Shino Jomoto, Billie Jonn, Heewoo Jun, Tomer Kaftan, Łukasz Kaiser, Ali Kamali, Ingmar Kanitscheider, Nitish Shirish Keskar, Tabarak Khan, Logan Kilpatrick, Jong Wook Kim, Christina Kim, Yongjik Kim, Jan Hendrik Kirchner, Jamie Kiros, Matt Knight, Daniel Kokotajlo, Łukasz Kondraciuk, Andrew Kondrich, Aris Konstantinidis, Kyle Kosic, Gretchen Krueger, Vishal

- Kuo, Michael Lampe, Ikai Lan, Teddy Lee, Jan Leike, Jade Leung, Daniel Levy, Chak Ming Li, Rachel Lim, Molly Lin, Stephanie Lin, Mateusz Litwin, Theresa Lopez, Ryan Lowe, Patricia Lue, Anna Makanju, Kim Malfacini, Sam Manning, Todor Markov, Yaniv Markovski, Bianca Martin, Katie Mayer, Andrew Mayne, Bob McGrew, Scott Mayer McKinney, Christine McLeavey, Paul McMillan, Jake McNeil, David Medina, Aalok Mehta, Jacob Menick, Luke Metz, Andrey Mishchenko, Pamela Mishkin, Vinnie Monaco, Evan Morikawa, Daniel Mossing, Tong Mu, Mira Murati, Oleg Murk, David Mély, Ashvin Nair, Reiichiro Nakano, Rajeev Nayak, Arvind Neelakantan, Richard Ngo, Hyeonwoo Noh, Long Ouyang, Cullen O’Keefe, Jakub Pachocki, Alex Paino, Joe Palermo, Ashley Pantuliano, Giambatista Parascandolo, Joel Parish, Emy Parparita, Alex Passos, Mikhail Pavlov, Andrew Peng, Adam Perelman, Filipe de Avila Belbute Peres, Michael Petrov, Henrique Ponde de Oliveira Pinto, Michael, Pokorny, Michelle Pokrass, Vitchyr H. Pong, Tolly Powell, Alethea Power, Boris Power, Elizabeth Proehl, Raul Puri, Alec Radford, Jack Rae, Aditya Ramesh, Cameron Raymond, Francis Real, Kendra Rimbach, Carl Ross, Bob Rotsted, Henri Roussez, Nick Ryder, Mario Saltarelli, Ted Sanders, Shibani Santurkar, Girish Sastry, Heather Schmidt, David Schnurr, John Schulman, Daniel Selsam, Kyla Sheppard, Toki Sherbakov, Jessica Shieh, Sarah Shoker, Pranav Shyam, Szymon Sidor, Eric Sigler, Maddie Simens, Jordan Sitkin, Katarina Slama, Ian Sohl, Benjamin Sokolowsky, Yang Song, Natalie Staudacher, Felipe Petroski Such, Natalie Summers, Ilya Sutskever, Jie Tang, Nikolas Tezak, Madeleine B. Thompson, Phil Tillet, Amin Tootoonchian, Elizabeth Tseng, Preston Tuggle, Nick Turley, Jerry Tworek, Juan Felipe Cerón Uribe, Andrea Vallone, Arun Vijayvergiya, Chelsea Voss, Carroll Wainwright, Justin Jay Wang, Alvin Wang, Ben Wang, Jonathan Ward, Jason Wei, CJ Weinmann, Akila Welihinda, Peter Welinder, Jiayi Weng, Lilian Weng, Matt Wiethoff, Dave Willner, Clemens Winter, Samuel Wolrich, Hannah Wong, Lauren Workman, Sherwin Wu, Jeff Wu, Michael Wu, Kai Xiao, Tao Xu, Sarah Yoo, Kevin Yu, Qiming Yuan, Wojciech Zaremba, Rowan Zellers, Chong Zhang, Marvin Zhang, Shengjia Zhao, Tianhao Zheng, Juntang Zhuang, William Zhuk, and Barret Zoph. 2024. [Gpt-4 technical report](#). *Preprint*, arXiv:2303.08774.
- Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. 2022. Training language models to follow instructions with human feedback. *Advances in neural information processing systems*, 35:27730–27744.
- Ellie Pavlick, Matt Post, Ann Irvine, Dmitry Kachaev, and Chris Callison-Burch. 2014. The language demographics of amazon mechanical turk. *Transactions of the Association for Computational Linguistics*, 2:79–92.
- Luiza Pozzobon, Beyza Ermis, Patrick Lewis, and Sara Hooker. 2023. [Goodtriever: Adaptive toxicity mitigation with retrieval-augmented models](#). *Preprint*, arXiv:2310.07589.
- Yizhong Wang, Yeganeh Kordi, Swaroop Mishra, Alisa Liu, Noah A. Smith, Daniel Khashabi, and Hannaneh Hajishirzi. 2023. [Self-instruct: Aligning language models with self-generated instructions](#). *Preprint*, arXiv:2212.10560.
- Jason Weston, Antoine Bordes, Sumit Chopra, Alexander M Rush, Bart Van Merriënboer, Armand Joulin, and Tomas Mikolov. 2015. Towards ai-complete question answering: A set of prerequisite toy tasks. *arXiv preprint arXiv:1502.05698*.
- Saizheng Zhang, Emily Dinan, Jack Urbanek, Arthur Szlam, Douwe Kiela, and Jason Weston. 2018. [Personalizing dialogue agents: I have a dog, do you have pets too?](#) In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2204–2213, Melbourne, Australia. Association for Computational Linguistics.
- Wenting Zhao, Xiang Ren, Jack Hessel, Claire Cardie, Yejin Choi, and Yuntian Deng. 2024. Wildchat: 1m chatgpt interaction logs in the wild. *arXiv preprint arXiv:2405.01470*.
- Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric P Xing, Hao Zhang, Joseph E. Gonzalez, and Ion Stoica. 2023. [Judging llm-as-a-judge with mt-bench and chatbot arena](#). *Preprint*, arXiv:2306.05685.

