

Measuring User Understanding in Dialogue-Based xAI Systems

Dimitry Mindlin^{a,*}, Amelie Sophie Robrecht^b, Michael Morasch^a and Philipp Cimiano^a

^aBielefeld University, Faculty of Technology, Semantic Computing

^bBielefeld University, Faculty of Technology, Social Cognitive Systems

ORCID (Dimitry Mindlin): <https://orcid.org/0000-0003-1367-4365>, ORCID (Amelie Sophie Robrecht): <https://orcid.org/0000-0001-5622-8248>, ORCID (Philipp Cimiano): <https://orcid.org/0000-0002-4771-441X>

Abstract. The field of eXplainable Artificial Intelligence (xAI) is increasingly recognizing the need to personalize and interactively adapt the explanation to reflect better users' explanation needs. While dialogue-based approaches to xAI have been proposed recently, the state-of-the-art in xAI is still characterized by what we call one-shot, non-personalized, and one-way explanations. In contrast, dialogue-based systems that can adapt explanations through interaction with a user promise to be superior to GUI-based or dashboard explanations as they offer a more intuitive way of requesting information. In general, while interactive xAI systems are often evaluated in terms of user satisfaction, there are limited studies that assess user's objective model understanding. This is in particular the case for dialogue-based xAI approaches. In this paper, we close this gap by carrying out controlled experiments within a dialogue framework in which we measure the understanding of users in three phases by asking them to simulate the predictions of the model they are learning about. By this, we can quantify the level of understanding of how the model works, comparing the state before and after the interaction. We further analyze the data to reveal patterns of how the interaction between groups with high vs. low understanding differs. Overall, our work thus contributes to our understanding of the effectiveness of xAI approaches.

1 Introduction

The field of xAI focuses on developing methods to render predictions or decisions by machine-learned models comprehensible for users, organizations, developers, and other stakeholders. [14]. Methods range from those that make the general inner workings of machine learning models transparent, such as the meaning of activated neurons in neural networks [40], to generating explanations of AI behavior for stakeholders using visualizations, knowledge extraction or example-based explanations [1].

Most approaches in xAI provide one-shot, non-personalized, and one-way explanations in the sense that they deliver a single explanation that is not adapted to the needs of a particular user and that is generated by the explaining system without any possibility for the user to modify the explanation, request further details, ask for an elaboration, etc. More recently, though, it has been recognized that a single explanation can not meet the needs of all users [37] and that users should be involved in a more participatory way in the explanation process, in particular in the decision of what exactly should be

explained and how [34]. This is especially important for lay users, who are affected by AI decisions and should receive accessible, understandable explanations without requiring advanced technical or domain knowledge.

One approach that allows users to directly shape the process of generating explanations is dialogue-based xAI systems. Dialogue-based xAI methods promote a dynamic, two-way interaction that allows for explanations to be personalized and contextualized, fostering a co-constructive process where explanations evolve through ongoing user interaction and feedback. This approach aligns more closely with natural human conversational norms and facilitates a deeper, more intuitive understanding of AI systems [36, 23]. By adapting to individual cognitive needs and feedback, dialogue-based methods enhance trust and effectively bridge the gap between complex AI functionalities and user expectations [16].

Studies in human-agent interactions, however, present mixed outcomes regarding interactive explanations. Some evidence suggests that while interactive explanations can improve objective understanding, they may also decrease efficiency and user satisfaction due to increased time demands [33]. Other research has shown that adaptive explanations do not consistently outperform non-adaptive ones in enhancing the users' understanding [3]. These findings underscore the complexities and varying results across different settings, emphasizing the importance of targeted studies to validate the practical benefits of dialogue-based xAI approaches.

Within conversational xAI, there is a significant lack of studies that quantify users' objective understanding of AI models, with a prevailing focus on collecting subjective feedback [24]. Our research aims to address this gap by using a "simulation task" [10] method to measure objective user understanding and the effectiveness of interactive explanatory approaches. In xAI research, this method involves participants simulating the behavior of the model on unseen examples, ideally complemented by confidence ratings and explanations for their predictions, providing deeper insights into their understanding [17].

We conducted a study on Prolific to investigate whether a dialogue-based setting yields a higher model understanding than a static setting. Our study design incorporates three phases to measure the user's model understanding after receiving xAI explanations under different conditions. First, in the initial test phase, users undertake the prediction task without exposure to model predictions or explanations, establishing a measure of their intuition ($U_{intuition}$).

* Corresponding Author. Email: dimitry.mindlin@uni-bielefeld.de.

Next, in the learning phase, participants interact with model predictions and explanations either through a static report (as a baseline condition) or an interactive setting where they can see explanations by asking questions. This phase aims to enhance their model understanding over time, monitored as U_{time} . The final phase assesses their model understanding (U_{model}). While U_{model} can be seen as a measure of deep enabledness [8], U_{time} measures how well participants understand and apply the learning of the explanations immediately after receiving them and align their task intuition ($U_{intuition}$) to understanding the model. This structured approach seeks to provide empirical evidence on whether advanced explanatory modalities significantly enhance the user’s model understanding, thus contributing to the optimal design and implementation of conversational xAI systems.

Our paper offers the following contributions:

1. We develop a controlled, dialogue-based experimental framework that allows users to interact with an explaining system over multiple turns by asking pre-defined questions selected from a GUI.
2. This framework was used to conduct experiments with 200 lay users on a simplified income prediction task, assessing their understanding of model decisions through a simulation task.
3. We compare the dialogue-based setting to a static setting, showing highly significant results that the dialogue-based approach yields a higher model understanding compared to a static explanation (Mann-Whitney-U-Test; $p < 0.004$).
4. We analyze and identify common question patterns that users ask to understand model decisions, comparing those who show the highest versus lowest model understanding.

