ConEX: A Context-Aware Framework for Enhancing Explanation Systems

Yasmeen Khaled and Nourhan Ehab
German University in Cairo, Cairo, Egypt

Keywords: Context, Context-Sensitive Explainability, User-Centered Explainability.

Abstract: Recent advances in Artificial Intelligence (AI) have led to the widespread adoption of intricate AI models, raising concerns about their opaque decision-making. Explainable AI (XAI) is crucial for improving transparency and trust. However, current XAI approaches often prioritize AI experts, neglecting broader stakeholder requirements. This paper introduces a comprehensive context taxonomy and *ConEX*, an adaptable framework for context-sensitive explanations. ConEX includes explicit problem-solving knowledge and contextual insights, allowing tailored explanations for specific contexts. We apply the framework to personalize movie recommendations by aligning explanations with user profiles. Additionally, we present an empirical user study highlighting diverse preferences for contextualization depth in explanations, highlighting the importance of catering to these preferences to foster trust and satisfaction in AI systems.

1 INTRODUCTION

There is no doubt that the field of Artificial Intelligence (AI) has witnessed remarkable advancements in recent decades, resulting in the widespread deployment of complex AI models across various domains. However, concerns about the black-box decision-making processes of these models have escalated. This has sparked a growing interest in Explainable AI (XAI), essential for enhancing trust and transparency within AI systems. Regulatory reforms, including the General Data Protection Regulation (GDPR) in Europe and initiatives like DARPA's Explainable AI research program in the USA, have further accelerated this interest (Gunning and Aha, 2019).

Explainability seeks to make AI outcomes understandable to users (Schneider and Handali, 2019). Unfortunately, current eXplainable AI (XAI) approaches often prioritize the needs of AI experts, neglecting a broader spectrum of stakeholders (Srinivasan and Chander, 2021). Different stakeholders have diverse expectations for explanation complexity and presentation formats. Recent research explores question-driven designs in XAI systems to better cater to users' requirements (Liao et al., 2021). The idea of delivering explanations in a conversational manner through dialogue interfaces has also emerged as a social process (Malandri et al., 2023). Some approaches focus on explainers extracting human-understandable input features (Apicella et al., 2022). However, these approaches tend to overlook the contextual dimension of explanations. An explanation may be intelligible in certain contexts but lack relevance in others. For instance, a system predicting diabetes risk may provide an explanation that a nine-year-old user will develop diabetes due to age, which is considered out-of-context and not aligned with factual knowledge about the disease. There is a critical need to conceptualize explainability considering context, audience, and purpose (Robinson and Nyrup, 2022). Moreover, existing approaches often propose new implementations for explainers instead of utilizing them in a user-centric manner without altering their core functionality. An explainer-agnostic approach to contextualizing explanations, adaptable to any existing explainer, appears to be absent.

The paper argues that constructing a *good* explanation involves two fundamental aspects: (1) explicit knowledge considered during problem-solving for transparency and (2) contextual knowledge surrounding the instance, providing a *frame of reference* for tailoring explanations. Currently, there's no comprehensive taxonomy outlining contextual knowledge elements for explaining instances (Brézillon, 2012).

To address these gaps, we present a novel taxonomy of context, offering guidance on key considerations for creating context-sensitive explanations. We then introduce a general framework, "*ConEX*"

Khaled, Y. and Ehab, N.
ConEX: A Context-Aware Framework for Enhancing Explanation Systems.
DOI: 10.5220/0012385300003636

In *Proceedings of the 16th International Conference on Agents and Artificial Intelligence (ICAART 2024) - Volume 3*, pages 699-706
ISBN: 978-989-758-680-4; ISSN: 2184-433X

as a roadmap for developing context-sensitive explanations using our context taxonomy with any state-of-the-art post-hoc explainer. Finally, we apply the *ConEX* design to build a prototype for context-sensitive explanations in movie recommendations.

This paper is structured as follows: Section 2 presents the proposed context taxonomy, Section 3 introduces the *ConEX* framework, and Section 4 showcases its application in movie recommendations. In Section 5, we present a user study measuring the effect of our prototype on different constructs of trust, followed by concluding remarks in Section 6.

