Ask More, Score Higher: Towards Robust Semantic Enrichment through Interactive Query Population and Richness Scoring

Yuan Tian[1], Yiru Chen[2], Rakesh R. Menon[2], Zifan Liu[2], Ting Cai[3], Fei Wu[2], Anudeep Chimakurthi[2], Prashanthi Ramamurthy[2], Sridevi Aishwariya Ganesan[2], Kun Qian[2], Yunyao Li[2]

[1]Purdue University
[2]Adobe
[3]University of Wisconsin-Madison
tian211@purdue.edu,
{yiruc, rakeshra, zifanl, feiw, achimakurthi, pramamur, sridevi, kunq, yunyaol}@adobe.com,
tingcai@cs.wisc.edu

Abstract

LLM-powered agents are increasingly being deployed for data-related tasks, including data import, data exploration, data visualization, and data analytics. However, their performance heavily depends on the clarity and completeness of data field semantics. Unfortunately, many field descriptions remain ambiguous or incomplete, as much of the essential context (e.g., the meaning of a customized field) originates from users' domain knowledge and is rarely documented publicly. This gap restricts the effectiveness of LLM-based agents in downstream tasks, such as entity linking. To bridge this gap, we introduce a novel **I**nteractive **SE**mantic **E**nrichment system (ISEE). Given a field description, ISEE evaluates its quality using a novel scoring system, efficiently gathers user knowledge, and collaboratively enriches the semantics with users.

Introduction and Related Work

LLM-powered agents are increasingly deployed in enterprise tasks such as entity linking (Shen, Wang, and Han 2015; Kolitsas, Ganea, and Hofmann 2018) and natural language querying of databases (Ning et al. 2023; Tian et al. 2023; Ning et al. 2024; Zhang et al. 2025). While academic environments assume well-documented or publicly available context, enterprise environments pose distinct challenges. Business databases often contain highly customized fields whose meaning is privately owned by domain experts. Consequently, when field descriptions are ambiguous or incomplete, LLMs face challenges in interpreting their meaning, leading to errors in downstream tasks. For example, when a user asks a natural language question, the entity-linking service is expected to accurately map the query to the corresponding mentioned fields. However, given the unclear semantics, LLMs can only guess and make mistakes.

Recently, semantic enrichment methods (Belsky, Sacks, and Brilakis 2016; Xue, Wu, and Lu 2021) are proposed to automatically make field description more meaningful and accurate. While these methods can mitigate this issue, they come with inherent limitations in practical enterprise scenarios. This is because these methods leverage the broad knowledge embedded in general-purpose LLMs, whereas domain-specific expertise remains inaccessible to the underlying LLMs. Without specific context, it is infeasible for LLMs to automatically infer the missing information and effectively enrich these fields.

To address this challenge, it is necessary to keep the domain experts in the loop and solicit domain knowledge. We present ISEE, a novel **I**nteractive **SE**mantic **E**nrichment system. ISEE follows human-in-the-loop system design principles (Rogers 2012; Tian and Zhang 2025; Tian et al. 2024, 2025), in which humans can efficiently understand the system's behavior and provide precise feedback. ISEE automatically evaluates the quality of existing descriptions, efficiently elicits user knowledge, and collaboratively enhances descriptions with domain experts.

System Overview

Figure 1 demonstrates the UI and workflow of ISEE. Starting with field descriptions from a data dictionary (Adobe Experience Data Model[1]), ISEE includes three iterative stages—*Evaluation*, *Clarification (Passive and Active)*, and *Enrichment*.

In the first stage (*Evaluation*), ISEE comprehensively evaluates the quality of the current description across five dimensions, including usability, informativeness, clarity, conciseness, and readability (Figure 1 ①). This enables users to understand the situation and identify areas for improvement.

In the second stage (*Clarification*), once users understand the quality of the current description, they can contribute related information by providing feedbacks. ISEE enables both passive feedback (Figure 1 ②) and active feedback (Figure 1 ③). Regarding passive feedback, ISEE generates specific and easy-to-answer clarification questions. The question may take the form of a multiple-choice question with straightforward options or an open-ended question. These questions are designed to efficiently elicit users' domain knowledge, helping to resolve ambiguity in the current description. In addition to passively answering clarification questions, ISEE also allows users to actively provide natural language (NL) queries or example data related to this field,