2 Related Work

We explore three lines of work that are relevant to our paper. First, we examine approaches that extend beyond one-shot explanations. Next, we review related work in the field of dialogue-based xAI approaches. Finally, we address studies related to quantifying the objective understanding of a model by the explainees, that is the receivers of an explanation.

2.1 Beyond Single Explanations

Several studies have demonstrated the effectiveness of providing multiple explanations rather than a single explanation.

For example, Baniecki et al. [4] found that providing multiple sequential explanations improved domain specialists’ understanding of incorrect model predictions compared to a single explanation. However, the order of these explanations was fixed, not accommodating individual preferences in selecting explanations.

Furthermore, Arora et al. [2] extended sentiment classification explanations by letting users alter sentences and see model confidence changes, demonstrating that global cues with feature attributions aid meaningful edits. While up to two explanations are shown, the authors do not allow participants to ask for explanations but provide explanations in a one-shot manner.

Research in Human-Robot Interaction (HRI) has revealed interesting insights into explanations as a communicative process. Based on strategies observed in human-human explanation settings, Robrecht and Kopp [32] introduced an adaptive explainer to explain a board game. Axelsson and Skantzé [3] have used adaptive explanation strategies to explain art in a gallery, and Buschmeier [7] has introduced an agent that can do cooperative time scheduling with the user.

Drawing on these findings, it is evident that conversational xAI systems are highly advantageous, facilitating personalized interactions and offering multiple complementary explanations tailored to individual needs and preferences.

2.2 Conversational xAI Systems

Research in conversational xAI so far has focused on tailoring explanations by collecting user feedback. Sokol and Flach [36] pioneered this with the “Glass-Box”, gathering feedback at a conference to identify key features for personalized explanations, which improve transparency in machine learning models [37]. Kuzba and Biecek [19] developed a conversational model to collect queries from the data science and R communities. While these studies motivate the importance of incorporating user needs to refine explanation systems in xAI, their approach lacked an objective assessment of the user’s model understanding, a gap our research addresses.

In the realm of conversational interfaces, both Malandri et al. [22] and Slack et al. [35] introduce innovations aimed at enhancing user interaction, such as improved intent understanding and dialogue management. However, they do not rigorously assess objective improvements in user understanding, instead relying on qualitative feedback. Although their assessments show a user preference for chat interfaces over dashboards, they fall short of objectively evaluating model understanding, highlighting a gap in demonstrating actual improvements in user comprehension within conversational xAI. Similarly, Feldhus et al. [13]’s application of conversational xAI to NLP model understanding also relies on qualitative evaluations, emphasizing the need for more objective assessment methods in the field.

The reviewed studies show advancements in conversational xAI systems with increased personalization and new capabilities. However, the absence of comparable objective evidence on user understanding points to potential risks of improvements in the wrong directions that could distract from increasing understanding and favor aspects such as the likability of the systems.

2.3 Measuring the Impact of Multiple xAI Explanations on Model Understanding

The previous section highlighted advancements in dialog systems within xAI, yet underscored the shortfall in rigorous evaluations of user understanding. While there are many studies comparing the influence of one-shot explanations on user understanding (see Wang and Yin [38] for an overview), there is only a limited amount dedicated to measuring the impact of providing multiple explanations and different interface types.

Cheng et al. [9] tested whether an interactive interface that allowed to access model predictions by customized changes would increase model understanding. Understanding was assessed through crafted questions that tested different aspects of understanding like asking about the influence of attribute changes on the model’s prediction or a simulation task. However, since the authors did not account for users’ intuition, they cannot conclusively attribute the observed understanding improvements to the explanations provided.

Hase and Bansal [15] isolate the effect of explanations by including a pre-assessment of the understanding before providing any explanations. Given this, they estimate the change in understanding by measuring the accuracy of simulation tasks and comparing different explanation approaches as well as a composite method that combines multiple explanations. Interestingly, they do not show clear trends that all explanations help to increase understanding. In particular, the



(a) Dialogue Explanations

(b) Testing Step

Figure 1: Experiment UI. a) shows a teaching step in the learning phase of the interactive condition. After a prediction is selected on the right, the chatbot and Question panel pop up to see the true model prediction and engage in explanations. b) shows a testing instance in the learning step, where the modifications are indicated.

composite explanation condition did not yield a significantly higher understanding. Lastly, the explanations were fixed and neither interactive nor selective—two elements we aim to explore in our research.

Baniecki et al. [4] assess model understanding by a different proxy task that involves selecting whether the model’s predictions were correct and which of the consecutive explanations was most helpful. These consecutive explanations are similar to a dialog, as they are interrelated and provide a drill-down approach. Again, since the order of these explanations was fixed, the authors did not evaluate the impact of customizing the order and quantity of explanations on understanding.

We identified a single study dedicated to objectively measuring the impact of an explanation interface [9], but they do not isolate the interface’s effect on understanding. Other studies that provided multiple explanations did so unidirectionally, limiting the comparability to a dialogue-based setting. We aim to fill this gap and propose an experiment setup to measure model understanding in a conversational xAI setting where we isolate the effect of the interface.

3 Methods

This section provides an overview of the experiment framework that we propose as well as the methods we employ to run it. The code and data can be found in our GitHub repository¹.