2 CONTEXT TAXONOMY

Context encompasses information characterizing the situation of an entity, including people, places, or objects relevant to user-system interactions (Dey, 2009). To reach context-sensitive explainability, we present a systematic taxonomy, decipted in Figure 1, that dissects various context dimensions serving as a conceptual roadmap for understanding context-sensitive explanations and positing that context comprises *static* and *dynamic* aspects.

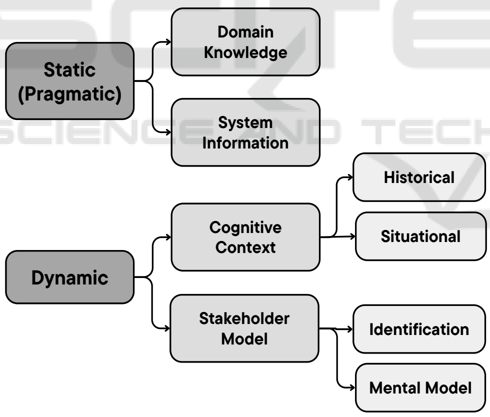

Figure 1: Proposed Context Taxonomy.

2.1 Static Aspects

Our exploration begins with *static context*, the unchanging foundation for the problem at hand, consisting of two key components: *domain knowledge* and *system information*.

Domain knowledge, curated by experts, constitutes a repository of facts within a domain, such as details about movies, actors, directors, and relevant information in a movie recommendation system. In a diagnostic AI system, it includes critical information about risk factors and disease interrelationships. This static context ensures that explanations align with the foundational knowledge of the domain.

System information encompasses definitions of system features and the significance of user interactions. For a movie rating system, it distinguishes the importance of a 5-star rating from a 4-star one and may include user attributes like age groups or postal codes, facilitating user clustering. These insights enable the system to cater to the unique needs of different user groups.

2.2 Dynamic Aspects

The dynamic context involves two entities: the *stakeholder model* and the *cognitive context*.

The *stakeholder model* introduces dynamic personas, acknowledging that stakeholders evolve with each interaction. It involves stakeholder identification, capturing attributes like role and age, and understanding each stakeholder's objectives for personalized explanations. The stakeholder's mental model, shaped by experiences, holds tacit knowledge derived from past interactions and contextual cues.

The *cognitive context* presents a dynamic user worldview, categorized into *situational* and *historical*. Situational cognitive context encompasses emotions, the user's companions, and real-world elements such as date, time, and weather, enriching explanations. Historical cognitive context delves into past interactions, forming the backdrop against which the current interaction unfolds. This taxonomy offers a comprehensive framework for integrating context into explanations, enhancing user-centric AI systems.

3 THE ConEX FRAMEWORK

ConEX, depicted in Figure 2, provides a general framework for generating context-sensitive explanations while decoupling the model from the explanation generation process, thus preserving its predictive accuracy. *ConEX* comprises two distinct modules: (1) a post-hoc explainer and (2) a context model, based on the context taxonomy detailed in Section 2. The outputs of these modules are combined to create context-sensitive explanations. *ConEX* serves as a guideline for enhancing existing systems with context-sensitive explainability or for designing new XAI systems. Below, we delve into the roles of each module.

3.1 Task Model and Explainer

The AI model undergoes initial training for a specific task within a defined domain. In our design, the task model focuses solely on excelling at its designated

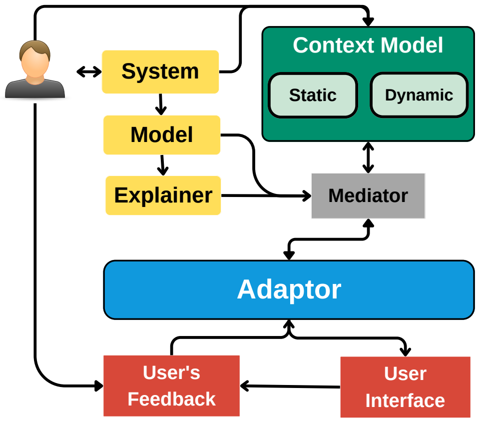

Figure 2: The ConEX Framework.

task without the added responsibility of explaining its predictions. This deliberate separation enables the optimization of the task model for accuracy without compromising explainability.

