ConvXAI: Delivering Heterogeneous AI Explanations via Conversations to Support Human-AI Scientific Writing

Hua Shen
huashen@umich.edu
University of Michigan
Pennsylvania State University
USA

Chieh-Yang Huang
chiehyang@psu.edu
Pennsylvania State University
USA

Tongshuang Wu
sherryw@cs.cmu.edu
Carnegie Mellon University
USA

Ting-Hao (Kenneth) Huang
txh710@psu.edu
Pennsylvania State University
USA

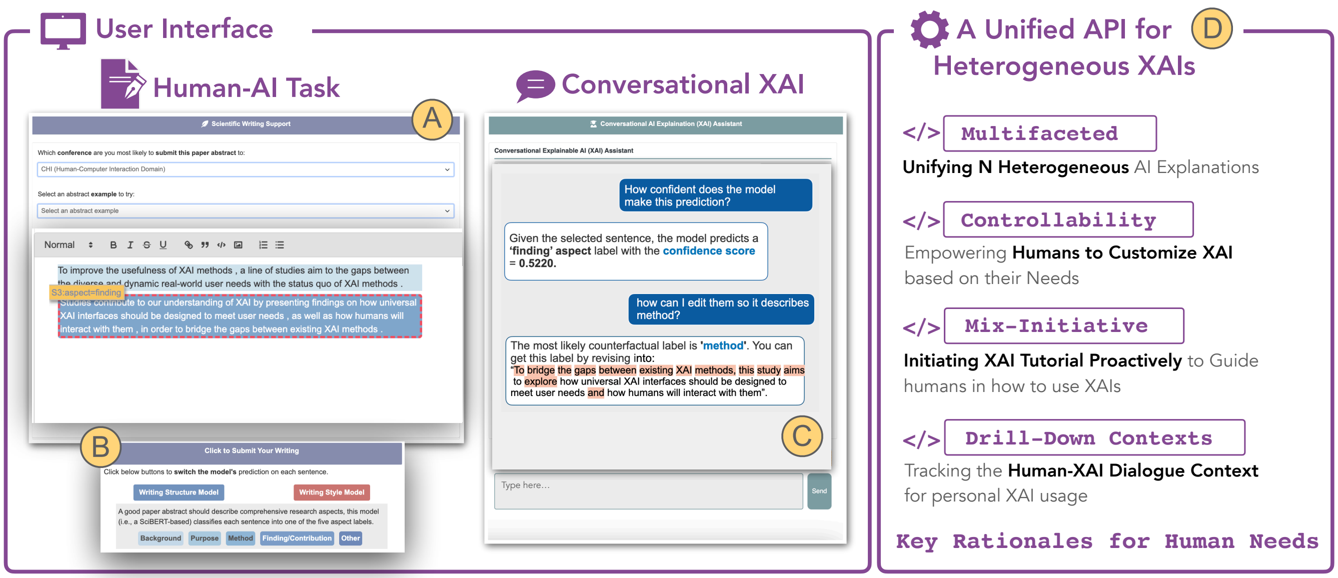

Figure 1: An overview of ConvXAI to support human-AI scientific writing with heterogeneous AI explanations via dialog. ConvXAI includes a front-end User Interface to (A) support human-AI collaborative task, (B) check AI models and predictions, and (C) inquire about heterogeneous AI explanations via dialogue. Also, ConvXAI involves a back-end deep learning server to generate AI predictions and explanations, which is embedded with (D) a unified API for generating heterogeneous AI explanations.

ABSTRACT

Despite a surge collection of XAI methods, users still struggle to obtain required AI explanations. Previous research suggests chatbots as dynamic solutions, but the effective design of conversational XAI agents for practical human needs remains under-explored. This paper focuses on Conversational XAI for AI-assisted scientific writing tasks. Drawing from human linguistic theories and formative studies, we identify four design rationales: "multifaceted", "controllability", "mix-initiative", "context-aware drill-down". We incorporate them into an interactive prototype, ConvXAI[1], which facilitates heterogeneous AI explanations for scientific writing through dialogue. In two studies with 21 users, ConvXAI outperforms a GUI-based baseline on improving human-perceived understanding and writing improvement. The paper further discusses the practical human usage patterns in interacting with ConvXAI for scientific co-writing[2].

CSCW '23 Companion, October 14–18, 2023, Minneapolis, MN, USA

ACM ISBN 979-8-4007-0129-0/23/10…$15.00
https://doi.org/10.1145/3584931.3607492

CCS CONCEPTS

• **Human-centered computing → Interactive systems and tools; Collaborative and social computing systems and tools.**

[1]See the ConvXAI system code at: https://github.com/huashen218/convxai.git.
[2]See a full paper of ConvXAI study at: https://arxiv.org/pdf/2305.09770.pdf.