3.1 User Interface and Dialogue System

The dialogue system is presented as a web application, depicted in Figure 1. It features an interaction paradigm where users see the attributes of an instance on the upper left, complete with detailed descriptions (available upon hovering on the ?). Users interact by selecting questions on the right; responses appear in the central chatbot window, which displays the conversation history (see 1a). In the

static condition, all explanations are presented within a report, categorized by topics, and organized in the same order as the questions outlined in the interactive condition, rather than in a dialogue window or questions panel. The interface where users are expected to make a prediction is shown in Figure 1b. It indicates modifications applied to the learned instance during the learning phase, though this is not shown in the initial and final testing phases.

In developing our dialogue-based xAI system, we chose not to include Natural Language Understanding (NLU) for two primary reasons. First, using a chatbot with predefined questions guarantees clear, error-free interactions by directly mapping queries to responses, though it limits dialogue flexibility. Second, implementing NLU like in TalkToModel [35] would necessitate extensive resources to process around 40,000 phrases for new datasets. This decision simplifies our system and sharpens our focus on evaluating user interaction and the practicality of our experimental design.

3.2 Experiment Framework

Participants begin the experiment by reviewing and agreeing to the study conditions, followed by providing demographic information and describing their familiarity with machine learning and AI. Afterward, they were shown descriptions and animations of the different tasks for their study condition. The structure of the experiment, depicted in Figure 2, consists of three distinct phases: the *initial test phase*, the *learning phase*, and the *final test phase*. Each phase is designed with a specific objective, collectively contributing to the assessment of the participant’s understanding.

The initial phase assesses participants’ intuition. In this phase, participants make predictions on instances using solely their pre-existing knowledge, without access to model predictions or explanations. This yields a score, $U_{intuition}$, calculated on the basis of ground truth, that allows to capture a baseline understanding before having received information about the model via explanations.

During the learning phase, participants partake in teaching and

¹ <https://github.com/dimitrymindlin/Measuring-User-Understanding-in-Dialogue-based-xAI-Systems>

testing steps. In the teaching step, they first predict outcomes for an instance to enhance engagement through cognitive forcing, a mechanism to delay the display of explanations [6]. This prediction is made before they are shown the model’s prediction and explanations. The conditions differ in the way the explanations are presented:

- **static condition:** Provides an expanded report that lists all the available explanations
- **interactive condition:** interface where users can select the question they would like to ask and see the sequence of questions and answers displayed as a chat window (see Figure 1a)

The testing step involves predicting outcomes for a modified instance with highlighted changes (See 1b), producing a score termed “understanding over time”, U_{time} , based on correct predictions. This cycle promotes reflection and immediate application of new knowledge, enhancing understanding of the model’s decision-making and addressing potential overconfidence, rather than expecting long-term retention [15]. After completing the learning phase, participants are asked to rate their subjective understanding and state their preferred types of questions and explanations. Additionally, we inform them of their performance by disclosing the number of test instances they correctly predicted during the phase.

In the final phase, participants are asked to predict how the model would classify the instances they encountered during the initial test phase, again without the aid of the model’s predictions or explanations. We then use the correctly classified instances to calculate the model understanding, U_{model} . This step assesses participants’ objective model understanding, facilitating a between-subject comparison across the two study conditions. This design allows us to distinctly measure and attribute model understanding to either the interactive or static settings.

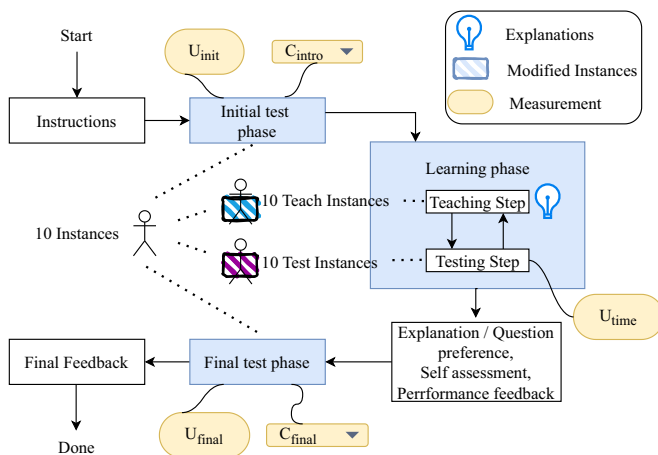


Figure 2: Flow diagram of the study steps: The blue phases require participants to make predictions for displayed instances. The yellow fields indicate assessments of understanding and confidence. While the initial test and final test phases involve the same instances, the instances for the learning phase are modifications of those.

3.2.1 Questions and Explanations

In a literature review, Mindlin et al. [24] highlighted the comprehensive Question Bank by Liao et al. [21] from conversational xAI studies, which we use as the primary framework for selecting our user queries.

We equipped our xAI system with common local explanations and feature statistics techniques, drawing on frameworks like TalkTo-Model [35] as the backend and incorporating further question collections from Liao et al. [21], Nguyen et al. [27], and Malandri et al. [22]. Inspired by the question bank of Nguyen et al. [27], we included *Anchor explanations* [31] to identify necessary conditions to keep the current prediction. Furthermore, we included *Ceteris Paribus explanations* [20] to demonstrate the impact of feature modifications on the model’s decisions, similarly to prior work by Nguyen et al. [27] and Malandri et al. [22]. We align the answers closely to the work of Nguyen et al. [27] but adapt them slightly for our binary classification setting.

We provide the participants with a range of general and feature-specific questions that explain the model’s prediction. Our questions are tailored for a general audience and are influenced by a recent study that gathered information through a user survey by Malandri et al. [22]. The final questions and their corresponding answer methods, presented in Figure 1, were evaluated for clarity and comprehensiveness in our prestudy, as detailed in Section 4.1. While the interactive group was shown explanations based on clicking on the questions, the static group saw all explanations at the same time listed via a static report.