While state-of-the-art post-hoc explainers excel at approximating task models, adapting explanations to diverse stakeholder needs and contexts can be challenging (Apicella et al., 2022). In our design, the post-hoc explainer is confined to providing fundamental explanations of the task model's results, emphasizing feature importance and decision rules. These explanations contribute to predictions and form the basis for subsequent context-sensitive explanations.

3.2 Context Model

The context model is tasked with supplying all relevant contextual information during ongoing interactions, utilizing the context taxonomy outlined in Section 2. It should provide pertinent domain knowledge components, consider stakeholders' expectations based on their role and mental models, and account for the complexity of explanation presentation and content tailored to situational and historical contexts. In essence, the context model enhances the understanding of the current interaction, regardless of the model's inner workings. It's crucial to note that the context model doesn't generate explanations; instead, it serves as a knowledge reservoir.

3.3 Mediator

The mediator functions as a gateway between the explainer, the context model, and the adaptor. It processes the output from the post-hoc explainer, queries the context model to retrieve contextual information about the instance of interest, and shapes the results to the format expected by the adaptor.

3.4 Adaptor

Contextual knowledge plays a crucial role in filtering and determining the relevant information to consider during a given interaction. Therefore, we contend that excluding irrelevant data from an explanation or supplementing it with external information should not be perceived as misleading. In different contexts, the comprehensive disclosure of the entire decision-making process may prove unnecessary. That is the role of the adaptor module.

The adaptor module is crucial in crafting the final explanation for the user. It integrates the primitive explanation from the explainer with relevant static and dynamic contexts through a two-step process: (1) Pragmatic fitting and (2) Dynamic fitting. In the first step, the primitive explanation aligns with domain knowledge. In the second step, the resulting explanation achieves a balance between complexity and comprehensibility, tailored to the user's identity, mental model, and situational context, facilitating informed decision-making.

Additionally, the Adaptor verifies the prediction itself. If the primitive explanation does not align with domain knowledge, the adaptor may conclude that the prediction should be excluded, ensuring the explanation does not mislead the user.

3.5 User Interface and Feedback

The final step involves presenting the explanation to the user and gathering feedback. The User Interface module delivers explanations in the format suggested by the Adaptor, which may include personalized text templates, images, etc. Users can then provide feedback on various aspects of the explanation. This feedback is invaluable for refining the explanation, aligning it with user preferences and needs, and ensuring it effectively serves its purpose.

Feedback not only improves the explanation but also empowers users by providing a sense of control over the explanation process. This, in turn, enhances user satisfaction and trust in the system.

4 AN APPLICATION: MOVIE RECOMMENDATION

In this section, we unveil our movie recommendation prototype, developed following the design guidelines of *ConEX* to create context-sensitive explanations, as depicted in Figure 3. It is worth noting that ConEX can be used in various other applications, but we chose this application due to data availability.

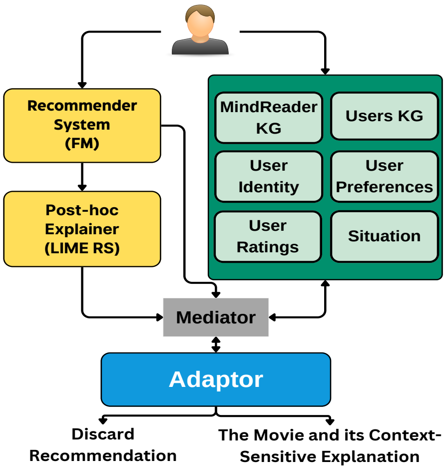

Figure 3: Prototype.

4.1 Task Model and Explainer

The first two blocks are the task model, which is a recommender system responsible for creating the predictions, and the post-hoc explainer used to extract primitive explanations about the recommended instance.