[1]https://github.com/adobe/xdm

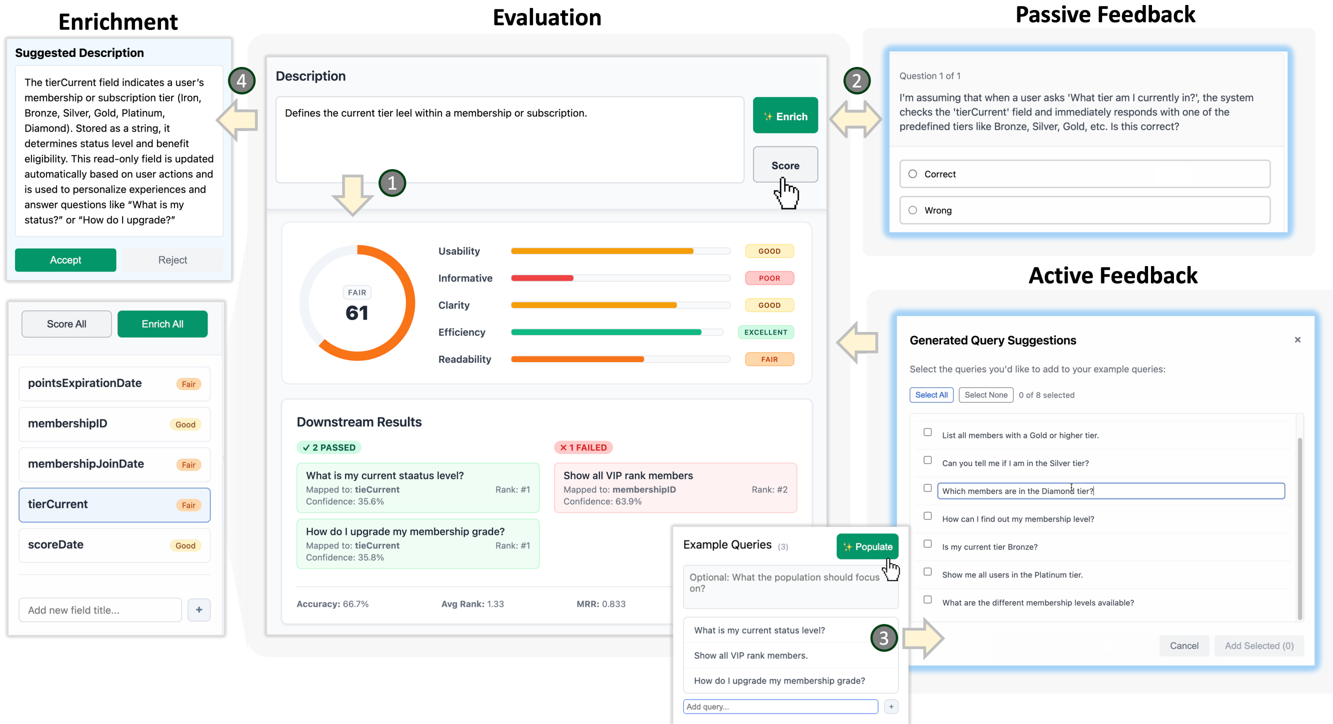

Figure 1: **UI and Workflow of ISEE. (1) Evaluate:** Users click *Score* to obtain a 0–100 quality score based on 5 dimensions and a dashboard of *Downstream Results* that reveal how the quality of the current description. **(2) Passive feedback:** The UI asks concise clarification questions (e.g., scope, attributes, logic); users answer with one click. **(3) Active feedback:** Users add example NL queries, or press *Populate* to review system-populated queries and select the relevant ones. **(4) Enrichment:** The system drafts a *Suggested Description* that the user can Accept/Reject or lightly edit. The loop can be repeated until the score or downstream evaluation meets the user's expectation.

which are important context as they make the system understand the purpose or usage scenario of this field. However, manually creating these examples can be demanding. To reduce effort, ISEE introduces a novel interactive query population feature that automatically generates queries under the given field. Our key insight is that while refining a textual description can be challenging, it is often easier for users to verify related queries. Thus, the system populates candidate queries, and the user only needs to validate or refine them rather than create them from scratch. The process is iterative: previously verified or user-provided queries are incorporated as context, enabling subsequent query population to better capture the intended semantics. To make the query population more controllable, users can further guide this process through an optional NL instruction. For example, a user might specify "focus on time evaluation," so ISEE only populates chronological queries. Without the additional instruction, ISEE generates diverse queries related to this field by default.