3.3 Measuring Users Model Understanding

According to Kulesza et al. [18], understanding is assessed through two mental models: functional and structural. Functional understanding enables the operation of a system, while structural understanding provides deeper insights into its workings. Similarly, Buschmeier et al. [8] distinguish between enabledness (similar to functional understanding) and comprehension (deeper knowledge of system mechanics). They also introduce a spectrum from shallow to deep for both comprehension and enabledness.

In this study, we define objective understanding as functional understanding, which is generally easier to measure than deep structural knowledge. We focus on assessing and measuring how well users can predict a model’s behavior after being exposed to the model prediction and explanations. This is often measured through prediction tasks, where users predict outputs for different instances, reflecting their grasp of critical attributes and system operations [39].

It has been shown that simulatability (prediction tasks) and decision-making in AI systems are influenced by factors such as the users’ background knowledge and the accuracy of model predictions and explanations [25]. Additionally, the complexity of the domain and task should be appropriately challenging [15, 33]. To manage these variables and ensure a controlled evaluation, several design decisions were made: Choosing a domain with intuitive features reduces the cognitive load associated with understanding the instances, allowing more time to focus on the explanations. To consider the participants’ intuition, we introduce an initial test phase that assesses task performance prior to learning. Next, we focus on explaining the model’s decision boundaries without differentiating between correct and incorrect predictions. Recognizing the difficulty in pre-assessing the complexity of prediction tasks, we apply the one-parametric Item Response Theory (IRT) [12] to assign performance scores to the participants. This theory helps adjust for varying difficulties across prediction instances post-experiment to more accurately measure the model understanding. Finally, in contrast to previous studies that separate the training and testing phases [15], our approach in part interweaves these phases by following each training instance with a testing instance, before the final test phase. This

Table 1: Question-Answer Mapping, Ordered as Presented in Experiment

Question	Method	Answer
Which attributes have the biggest impact on the model’s prediction for the current person	Lime [29]	“Here are the 3 most important attributes for the current prediction: Investment Outcome is the most important attribute...”
Which attributes have the smallest impact on the model’s prediction for the current person	Lime [29]	“Work Life Balance is the least important attribute...”
What is the strength of each attribute for the current model prediction?	Lime [29]	feature contributions plot
To switch the model’s prediction, which attributes would need to be different?	DICE [26]	“Here are possible scenarios that would change the prediction to under 50k: 1. Changing Occupation to Specialized and ...”
What is the scope of change permitted to still get the same prediction?	Anchors [30]	“If you keep these conditions: ... the prediction will stay the same.”
What happens if the value of <i>Investment Outcome</i> is changed, keeping all the other attributes the same?	Ceteris Paribus [20]	“Changing Investment Outcome to <i>Major Gain (above 5k\$)</i> will switch the prediction to over 50k.”
How are the different values of <i>Age</i> distributed in the dataset?	feature statistics	“Age ranges from 17 to 90 with a mean of 38,42.” or barplot of categories

strategy aims to improve information retention by ensuring the application of the learned information, addressing the challenge of participants not retaining information across distinct phases.

Instance Selection. In simulation tasks, our preliminary study (Section 4.1) indicates that test instances should not be markedly different from the training ones to maintain task relevance. However, overly similar instances might simplify the task excessively. For the main instances that are selected for the initial and final test, we use a balanced set of 50% low and high-income people with equally distributed values of the most important feature. For the teaching and testing instances, we use a balanced approach by modifying the features of the previously selected instances. The test instance is a further modification of the teaching instance, resulting in a closely related instance, ensuring it likely remains within the same decision region. This method, similar to the approach described in Hase and Bansal [15], involves randomly modifying 1 to 3 features of the training instances to create teaching and testing instances. Our approach considers the true feature value distribution and maintains a balanced ratio between modifications that change the model prediction and those that do not.

4 Experiments




Figure 3: Statistical analysis of difference in intuition and model understanding across conditions. No significant difference in intuition and highly significant difference in model understanding.

4.1 Pre-study

In preparation for the main study, a pre-study with 80 computer science bachelor students was conducted. Feedback from this pre-study

informed refinements to our experimental design, detailed in Section 3.2. Three changes in the study design were made based on the pre-study results. The dataset domain should be easily understandable by the participants, as the (1) medical feature names in the Pima Indian Diabetes dataset [11] posed challenges for participants. Some students commented that they found themselves searching online for additional information or resorting to guesswork to make predictions. Furthermore, the (2) exclusively numerical nature of the dataset’s features also presented difficulties. Participants struggled to recall and apply information learned during the learning phase to new instances, especially since numerical thresholds varied across instances and were hard to memorize. Therefore, we chose the Adult dataset from the UCI repository [11] for its intuitive, mostly categorical features, simplifying user interaction. Lastly, our initial choice of using (3) randomly selected instances as test cases during the learning phase was ineffective. The high degree of variance between these instances and the teaching instances hindered the application of learned insights. Instead, the main study uses test instances that are slight modifications of the teaching instances to enhance the learning experience.

4.2 Datasets and Model

We utilize the *Adult* dataset from the UCI Machine Learning Repository [11], which comprises 15,682 labeled records of individuals. These records are categorized based on whether the individuals’ annual income exceeds \$50,000 or is below this threshold. We use a similar data processing as Ribeiro et al. [31] but only include a subset of features as seen in the examples in Figure 1a. We trained a Random Forest Classifier using the sci-kit-learn library [28] (ROC-AUC: train: 0.914, test: 0.9).