4.1.1 Recommender System

The recommender system in our prototype utilized *Factorization Machines* (FM) (Rendle, 2010). We trained the model using the pyFM library with 50 latent factors and 10 iterations as the stop criteria on user-movie interactions, extended with movie genre features. Other parameters followed default values as per (Nóbrega and Marinho, 2019).

Training data was derived from the well-known MovieLens 1M dataset, comprising user ratings on a 5-star scale for movies (Harper and Konstan, 2015). Ratings were filtered by interaction frequency, considering users with at least 200 interactions. The resulting data was chronologically split (based on rating timestamps) into 70% training and 30% testing. Relevant movies were those rated 3.5 or above. Accuracy results for the FM model were as follows: *Precision@10*: 0.57, *Recall@10*: 0.10, and *RMSE*: 1.21.

4.1.2 Post-Hoc Explainer

LIME for Recommender Systems (LIME-RS) (Nóbrega and Marinho, 2019) was chosen as the post-hoc explainer. LIME-RS, a local post-hoc explainer, provides feature-based explanations for recommendations, drawing inspiration from the concept of LIME (Ribeiro et al., 2016).

In contrast to LIME, LIME-RS generates samples by fixing the user and sampling movies based on their empirical distribution, rather than perturbing data points. A ridge regression model is then trained on the data, outputting the top-n most important features as explanations, with n set to be 20 (*20 genres*).

To measure the fidelity of the ridge regression model with respect to the recommender, we utilized the Model Fidelity metric (Peake and Wang, 2018), as explained in Equation 1. The model was trained on the top 30 predictions for each user from a list of predictions for all items. The average global fidelity was *0.429*, indicating that it can retrieve 42.9% of items.

$$ModelFidelity = \frac{|Explainable \cap Recommended|}{|Recommended|} \quad (1)$$

4.2 Context Model

As shown in the previous results, there is room for improvement in both the recommender's accuracy and the explainer's fidelity. However, using the context model, the recommendations and the explanations can still be grounded in their correct context. Below we will describe how the context model was built for this prototype based on the taxonomy in Section 2.

In the static context, we utilize knowledge graphs (KGs) hosted on Neo4j to represent domain knowledge and user-system information. The MindReader KG (Brams et al., 2020) serves as our source of movie-related entities, covering movies, actors, directors, genres, subjects, decades, and companies. It consists of 18,133 movie-related entities built using a subset of 9,000 movies from the MovieLens dataset, sufficient for our prototype. For system information, user connections and demographic data are represented in a knowledge graph using the demographic data CSV file from MovieLens 1M. GraphXR is employed to create the knowledge graph.

In the situational context of the dynamic context, three factors—mood, company, and time of day are considered. Mood options include happy, sad, angry, and neutral; company options include partner, family, and alone; time of day is categorized as morning or night. Different combinations of these situational factors were randomized and presented to the system during the testing phase. The historical context of a user includes all their ratings of movies in the training dataset. As MovieLens lacks situational data accompanying the ratings, the historical context is represented as a set of {`user_id`, `movie_id`, `rating`} records without incorporating situational information.

The stakeholder model's identification is implemented using a class Person with two sub-classes: lay users (type: user) and developers (type: developer). Both have system-specific attributes, and lay users have additional user-specific attributes like age. This can be expanded with other stakeholder types.

For the stakeholder mental model, we focus on the mental models of lay users. Users' decision-making is assumed to depend on the genres and actors of a movie. The mental model is represented based on genre and actor preferences derived from users' historical cognitive context. Frequent pattern-growth (Han et al., 2004) (FP-Growth) is employed on the genres and actor lists extracted from the user's top-rated movies to mine frequent itemsets. A subset of the inferred genre and actor preferences of user *1004* is shown in Table 1 and Table 2, respectively.

Table 1: User 1004 Genre Preference.

support	**itemsets**
0.50	(Action)
0.41	(Comedy)
0.30	(Adventure)

Table 2: User 1004 Actor Preference.

support	**itemsets**
0.08	(Harrison Ford)
0.05	(Mary Ellen Trainor)
0.04	(Mel Gibson)

The higher the support the itemset has, the more impact it has on the user's decision. This helps determine which itemset to choose in an explanation to further personalize it. Hypothetical explanation content and presentation shapes (text or images) preferences were randomly assigned to the users' mental models.