KEYWORDS

Explainable AI, Conversational AI, Scientific Writing Support

ACM Reference Format:
Hua Shen, Chieh-Yang Huang, Tongshuang Wu, and Ting-Hao (Kenneth) Huang. 2023. ConvXAI: Delivering Heterogeneous AI Explanations via Conversations to Support Human-AI Scientific Writing. In *Computer Supported Cooperative Work and Social Computing (CSCW '23 Companion), October 14–18, 2023, Minneapolis, MN, USA.* ACM, New York, NY, USA, 4 pages. https://doi.org/10.1145/3584931.3607492

1 INTRODUCTION

Despite the potential of a surge collection of eXplainable AI (XAI) methods, a number of studies show that applying state-of-the-art XAI methods to real-world human tasks could not always help users better simulate model predictions, understand AI model mistakes, etc [6, 9]. To resolve the issues, researchers have explored the mismatch between real-world user demands and existing XAI methods. Shen and Huang [7], for instance, compare practical user questions with over 200 XAI studies and identify a bias in current methods towards certain types of XAI questions, neglecting others. Additionally, users also tend to have *multiple*, *dynamic* and sometimes *interdependent* questions on AI explanations [3, 4]. Addressing this array of questions necessitates an integration of heterogeneous AI explanations. Prior work has envisioned the concept of "explainability as a dialogue" to accommodate diverse user needs and mitigate cognitive load [3]. However, there is a dearth of exploration regarding the design of conversational XAI systems to meet practical user needs and understand user reactions.

In this paper, we investigate the potential of conversational XAI in the context of practical human-AI scientific writing, where we propose a conversational XAI system, ConvXAI. ConvXAI incorporates ten types of AI explanations into a unified dialog interface that empowers users to interactively ask various XAI questions about the AI predictions. Particularly, we augment ConvXAI with four design rationales collected from empirical formative studies with 7 users of diverse backgrounds and theoretic linguistic properties of human conversation: address various user questions ("multi-faceted"), actively provide XAI tutorials and suggestions ("mix-initiative"), empower users to dig into AI explanations ("context-aware drill-down"), and make flexible customization with details on-demand ("controllability").

We conducted two within-subject user studies with 21 users to compare with SelectXAI– the traditional GUI-based XAI system that displays all XAIs in a collapsible manner. Results show that users perceived ConvXAI to be more useful in understanding AI writing feedback and improving human writings. The results also validated the less cognitive load and effectiveness of the four user-oriented design principles. This work contributes insights into the design space of useful XAI in practice, reveals humans' XAI usage patterns and identifies opportunities for future XAI works.

2 CONVXAI SYSTEM

ConvXAI is a novel system to facilitate human-AI collaboration by providing conversational AI explanations, which empowers users to request heterogeneous AI explanations via a unified dialog interface. We apply ConvXAI on human-AI scientific writing task in this study and briefly describe its design and implementation below.

2.1 Delivering Heterogeneous Explanations via a Unified Dialogue Interface

Users ask for diverse XAI questions covering the whole AI lifecycle (*e.g.,* data, modeling, and evaluation) to understand the AI system [4, 7], while preferring the interactive interface to be concise for minimal cognitive load [5]. Therefore, we design ConvXAI to deliver heterogeneous explanations via a unified conversational interface. Specifically, we develop this conversational XAI pipeline based on the task-oriented Dialogue-State Architecture [1], where it mainly consists of four modules. Firstly, a *Natural Language Understanding* module parses any XAI user question into a pre-defined XAI user intent (*e.g.,* "explain top-k features", "explain similar examples"), which is then mapped to a specific XAI function out of 10 XAI types elaborated in Figure 2 (Table). Secondly, the *AI Explainers* module will generate AI explanations using cutting-edge XAI algorithms corresponding to the above XAI user intent. The third module, *Natural Language Generation*, then converts the generated explanation to a free-text response via the pre-defined natural language templates, and sends the response back to users via the dialogue interface. On top of the pipeline, we include a Global XAI State Tracker, to record users' turn-based conversational interactions, including user intent transitions and the users' customization on AI explanations. Overall, we design the conversational XAI pipeline to be model agnostic and XAI algorithm agnostic, which enables ConvXAI to be generalized to various AI and XAI methods.