4.3 Participant Details

We ran the study on Prolific with 200 participants in April 2024, with an average payment of 10.57£ per hour and a median completion time of 26 minutes. The total cost of running the study was 1.206£. Participants were assigned to one of two conditions: interactive or static. To ensure high-quality results, we recruited participants who had an approval rate of at least 99%, resided in the UK or US, and were native English speakers. To ensure engagement, we implemented three mechanisms. First, we incentivized participants by offering additional payments to those who scored in the top 10%

in the final test. Second, we introduced a random variable named “Work Life Balance” into the instances. This variable was consistently presented as the least important, contributing neither to the final prediction nor included in the anchors or counterfactual explanations. It served as an attention or engagement check by filtering out participants who justified their choices in the final test based on this irrelevant attribute. Finally, we limited our analysis to participants who completed the entire study, passed at least one Instructional Manipulation Check (IMC), and did not finish the study in less than two standard deviations below the mean study duration.

After preprocessing the data, we were left with 107 users of which 54 were in the static condition and 53 in the interactive condition. This sample size was deemed sufficient to achieve a power of 0.8 for statistical analysis. The age of participants ranged from 19 to 70 years, with an average age of 40.77 years and a standard deviation of 12.05. Regarding gender, the study included 62 males and 43 females. The self-reported knowledge levels of applications of Machine Learning among respondents are as follows: very low (10), low (44), moderate (43), highly knowledgeable (8), high proficient (2). Given that most respondents rated their knowledge as moderate or lower, where “moderate” reflects a basic understanding of AI applications, our study targets lay users rather than experts or professionals in the field.

5 Results

5.1 Impact of Conditions on Understanding and Confidence

To address variations in prediction difficulties across instances, we employed a one-parameter IRT model to assign performance scores to participants. Item difficulties were calculated separately for intuition and model understanding. To obtain both scores and account for the statistical nature of the IRT, we averaged the individual’s assigned performances across 100 runs. The intuition $U_{intuition}$ across groups was not significantly different (Mann-Whitney U: 1387; p : 0.39), and allows us to attribute the model understanding to the condition. We assessed the effectiveness of the two conditions by comparing the model understanding scores, U_{model} (Figure 3). The boxplots representing U_{model} show that the interactive group had highly significant results and a higher understanding than the static group ($U=998$; $p=0.004$).

We did not observe significant differences in understanding over time, subjective understanding, or confidence in the final predictions.

5.2 Impact of Question Selection on Understanding

We analyze the questioning behavior of participants in the interactive group with high and low U_{model} to grasp which question sequence might lead to higher model understanding.

We define high performers as participants in the top 15% of U_{model} scores and low performers as those in the bottom 15%, balancing the groups to 15 users each. Figure 4 reveals that the “high” group engaged particularly more in the *Ceteris Paribus* explanation and somewhat more in *Feature Ranges*, both of which involve exploring individual features. They also engaged frequently, with *Most Important Features* being their second most utilized question to gain an overview of which features are crucial. Interestingly, both groups rated *Ceteris Paribus* explanations as the least useful in the Question Preferences step.

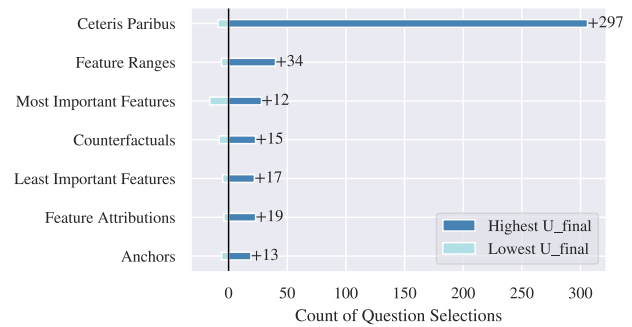


Figure 4: Question selection comparing users with highest vs lowest model understanding.

Exploratory Analysis. Furthermore, we apply process mining techniques to investigate question sequences, treating each dialogue as a distinct process, as suggested by Booshehri et al. [5]. Figure 5b presents the process graphs created using Celonis² software for participants who showed the highest U_{model} , while Figure 5a displays those for participants with the least scores.

Participants with the highest model understanding typically began by asking about the most important features before delving into *Ceteris Paribus* explanations. This approach suggests that understanding the key features provides a foundation before exploring individual feature impacts. In contrast, the least performing participants more often started directly with *Ceteris Paribus* explanations, lacking this initial overview from more general questions.

High performers explored different types of explanations in varied sequences, deepening their understanding of the model. In contrast, participants with lower model understanding used shorter, more direct sequences, which may have led them to miss important insights.