Lastly, by incorporating user feedback, ISEE suggests an updated field description (Figure 1 (4)), which users can further refine as needed.

This iterative loop continues until the description achieves a satisfactory level of quality. We use OpenAI's GPT-4o (OpenAI 2024) as the backend model backend model for generating clarification questions and suggested descriptions. We use `text-embedding-3-small` [2] as the embedding model for richness and downstream evaluation.

Conclusion and Future Work

We present ISEE, an interactive semantic enrichment system that enhances enterprise field descriptions through description scoring, query population, and clarification questions. This work demonstrates that interactive enrichment is an effective approach for capturing missing domain knowledge and benefiting enterprise data-related applications.

In future work, we plan to extend the scoring system with additional metrics that capture broader aspects of quality, as well as integrate a wider range of downstream tasks to make the richness scoring more comprehensive. When there are multiple metrics, users can choose which to focus on by adjusting weigts over different downstream tasks and metrics. We also aim to add visualization capabilities, such as embedding-based semantic maps, to provide users with more interpretable feedback. These enhancements will further improve transparency and effectiveness, enabling ISEE to better support diverse enterprise contexts.

[2] https://platform.openai.com/docs/models/text-embedding-3-small

Acknowledgments

We appreciate all the participants in the user study for their valuable comments. This work was supported by Adobe during the first author's internship.

References

Belsky, M.; Sacks, R.; and Brilakis, I. 2016. Semantic enrichment for building information modeling. *Computer-Aided Civil and Infrastructure Engineering*, 31(4): 261–274.

Kolitsas, N.; Ganea, O.-E.; and Hofmann, T. 2018. End-to-End Neural Entity Linking. arXiv:1808.07699.

Ning, Z.; Tian, Y.; Zhang, Z.; Zhang, T.; and Li, T. J.-J. 2024. Insights into Natural Language Database Query Errors: from Attention Misalignment to User Handling Strategies. *ACM Trans. Interact. Intell. Syst.*, 14(4).

Ning, Z.; Zhang, Z.; Sun, T.; Tian, Y.; Zhang, T.; and Li, T. J.-J. 2023. An Empirical Study of Model Errors and User Error Discovery and Repair Strategies in Natural Language Database Queries. In *Proceedings of the 28th International Conference on Intelligent User Interfaces*, IUI '23, 633–649. New York, NY, USA: Association for Computing Machinery. ISBN 9798400701061.

OpenAI. 2024. GPT-4o System Card. arXiv:2410.21276.

Rogers, Y. 2012. *HCI theory: classical, modern, and contemporary*, volume 14. Morgan & Claypool Publishers.

Shen, W.; Wang, J.; and Han, J. 2015. Entity Linking with a Knowledge Base: Issues, Techniques, and Solutions. *IEEE Transactions on Knowledge and Data Engineering*, 27(2): 443–460.

Tian, Y.; Kummerfeld, J. K.; Li, T. J.-J.; and Zhang, T. 2024. SQLucid: Grounding Natural Language Database Queries with Interactive Explanations. arXiv:2409.06178.

Tian, Y.; Lee, D.; Wu, F.; Mai, T.; Qian, K.; Sahai, S.; Zhang, T.; and Li, Y. 2025. Text-to-SQL Domain Adaptation via Human-LLM Collaborative Data Annotation. In *Proceedings of the 30th International Conference on Intelligent User Interfaces*, IUI '25, 1398–1425. New York, NY, USA: Association for Computing Machinery. ISBN 9798400713064.

Tian, Y.; and Zhang, T. 2025. Selective Prompt Anchoring for Code Generation. arXiv:2408.09121.

Tian, Y.; Zhang, Z.; Ning, Z.; Li, T. J.-J.; Kummerfeld, J. K.; and Zhang, T. 2023. Interactive Text-to-SQL Generation via Editable Step-by-Step Explanations. In Bouamor, H.; Pino, J.; and Bali, K., eds., *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, 16149–16166. Singapore: Association for Computational Linguistics.

Xue, F.; Wu, L.; and Lu, W. 2021. Semantic enrichment of building and city information models: A ten-year review. *Advanced Engineering Informatics*, 47: 101245.

Zhang, T.; Qian, K.; Sahai, S.; Tian, Y.; Garg, S.; Sun, H.; and Li, Y. 2025. Evoschema: Towards Text-to-SQL Robustness against Schema Evolution. *Proc. VLDB Endow.*, 18(10): 3655–3668.