4.3 Generating Context-Aware Explanations

We now illustrate how a context-sensitive explanation unfolds for a specific instance. In this case, we're focusing on the recommendation of the movie *Lethal Weapon 4* for a user, specifically *User 1004*, whose preferences are listed in Table 1 and Table 2.

4.3.1 Generating Primitive Explanations

During the recommendation process, the movie "Lethal Weapon 4" is suggested to user "1004." LIME-RS is then used to generate a primitive feature importance explanation. This explanation covers positive and negative attributions for all genres in the training data. As depicted in Figure 4, the top positive genre is Film-Noir, which is not present in the genres of the movie *Lethal Weapon 4*. In fact, in 88.4% of the cases when running LIME-RS on the testing dataset,

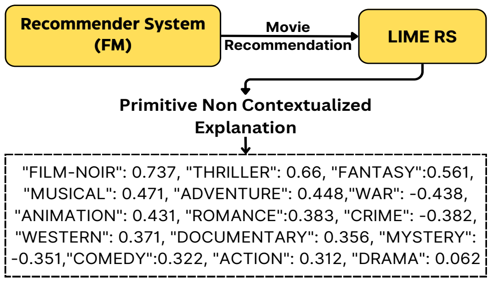

Figure 4: Primitive Explanation.

the highest positively attributing genre did not exist in the genres of the corresponding movie. This discrepancy arises because LIME-RS does not consider context and treats all genres as separate features, including both movie-related and non-movie-related genres (Nóbrega and Marinho, 2019).

Presenting this genre as an explanation to a lay user may not be meaningful, as this particular genre might not even exist in the movie. Therefore, contextualizing the primitive explanation becomes crucial to make it comprehensible to the lay user.

4.3.2 Fetching the Requested Instance Context

When generating context-sensitive explanations, two vital elements demand attention: the movie itself and the user. Through the mediator, we retrieve various context pieces, which encompass:

1. Movie *Lethal Weapon 4* genres: Comedy, Crime, Action, and Drama.
2. Movie *Lethal Weapon 4* actors which are retrieved from MindReader KG: Mel Gibson, Danny Glover, Mary Ellen Trainor, etc.
3. User *1004* identification: Type: lay user, age: 25-34, gender: male, occupation: clerical/admin.
4. User *1004* genre and actor preferences from the user's mental model: presented in Tables 1 and 2.
5. User *1004* preferred explanation content and presentation shape from the user's mental model: Level 2 - Image (this means the user prefers explanations displayed as images).
6. User *1004* randomized current situation: (Sad, With Partner, In the Morning).

4.3.3 Contextualized Explanation Generation

The mediator converts the data types of the primitive explanation and the context data to the ones expected by the adaptor. Subsequently, the adaptor takes the reins in crafting the ultimate context-sensitive explanation, using a three-level approach, with each level building upon the preceding one.

Level 1: Pragmatic Fitting

In the initial contextualization step, termed *Pragmatic Fitting*, the primitive explanation is harmonized with domain knowledge. Specifically, the domain knowledge comprises the actual movie genres acquired through the mediator. The process involves extracting genres from the movie *Lethal Weapon 4* in the primitive explanation (Figure 4). Subsequently, genres with negative attributions are filtered out, leaving only the positive ones. The outcome is a subset of genres, as demonstrated in Figure 5.

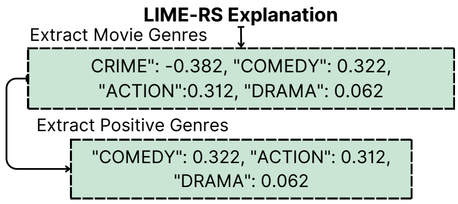

Figure 5: Pragmatic Fitting of Movie *Lethal Weapon 4*.