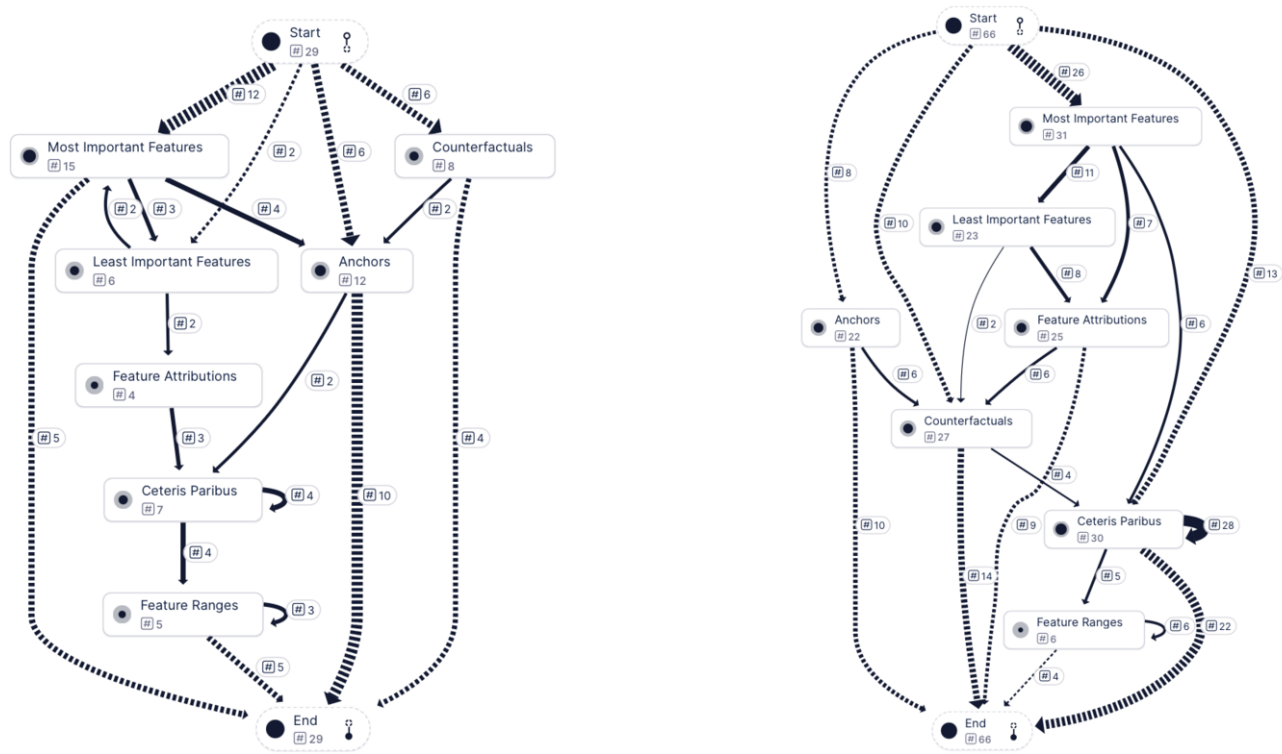
Correlation Analysis. In our final analysis, we used Pearson correlation coefficients to examine the relationships between variables and model understanding. Our primary focus was to determine whether asking specific types of questions correlates with higher model understanding. In addition to analyzing each question separately, we grouped the questions into two categories: feature-specific questions (*Ceteris Paribus* and *Feature Ranges*) and general questions that do not target a specific feature. We observed a strong positive correlation between feature-specific questions and model understanding ($r=0.51$).

6 Discussion

Our proposed experiment setup allows us to elicit user understanding after interacting with a dialogue-based xAI tool. We investigate the relative effectiveness of interactive versus static interactions in helping lay users understand machine learning models. Our results include:

- Our three-phase experiment framework effectively measures objective user understanding through prediction tasks comparing different interface designs.
- Interactive explanations significantly increase objective model understanding, with no substantial change in subjective understanding reported by users.
- Participants who achieved the highest understanding gains explored individual attributes more thoroughly by asking feature-specific questions and, overall, selected more questions.

² <https://www.celonis.com>



(a) Participant with lowest model understanding

(b) Participant with highest model understanding

Figure 5: Question Sequences as process mining graphs for participants with lowest and highest model understanding.

Higher Model Understanding in Conversational Settings. The results show that participants in the interactive condition achieve higher model understanding by selectively engaging with explanations, in contrast to the static condition where all explanations are presented simultaneously. This interactivity likely promotes deeper cognitive processing as participants actively choose which questions to ask, facilitating the co-construction of explanatory dialogues. Additionally, participants with higher understanding tended to ask more questions overall, focusing more on feature-specific questions and continuing the dialogue beyond general questions. This behavior reflects an effective educational strategy, aligning with the recommendation by Baniecki et al. [4] to begin with overviews of feature importances, followed by deeper exploration of key features through *Ceteris Paribus* explanations or *feature distributions*. These findings highlight the value of xAI systems that promote layered exploration of AI decisions, enabling deeper understanding through progressive questioning.

Resulting Follow-up Questions. Our findings prompt several key questions to further enhance dialogue-based xAI and user understanding: Which aspects of interactive explanations—such as the freedom to explore, or depth of engagement—most effectively enhance understanding? Additionally, what is the optimal sequence for delivering explanations to ensure a comprehensive understanding of AI models? Should this sequence be structured and guided, starting from general to specific, or be driven by the explainee’s inquiries?

Limitations. Our study’s limitations include our focus and result interpretation for lay users. While its controlled design, using predefined questions, does not mirror real-world dialogues where users ask varied questions, this yields a baseline dialogue system. Lastly, the

online study format can lead to participants rushing to finish, which we mitigated by excluding fast and inattentive participants.

7 Conclusion

Our study provides a systematic analysis of interactive explanatory systems, aiming to capture statistically significant differences in objective model understanding within a conversational setting. The results confirm the effectiveness of interactive, dialogue-based explanations over static ones in enhancing the understanding of machine learning models among lay users ($p < 0.004$). Our experimental setup, encompassing three phases, effectively assessed users’ objective understanding. We found that users engaging more deeply with feature-specific questions exhibited the highest understanding. In contrast, static explanations yielded lower model understanding. These results underscore the value of interactive explanations in xAI tools, fostering higher model understanding and engagement with explanations.