2.2 Enabling User-oriented Multifaceted XAI with Interactive Customization

By combining the feedback from a formative study with 7 users of diverse backgrounds and the human conversational linguistic theories [2], we embed four user-oriented design rationales into ConvXAI. The ConvXAI (Figure 2) requires to be *Multifaceted*, *Mixed-initiative*, *Context-aware drill-down*, and *Controllability*. Concretely, ConvXAI enables users to ask ten types of **multi-faceted** AI explanations (shown in Figure 2 (Table)) in the conversation input panel (Figure 2A$_1$). To design ConvXAI to be **mixed-initiative**, we start the explanation dialog with a review summary of the AI writing structure model and style model's outputs. The users can select any one sentence in this review list to ask instance-wise XAI questions and start a conversation session on the sentence. Uniquely, to make it serve as proactive guidance towards more sophisticated XAI methods, ConvXAI adds an additional explanation type, *understand suggestion* – to explain AI suggestions and provide brief XAI tutorials. Also, ConvXAI initiates a prompt message *"to improve..."* (Figure 2A$_2$) with a subset of relevant XAIs, based on "guessing" what users would want to improve their writing at this point. To enable **context-aware drill down** (Figure 2A$_3$), ConvXAI leverages Global XAI State Tracker to store the history of user and XAI agent dialog, and generate the XAI response based on the previous dialog accordingly. Still, given the default XAI responses may not satisfy users' needs on customizing their own explanations in some cases, we, therefore, make the XAI agent proactively present hints for human **controllablity**, *e.g.,* "would you like to..." (at the bottom of Figure 2A$_3$. Figure 2A$_4$) for humans to customize XAIs based on their needs. By embedding these user-oriented XAI rationales into ConvXAI design, the system can thus be more useful for humans in practical AI-assisted writing tasks.

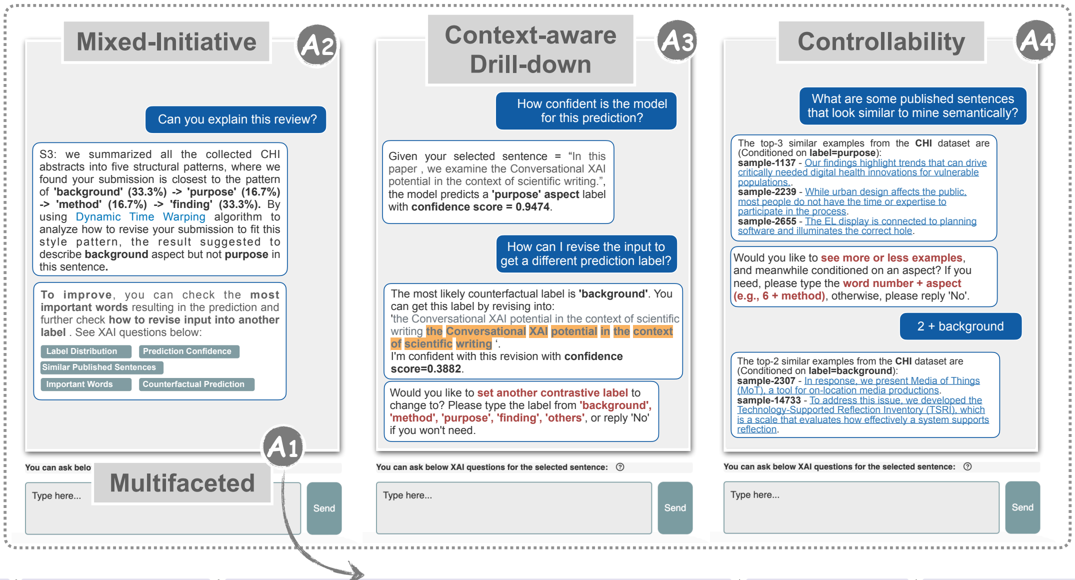

Stage	XAI Goal	User Question Samples	XAI Formats	Algorithm
1	Understand Data	1.What data did the system learn from?	Data Statistics	Data Sheets
		2.What's the range of the style quality scores?		
		3.How are the structure labels distributed?		
	Understand Model	4.What kind of models are used?	Model Description	Model Card
	Understand Instance	5.How confident is the model for this prediction?	Prediction Confidence	Model probability score
		6.What are some published sentences similar to mine semantically?	Similar Examples	NN-DOT
		7.Which words in this sentence are most important for prediction?	Feature Attribution	Integrated Gradient
	Improve Instance	8.How can I revise the input to get a different prediction label?	Counterfactual	GPT3 In-context Learning
2	Understand Data	9.What's the statistics of the sentence lengths?	Data Statistics	Data Sheets
	Understand Suggestion	10.Can you explain this sentence review?	XAI Tutorial	Template

Figure 2: An overview of four design rationales tailored for human use needs in ConvXAI. The ConvXAI dialogue flows are designed to follow the four principles of *"multifaceted"* (A_1), *"mixed-initiative"*(A_2), *"context-aware drill-down"*(A_3) and *"controllability"*(A_4). Particularly, for *"multifaceted"* (A_1), ConvXAI covers ten types of user questions (*i.e.,* Data Statistic, Feature Attribution, etc.) serving to five different XAI goals. Stage (1) shows eight XAIs included in the formative study, and Stage (2) indicates two added XAIs in ConvXAI.