This step results in a subset of genres that positively contributed to the recommendation, and are consistent with the movie's actual genres. These genres serve as the basis for explanations directly connecting the recommendation to these genres. For instance, one can select the highest attributing genre for explanation, as illustrated in Figure 6.

-Movie Lethal Weapon 4 was recommended to you because it has the genre Comedy.

Figure 6: Pragmatically Fit Explanation.

Level 2: Dynamic Fitting

In the second level of contextualization, the aim is to enhance pragmatic alignment by incorporating additional user-related context, excluding situational details. The explanation intends to guide the user's decision on whether to watch the movie, aligning with their decision-making process. To identify the genre with the highest expected positive impact on the user's decisions, we filter out genres present in the user's preferences (as detailed in Table 1) from the pragmatic fitting results. We then aggregate the attribution of the filtered genres with the support these genres have in the user's preferences as shown in Figure 7. The selected genre with the highest expected positive impact on the user's decisions is *Action*.

Despite Comedy being the genre with the highest attribution in the pragmatically fit result, the genre Action was chosen because it has a higher impact on the user's preferences. Genre Action did have a positive attribution in the pragmatically fit result, thus choosing it over Comedy is not misleading the explanation but rather picking the most relevant piece of

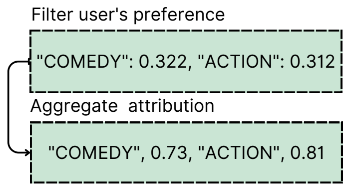

Figure 7: Dynamic Fitting of Movie *Lethal Weapon 4*.

-Movie Lethal Weapon 4 was recommended to you because it has the genre Action and 50% of the movies you highly rated had the same genre.

Figure 8: Dynamically Fit Explanation.

information to the user. The resulting explanation is detailed in Figure 8.

Additional information that is not included in the training of the recommender can be fetched from the context and used to enrich the dynamically fit explanation, as seen in Figure 9. The actors of *Lethal Weapon 4*, fetched in section 4.3.2, are matched against the user's actor preferences in Table 2. Since Mary Ellen Trainor has a higher support value, she was chosen to support the explanation.

Moreover, additional insights about similar users can be inferred from the system information. For instance, the percentage of users in the same age group as user *1004* who highly rated *Lethal Weapon 4* can be derived. This external knowledge augments the explanation, bringing it closer to the user's context.

-Mary Ellen Trainor is starring in this movie and this actor starred in 5% of the movies you highly rated before.

-11% of the users that belong to your age group rated this movie 4 stars or above.

Figure 9: Extra Information to support the Explanation.

Level 3: Incorporating Situation

The third and final level focuses on contextualizing the outcome of Level 2 with situational information. When requesting an explanation, the user inputs their current state, including mood, company, and time of day. For testing purposes, randomized user states were employed.

The goal is to further personalize the interaction, using the resulting genre (and actor if available) from Level 2 as the base information. The situational details are then incorporated to create different paragraphs of explanations, each with different tones and information to match the user's current state. To achieve this, the DeepAI text generator API, backed by a large-scale unsupervised language model, is used to generate paragraphs of text. Thus, the explanation of *Lethal weapon 4* for user *1004* should include genre Action and actress Mary Ellen Trainer.

Using the state of the user *1004* which is *(Sad, With Partner, In the Morning)*, the API outputs a result like the one shown in Figure 10.

> ***If you want an action-packed movie to lift your mood, Lethal Weapon 4 is the perfect choice. With its intense fight scenes and thrilling plot, it's sure to keep you on the edge of your seat. Plus, the talented cast, including the late Mary Ellen Trainer, will keep you engaged from start to finish. Trust me, watching this movie in the morning with your partner is the perfect way to shake off the blues and start your day off right. Don't miss out!***

Figure 10: DeepAI API result Example.

4.3.4 Explanation Reception and Feedback

The final step involves presenting the explanation to users and gathering their feedback. Suppose that User *1004* prefers Level 2 contextualization with image-shaped presentation. Therefore, the information derived from Level 2 can be shown in Figure 11.