Acknowledgements

This research is funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation): TRR 318/1 Subproject B01, A dialog-based approach to explaining machine learning models, 2021 – 438445824.

Ethics Statement

All procedures performed in this study were approved by the Ethics Committee of the Bielefeld University. Informed consent was obtained from all individual participants involved in the study.

References

- [1] A. Adadi and M. Berrada. Peeking inside the black-box: A survey on explainable artificial intelligence (XAI). *IEEE Access*, 6:52138–52160, 2018.
- [2] S. Arora, D. Pruthi, N. M. Sadeh, W. W. Cohen, Z. C. Lipton, and G. Neubig. Explain, edit, and understand: Rethinking user study design for evaluating model explanations. In *Thirty-Sixth AAAI Conference on Artificial Intelligence, AAAI 2022, Thirty-Fourth Conference on Innovative Applications of Artificial Intelligence, IAAI 2022, The Twelfth Symposium on Educational Advances in Artificial Intelligence, EAAI 2022 Virtual Event, February 22 - March 1, 2022*, pages 5277–5285. AAAI Press, 2022. doi: 10.1609/aaai.v36i5.20464. URL <https://doi.org/10.1609/aaai.v36i5.20464>.
- [3] N. Axelsson and G. Skantze. Modelling adaptive presentations in human-robot interaction using behaviour trees. In *Proceedings of the 20th Annual SIGdial Meeting on Discourse and Dialogue*, pages 345–352, Stockholm, Sweden, 2019. Association for Computational Linguistics. doi: 10.18653/v1/W19-5940. URL <https://www.aclweb.org/anthology/W19-5940>.
- [4] H. Baniecki, D. Parzych, and P. Biecek. The grammar of interactive explanatory model analysis. *Data Mining and Knowledge Discovery*, pages 1–37, 2023.
- [5] M. Booshehri, H. Buschmeier, M. Alshomary, K. Rohlfing, H. Wachsmuth, and P. Cimiano. Modeling explanations as processes: A new analytical framework accounting for relational and structural patterns in explanatory dialogues. *Journal of Artificial Intelligence*, 2024. doi: 10.5281/zenodo.13151350. Under review. Available as preprint on Zenodo.
- [6] Z. Bućinca, P. Lin, K. Z. Gajos, and E. L. Glassman. Proxy tasks and subjective measures can be misleading in evaluating explainable AI systems. In *IUI*, pages 454–464. ACM, 2020.
- [7] H. Buschmeier. *Attentive Speaking: From listener feedback to interactive adaptation*. PhD thesis, Bielefeld University, 2017.
- [8] H. Buschmeier, H. M. Buhl, F. Kern, A. Grimminger, H. Beierling, J. B. Fisher, A. Groß, I. Horwath, N. Klowitz, S. Lazarov, M. Lenke, V. Lohmer, K. J. Rohlfing, I. Scharlau, A. Singh, L. Terfloth, A. Vollmer, Y. Wang, A. Wilmes, and B. Wrede. Forms of understanding of xai-explanations. *CoRR*, abs/2311.08760, 2023.
- [9] H. F. Cheng, R. Wang, Z. Zhang, F. O’Connell, T. Gray, F. M. Harper, and H. Zhu. Explaining decision-making algorithms through UI: strategies to help non-expert stakeholders. In S. A. Brewster, G. Fitzpatrick, A. L. Cox, and V. Kostakos, editors, *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems, CHI 2019, Glasgow, Scotland, UK, May 04-09, 2019*, page 559. ACM, 2019. doi: 10.1145/3290605.3300789. URL <https://doi.org/10.1145/3290605.3300789>.
- [10] F. Doshi-Velez and B. Kim. Towards a rigorous science of interpretable machine learning. *arXiv preprint arXiv:1702.08608*, 2017.
- [11] D. Dua, C. Graff, et al. Uci machine learning repository. 2017.
- [12] S. E. Embretson and S. P. Reise. *Item response theory*. Psychology Press, 2013.
- [13] N. Feldhus, Q. Wang, T. Anikina, S. Chopra, C. Oguz, and S. Möller. Interrolang: Exploring NLP models and datasets through dialogue-based explanations. In *EMNLP (Findings)*, pages 5399–5421. Association for Computational Linguistics, 2023.
- [14] R. Guidotti, A. Monreale, S. Ruggieri, F. Turini, F. Giannotti, and D. Pedreschi. A survey of methods for explaining black box models. *ACM Comput. Surv.*, 51(5):93:1–93:42, 2019.
- [15] P. Hase and M. Bansal. Evaluating explainable AI: which algorithmic explanations help users predict model behavior? In D. Jurafsky, J. Chai, N. Schluter, and J. R. Tetreault, editors, *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL 2020, Online, July 5-10, 2020*, pages 5540–5552. Association for Computational Linguistics, 2020. doi: 10.18653/v1/2020.acl-main.491. URL <https://doi.org/10.18653/v1/2020.acl-main.491>.
- [16] R. R. Hoffman, G. Klein, and S. T. Mueller. Explaining explanation for “explainable ai”. In *Proceedings of the human factors and ergonomics society annual meeting*, volume 62, pages 197–201. SAGE Publications Sage CA: Los Angeles, CA, 2018.
- [17] R. R. Hoffman, S. T. Mueller, G. Klein, and J. Litman. Measures for explainable AI: explanation goodness, user satisfaction, mental models, curiosity, trust, and human-ai performance. *Frontiers Comput. Sci.*, 5, 2023.