2.3 Applying Conversational XAI to Human-AI Scientific Writing

We apply ConvXAI to human-AI collaborative writing on scientific papers, as the human-AI co-writing process consumes complex cognitive loads that can potentially inspire users to enquire more AI explanations. As shown in Figure 1, humans can submit their drafts (Figure 1A) to the editor to get AI models' feedback (Figure 1B). Then they can leverage the conversational AI explanations to *understand the writing models' feedback* (Figure 1C) (*i.e.,* including AI model integrated feedback), and further *improve and resubmit their writings iteratively*. Here, we incorporate two AI models to provide AI predictions on writing structure and style, respectively, and further integrate them into writing reviews.

3 USER STUDIES

We conducted two within-subjects human evaluation studies with 13 participants in the Task One and 8 participants rejoining in the Task Two. The users were recruited from university mailing list and required to have research writing experience. We asked users to compare ConvXAI against SelectXAI, a GUI-based XAI system that statically displays all the XAI formats at one-time. The user study aimed to investigate *if* ConvXAI *can help users better understand the AI writing feedback and further improve the writing artifacts accordingly*. We asked each participant to edit two paper abstracts with the help of ConvXAI and SelectXAI, respectively. Participants were then asked to rate their experience using 5-point Likert scale in the survey. We particularly designed the Task One to be an open-ended writing task to evaluate the effectiveness of user-oriented design in the system, and Task Two as a well-defined writing task to investigate how systems can help users improve their scientific writing process and output in practice [8].

We summarize participants' ratings on the two systems, ConvXAI and SelectXAI, in Figure 3. We performed the non-parametric Wilcoxon signed-rank test to compare users' nominal Likert Scale ratings and found that participants self-perceived ConvXAI to **help them to better understand why their writings were given the corresponding reviews** (ConvXAI 4.07 ± 1.18 vs. SelectXAI 3.69 ± 1.37, $p = 0.036$, Figure 3A). . They also felt that **ConvXAI helped them more in improving their writing** ($4{\pm}0.91$ vs. $3.53{\pm}0.77$, $p =$

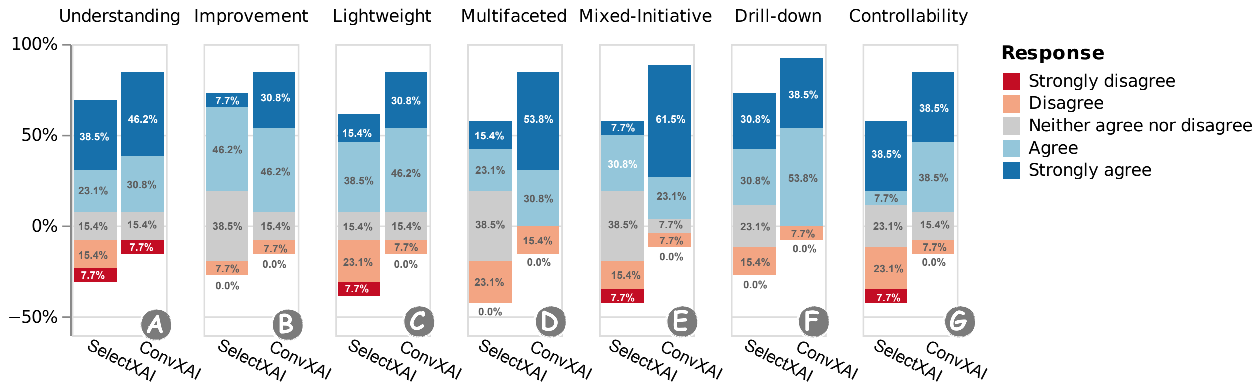

Figure 3: Analyses on users' self-ratings on their experiences playing with ConvXAI and SelectXAI in Task One. They self-rated ConvXAI to be better on all dimensions, and most significantly on the usefulness of mix-initiative and multifaceted functionality.

0.019, Figure 3B). The helpfulness are likely because participants can more effectively find answers to their diverse questions.