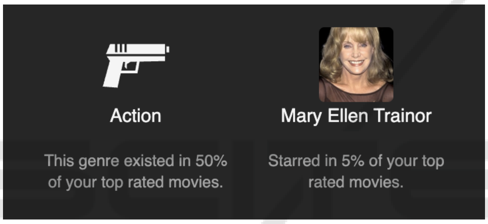

Figure 11: Image presentation for level 2 contextualizations.

Real users can have the option to choose their preferred level of contextualization, content, and presentation format for explanations. The collected feedback is then utilized to refine the explanation style to better match user expectations. Additionally, for developers debugging recommendation instances, the prototype logs primitive explanations and results from each contextualization level. This comprehensive log allows for a detailed analysis, facilitating system optimization. By tailoring explanations to both users and developers, our system provides a versatile and adaptable approach to context-sensitive insights, meeting the specific needs of each user type.

4.3.5 Contextualization Failures

At Level 1 of the contextualization process, scenarios may arise where no genres have positive attributions in the primitive explanation. In such cases, excluding the recommendation is advisable to avoid presenting a misaligned, misleading explanation due to negative attributions. Additionally, if a user prefers Level 2 contextualization but none of the pragmatically aligned genres match their preferences, the recommendation can either be omitted or explained using Level 1. Essentially, the role of the adaptor extends beyond contextualization to act as a truth-checker, filtering out incoherent explanations and enhancing the trustworthiness of the system.

5 EMPIRICAL USER STUDY

In this paper, we hypothesized that diverse stakeholders require tailored explanations to address the same problem. Moreover, we posited that individuals within the same stakeholder category, particularly lay users, prioritize distinct facets of an explanation.

To substantiate our theory, we conducted an empirical user study that investigated the preferences of lay users regarding contextualization levels 1, 2, and 3 with respect to different dimensions of trust. The dimensions of trust were adapted from the work of Berkovsky et al. (Berkovsky et al., 2017), originally designed for assessing trust in various recommender systems. We rephrased these dimensions to pertain to explanations, as seen below:

1. **Competence:** I think the explanation that is most knowledgeable about the movies is...
2. **Integrity:** The explanation that provides the most honest and unbiased reasons is...
3. **Benevolence:** The explanation that reflects my interests in the best way is...
4. **Transparency:** The explanation that helps me understand the recommendation reasons the best is...
5. **Re-Use:** To select my next movie, I would use...
6. **Overall:** The most trustworthy explanation is...

We conducted a study using a survey featuring three movies, each with corresponding explanations labeled as Explanation A (level 1), Explanation B (level 2), and Explanation C (level 3). Participants were introduced to explanations generated by our prototype and asked to respond to six multiple-choice questions assessing the alignment of explanations with the six trust dimensions.

A total of 136 participants completed the survey, with results illustrated in Figure 12. Notably, Explanation A received the least favorability across all trust dimensions, indicating a preference for more personalized explanations. Explanation B scored highest in Integrity, Transparency, and Overall trustworthiness, suggesting a preference for personalized, quantitatively structured explanations. Conversely, Explanation C excelled in Competence, Benevolence, and

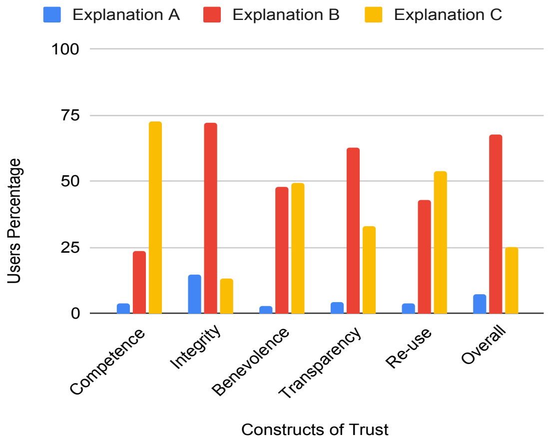

Figure 12: Experiment Results.