- [18] T. Kulesza, M. M. Burnett, W. Wong, and S. Stumpf. Principles of explanatory debugging to personalize interactive machine learning. In *IUI*, pages 126–137. ACM, 2015.
- [19] M. Kuzba and P. Biecek. What would you ask the machine learning model? identification of user needs for model explanations based on human-model conversations. In *PKDD/ECML Workshops*, volume 1323 of *Communications in Computer and Information Science*, pages 447–459. Springer, 2020.
- [20] M. Kuźba, E. Baranowska, and P. Biecek. pyceterisparibus: explaining machine learning models with ceteris paribus profiles in python. *Journal of Open Source Software*, 4(37):1389, 2019.
- [21] Q. V. Liao, D. Gruen, and S. Miller. Questioning the ai: informing design practices for explainable ai user experiences. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, pages 1–15, 2020.
- [22] L. Malandri, F. Mercorio, M. Mezzanatica, and N. Nobani. Convxai: a system for multimodal interaction with any black-box explainer. *Cogn. Comput.*, 15(2):613–644, 2023. doi: 10.1007/s12559-022-10067-7. URL <https://doi.org/10.1007/s12559-022-10067-7>.
- [23] T. Miller. Explanation in artificial intelligence: Insights from the social sciences. *Artificial intelligence*, 267:1–38, 2019.
- [24] D. Mindlin, F. Beer, L. N. Sieger, S. Heindorf, P. Cimiano, E. Esposito, and A.-C. Ngonga-Ngomo. Beyond one-shot explanations: A systematic literature review of dialogue-based xai approaches. 2024.
- [25] K. Morrison, P. Spitzer, V. Turri, M. Feng, N. Kühl, and A. Perer. The impact of imperfect XAI on human-ai decision-making. *Proc. ACM Hum. Comput. Interact.*, 8(CSCW):1–39, 2024.
- [26] R. K. Mothilal, A. Sharma, and C. Tan. Explaining machine learning classifiers through diverse counterfactual explanations. In *FAT**, pages 607–617. ACM, 2020.
- [27] V. B. Nguyen, J. Schlötterer, and C. Seifert. From black boxes to conversations: Incorporating XAI in a conversational agent. In *xAI (3)*, volume 1903 of *Communications in Computer and Information Science*, pages 71–96. Springer, 2023.
- [28] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, et al. Scikit-learn: Machine learning in python, journal of machine learning re-search, 12. 2011.
- [29] M. T. Ribeiro, S. Singh, and C. Guestrin. “why should I trust you?”: Explaining the predictions of any classifier. In *KDD*, pages 1135–1144. ACM, 2016.
- [30] M. T. Ribeiro, S. Singh, and C. Guestrin. Anchors: High-precision model-agnostic explanations. In *AAAI*, pages 1527–1535. AAAI Press, 2018.
- [31] M. T. Ribeiro, S. Singh, and C. Guestrin. Anchors: High-precision model-agnostic explanations. In *AAAI*, pages 1527–1535. AAAI Press, 2018.
- [32] A. S. Robrecht and S. Kopp. SNAPE: A sequential non-stationary decision process model for adaptive explanation generation. In *ICAART (1)*, pages 48–58. SCITEPRESS, 2023.
- [33] A. S. Robrecht, M. Rothgänger, and S. Kopp. A study on the benefits and drawbacks of adaptivity in ai-generated explanations. 2023.
- [34] K. J. Rohlfing, P. Cimiano, I. Scharlau, T. Matzner, H. M. Buhl, H. Buschmeier, E. Esposito, A. Grimminger, B. Hammer, R. Hüb-Umbach, I. Horwath, E. Hüllermeier, F. Kern, S. Kopp, K. Thommes, A. N. Ngomo, C. Schulte, H. Wachsmuth, P. Wagner, and B. Wrede. Explanation as a social practice: Toward a conceptual framework for the social design of AI systems. *IEEE Trans. Cogn. Dev. Syst.*, 13(3):717–728, 2021. doi: 10.1109/TCDS.2020.3044366. URL <https://doi.org/10.1109/TCDS.2020.3044366>.
- [35] D. Slack, S. Krishna, H. Lakkaraju, and S. Singh. Explaining machine learning models with interactive natural language conversations using talktomodel. *Nature Machine Intelligence*, pages 1–11, 2023.
- [36] K. Sokol and P. A. Flach. Glass-box: Explaining AI decisions with counterfactual statements through conversation with a voice-enabled virtual assistant. In J. Lang, editor, *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI 2018, July 13-19, 2018, Stockholm, Sweden*, pages 5868–5870. ijcai.org, 2018. doi: 10.24963/ijcai.2018/865. URL <https://doi.org/10.24963/ijcai.2018/865>.
- [37] K. Sokol and P. A. Flach. One explanation does not fit all: The promise of interactive explanations for machine learning transparency. *CoRR*, abs/2001.09734, 2020.
- [38] X. Wang and M. Yin. Effects of explanations in ai-assisted decision making: principles and comparisons. *ACM Transactions on Interactive Intelligent Systems*, 12(4):1–36, 2022.
- [39] D. S. Weld and G. Bansal. Intelligible artificial intelligence. *CoRR*, abs/1803.04263, 2018. URL <http://arxiv.org/abs/1803.04263>.
- [40] F. Xu, H. Uszkoreit, Y. Du, W. Fan, D. Zhao, and J. Zhu. Explainable AI: A Brief Survey on History, Research Areas, Approaches and Challenges. In *Natural Language Processing and Chinese Computing*, volume 11839 of *LNCIS Sublibrary*, pages 563–574. Springer International Publishing, 2019. doi: 10.1007/978-3-030-32236-6_51.