A Terms of Use

Inspired by the release of the ChatGPT, the open-source community recently began to develop open access models with increased transparency about their development. The next challenge for democratizing large language models is data.

This extension collects the conversations you are having with open large language models (“chatbots”). By using this extension, you are giving your permission to contribute your conversations’ content (both your side of the conversation, and the model’s) for creating an open-license chat-bot conversations dataset, a valuable resource for the open-source community. The conversations will be released with the most permissive license that is allowed by the specific model. This dataset will be a valuable resource for both model developers and researchers. Specifically, we plan to use this dataset to study and improve the nature of human-model interaction.

The extension supports a couple of chatbots demos, mostly within Huggingface Spaces

(<https://huggingface.co/spaces>). You will see a banner on the top of the demo page indicating it. You can choose not to share a particular conversation by clicking the ‘do not share’ button. As an additional precaution, the conversations are not posted to the database immediately. You can see the conversations from the last 24 hours in the extension popup window and remove them. To stop sharing your conversations permanently, please disable or remove the extension. Note that removing the extension does not delete the conversations you have already made.

Along with the conversation’s content, we are collecting the URL (to identify the model), GMT time and an anonymous user-id. Optionally, you can fill some demographic data (age, location, gender) and rate your satisfaction. We are not collecting any identifying metadata (such as IP address, local time, browser type, etc.). However, it is possible that you will be identified by the content of your conversations. Therefore, please avoid sharing conversations with Identifying/sensitive content (names, e-mail addresses, etc.), as the content of your conversations will be publicly released. If you accidentally shared the content of a conversation you prefer to keep private, please fill the contact form so we will remove it (available in the extension popup). You can ask to remove all your conversations at any time, but please note that after the dataset was already released it is very likely that someone has already downloaded and saved an old version of it. You are encouraged to use this form also for reporting conversations that are copyrighted, defamatory, threatening to others, violating of others’ privacy, or that you view as harmful if released.

Please be advised that this extension is independently developed by us, and while we have put our best efforts into ensuring a smooth experience, it’s important to note that there might be bugs or unforeseen issues. Your feedback is valuable to us, so please feel free to report any issues you may encounter.

The research is conducted by Shachar Don-Yehiya, Leshem Choshen and Omri Abend at the Hebrew University.

For more questions, please contact us at shareLM.project@gmail.com.

Participation is from age 18 and over only.

Participation is voluntary. Thank you for your contribution!