Re-use, implying that utilizing the DeepAI API enriched explanations with additional movie-related information, making them more adaptable and likely to be chosen. These findings support our hypothesis that diverse users prefer different explanations, underscoring the need to provide users with control over their choice to accommodate varied informational needs.

6 CONCLUDING REMARKS

In this paper, we have introduced a comprehensive taxonomy of context that finds relevance across diverse domains and systems. To practically realize context-sensitive explanations, we have presented *ConEX*, a general framework founded on our context conceptualization, along with the incorporation of a post-hoc explainer. We presented an application of ConEX that leverages context-sensitive explanations to enhance the personalization of movie recommendations. Additionally, we conducted a user study to demonstrate that context-sensitive explanations enhance user trust and satisfaction empirically. Future work in this domain can include research into automated situation recognition, reducing users' input, and tracking their current state. Moreover, addressing the challenge of temporal changes in user preferences and maintaining the context model's accuracy over time is a promising avenue for future research.

REFERENCES

Apicella, A., Giugliano, S., Isgrò, F., and Prevete, R. (2022). Exploiting auto-encoders and segmentation methods for middle-level explanations of image classification systems. *Knowledge-Based Systems*, 255:109725.

Berkovsky, S., Taib, R., and Conway, D. (2017). How to recommend? user trust factors in movie recommender systems. In *Proceedings of the 22nd International Conference on Intelligent User Interfaces*, IUI '17, page 287–300, New York, NY, USA. Association for Computing Machinery.

Brams, A. H., Jakobsen, A. L., Jendal, T. E., Lissandrini, M., Dolog, P., and Hose, K. (2020). Mindreader: Recommendation over knowledge graph entities with explicit user ratings. In *Proceedings of the 29th ACM International Conference on Information & Knowledge Management*, page 2975–2982, New York, NY, USA. Association for Computing Machinery.

Brézillon, P. (2012). Context in artificial intelligence: I. a survey of the literature. *COMPUTING AND INFORMATICS*, 18(4):321–340.

Dey, A. (2009). Explanations in context-aware systems. pages 84–93.

Gunning, D. and Aha, D. (2019). Darpa's explainable artificial intelligence (xai) program. *AI Magazine*, 40(2):44–58.

Han, J., Pei, J., Yin, Y., and Mao, R. (2004). Mining frequent patterns without candidate generation: A frequent-pattern tree approach. *Data Mining and Knowledge Discovery*, 8(1):53–87.

Harper, F. M. and Konstan, J. A. (2015). The movielens datasets: History and context. *ACM Trans. Interact. Intell. Syst.*, 5(4).

Liao, Q. V., Pribic, M., Han, J., Miller, S., and Sow, D. (2021). Question-driven design process for explainable AI user experiences. *CoRR*, abs/2104.03483.

Malandri, L., Mercorio, F., Mezzanzanica, M., and Nobani, N. (2023). Convxai: a system for multimodal interaction with any black-box explainer. *Cogn. Comput.*, 15(2):613–644.

Nóbrega, C. and Marinho, L. (2019). Towards explaining recommendations through local surrogate models. In *Proceedings of the 34th ACM/SIGAPP Symposium on Applied Computing*, SAC '19, page 1671–1678, New York, NY, USA. Association for Computing Machinery.

Peake, G. and Wang, J. (2018). Explanation mining: Post hoc interpretability of latent factor models for recommendation systems. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, KDD '18, page 2060–2069, New York, NY, USA. Association for Computing Machinery.

Rendle, S. (2010). Factorization machines. In *2010 IEEE International Conference on Data Mining*, pages 995–1000.

Ribeiro, M. T., Singh, S., and Guestrin, C. (2016). "why should i trust you?": Explaining the predictions of any classifier.

Robinson, D. and Nyrup, R. (2022). Explanatory pragmatism: A context-sensitive framework for explainable medical ai. *Ethics and Information Technology*, 24(1).

Schneider, J. and Handali, J. (2019). Personalized explanation in machine learning. *CoRR*, abs/1901.00770.

Srinivasan, R. and Chander, A. (2021). Explanation perspectives from the cognitive sciences—a survey. In *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence*, IJCAI'20.