Besides their promising self-reflection, 3 out of 13 participants actually edited and iterated their abstracts in ConvXAI. They all successfully addressed the AI-raised issue (*i.e.,* the corresponding suggestion disappeared when they re-evaluated the edited version). However, the other 10 participants showed low incentive to revise the published abstracts. Through interviews, we also summarize some challenges they faced in interacting with the current ConvXAI. Through the study observations and free-form question interviews with users, we obtained that 9 out of 13 participants prefer to use ConvXAI than SelectXAI system for improving their scientific writing. We conjecture that this might primarily result from ConvXAI's ability to answer user questions more *sufficiently*, *efficiently*, and *diversely*. More specifically, the benefit comes from three dimensions: first, ConvXAI reduces users' cognitive load digesting the available information; second, ConvXAI enables users to pinpoint the XAI questions efficiently. third, ConvXAI provides sufficient AI explanations crafted for user need.

For Task Two, we evaluate participants' scientific writing performance quantitatively in terms of *productivity* and *writing performance* (*i.e.,* how many changes have been made and whether the improved writing outputs are scored better). Akin to Task1, we also qualitatively assess participants' *perceived usefulness* with 5 points likert scale from the post-survey.

We can observe that, by comparing with *Original* scores, **both ConvXAI and SelectXAI are useful for humans to improve their auto-metric writing performance**, including the "Grammarly", "Model Quality", and "Model Structure" scores. Furthermore, ConvXAI specifically outperforms SelectXAI on Grammarly and writing quality metrics, indicating that **ConvXAI can potentially help users to write better grammar-based and style-based sentences** in scientific abstracts than SelectXAI. On the other hand, the human editor's evaluation shows inconsistent results, where **ConvXAI and SelectXAI can both improve the writing Structure** evaluations, but not in the Quality metric. To probe the inconsistency between human and auto-metric evaluations, we further compute the Pearson correlation between the model scores and the human ratings and find that both quality and structure are negatively correlated or not correlated (quality: -0.0311 and structure: -0.1150), showing that there is a misalignment between humans and models.

Therefore, we posit that both universal XAI systems, including ConvXAI and SelectXAI, are useful to improve human writing performance under auto-metric evaluations. Particularly, ConvXAI can outperform SelectXAI in terms of grammar and style-based writing quality. Besides, as the human is not aligned with model evaluations based on Pearson correlations, the improvement failed in the human quality metric. This negative finding actually provides valuable insights into the importance of aligning the human judgment and model objective in AI tasks, so that users can use the systems to effectively reach both improvement goals. In the post-survey, we also ask users to rate their perception of system usefulness in terms of assisting their abstract writing. We particularly measured the users' perceived usefulness on "Overall Writing", "Writing Structure", and "Writing Quality" improvement. We can see participants perceived ConvXAI to be 1 (out of 5) point higher than SelectXAI in terms of use on all writing aspects.

REFERENCES

[1] Yossi Adi, Einat Kermany, Yonatan Belinkov, Ofer Lavi, and Yoav Goldberg. 2017. Fine-grained Analysis of Sentence Embeddings Using Auxiliary Prediction Tasks. In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings.*

[2] Dan Jurafsky. 2000. *Speech & language processing.* Pearson Education India.

[3] Himabindu Lakkaraju, Dylan Slack, Yuxin Chen, Chenhao Tan, and Sameer Singh. 2022. Rethinking Explainability as a Dialogue: A Practitioner's Perspective. *arXiv preprint arXiv:2202.01875* (2022).

[4] Q Vera Liao, Daniel Gruen, and Sarah Miller. 2020. Questioning the AI: informing design practices for explainable AI user experiences. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems.* 1–15.

[5] Forough Poursabzi-Sangdeh, Daniel G Goldstein, Jake M Hofman, Jennifer Wortman Vaughan, and Hanna Wallach. 2018. Manipulating and measuring model interpretability. (2018).

[6] Hua Shen and Ting-Hao Huang. 2020. How Useful Are the Machine-Generated Interpretations to General Users? A Human Evaluation on Guessing the Incorrectly Predicted Labels. In *Proceedings of the AAAI Conference on Human Computation and Crowdsourcing*, Vol. 8. 168–172.

[7] Hua Shen and Ting-Hao'Kenneth' Huang. 2021. Explaining the Road Not Taken. *ACM CHI 2022 Workshop on Human-Centered Explainable AI* (2021).

[8] Hua Shen and Tongshuang Wu. 2023. Parachute: Evaluating interactive human-lm co-writing systems. *arXiv preprint arXiv:2303.06333* (2023).

[9] Hua Shen, Tongshuang Wu, Wenbo Guo, and Ting-Hao Huang. 2022. Are Shortest Rationales the Best Explanations for Human Understanding?. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics.* Association for Computational Linguistics.