

More Than “Means to an End”: Supporting Reasoning with Transparently Designed AI Data Science Processes

Venkatesh Sivaraman
Patrick Vossler
venkatesh.sivaraman@ucsf.edu
patrick.vossler@ucsf.edu
Weill Cancer Hub West
UC San Francisco
San Francisco, CA, USA

Adam Perer
adamperer@cmu.edu
Carnegie Mellon University
Pittsburgh, PA, USA

Julian Hong
Jean Feng
julian.hong@ucsf.edu
jean.feng@ucsf.edu
Weill Cancer Hub West
UC San Francisco
San Francisco, CA, USA

Abstract

Generative artificial intelligence (AI) tools can now help people perform complex data science tasks regardless of their expertise. While these tools have great potential to help more people work with data, their end-to-end approach does not support users in evaluating alternative approaches and reformulating problems, both critical to solving open-ended tasks in high-stakes domains. In this paper, we reflect on two AI data science systems designed for the medical setting and how they function as tools for thought. We find that success in these systems was driven by constructing AI workflows around intentionally-designed *intermediate artifacts*, such as readable query languages, concept definitions, or input-output examples. Despite opaqueness in other parts of the AI process, these intermediates helped users reason about important analytical choices, refine their initial questions, and contribute their unique knowledge. We invite the HCI community to consider when and how intermediate artifacts should be designed to promote effective data science thinking.

ACM Reference Format:

Venkatesh Sivaraman, Patrick Vossler, Adam Perer, Julian Hong, and Jean Feng. 2026. More Than “Means to an End”: Supporting Reasoning with Transparently Designed AI Data Science Processes. In *Proceedings of CHI’26 Workshop on Tools for Thought (TfT Workshop at CHI’26)*. ACM, New York, NY, USA, 5 pages. <https://doi.org/10.1145/nnnnnnn.nnnnnnn>

1 Introduction

Generative AI systems built on large language models (LLMs) have strong potential to accelerate scientific progress. In data science, they can do so by lowering the barrier to accessing sophisticated techniques like data visualization, predictive modeling, and natural language processing [17, 20, 21]. However, AI data science tools suffer similar pitfalls as their counterparts in AI-assisted software development: they may not actually save users time [1], and they can introduce subtle errors that are difficult to spot but can undermine the validity of results [18, 19, 26]. When applying autonomous or agentic AI systems to societally-relevant data science problems, we believe that human reasoning is especially critical to success yet is currently under-supported.

To understand how modern AI data science tools may fail to fulfill their potential, let us consider Art, a (fictitious) clinician researcher who wants to understand why some patients with a certain type of cancer respond poorly to treatment. AI-based tools such as EHR-Agent [24] and even off-the-shelf LLMs hold the possibility to supercharge Art’s data science capabilities, especially given his lack of coding expertise. While previously he would have to find a data science expert at his institution who would be willing to work with him over a period of months, he can now ask the AI system to extract the relevant data from his institution’s health record database and train models automatically. In no more than a few weeks, Art can run a large set of models and find one with strong predictive performance.

However, there is one problem with Art’s new solution: he has no idea whether his results look so good because of a real effect, or because the AI agent misinterpreted his request and solved a subtly different problem! The system wrote hundreds of lines of code to perform its analysis, but since Art has very little data science expertise, he can’t be sure whether the code is correct. Even the short explanation that the system provided gives no indication of a potential misunderstanding that could skew the results. As a well-intentioned researcher, Art doesn’t feel comfortable publishing an analysis that he doesn’t fully understand, so he spends several more weeks consulting data science experts, painstakingly validating the LLM-generated code, and re-running the tool in search of possible mistakes.

Art’s predicament mirrors findings in the emerging literature on AI-assisted development [1, 18, 19]. For example, one study showed that experienced developers took longer to complete tasks with AI coding tools than without them, largely due to time spent verifying outputs [1]. Expert data analysts have begun to develop strategies for prompting generative AI so that the results are easier to verify [2, 5], but these techniques may not be obvious to those with less data expertise. Indeed, a study of scientists using LLMs to program found that their primary verification strategies—running code and inspecting output—failed to catch errors with scientific ramifications, and that some participants chose not to attempt to understand the generated code at all [19].

The most obvious solution to this problem, and one that has been explored from the early days of explainable AI [14] to modern multi-agent systems [3, 12], is to provide greater visibility into how the AI system arrived at its answer. However, as AI systems for data science have become more complex and capable [7, 9, 13, 24], the paths they take to produce outputs have increasingly been



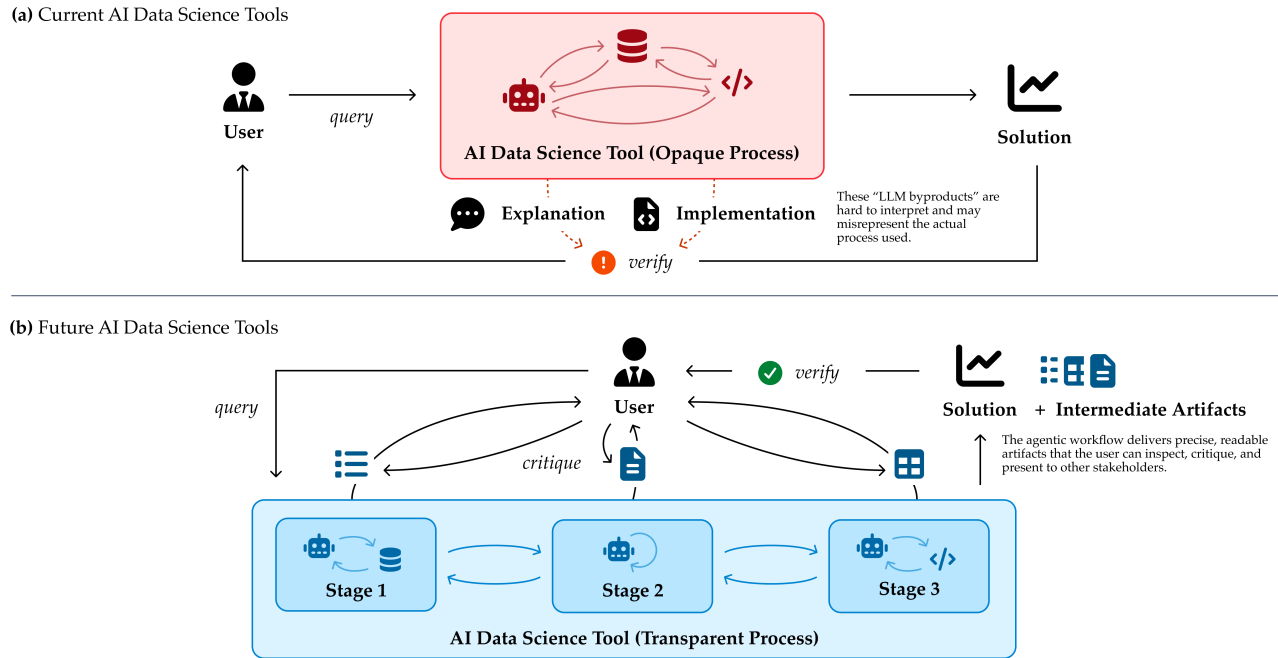


Figure 1: While current AI data science tools use complex, opaque processes that are difficult for users to verify and control (a), we propose that AI data science workflows be explicitly designed around intermediate artifacts that can foster data science reasoning and allow users to steer execution (b). How to choose the stages and intermediate artifacts to structure this transparent process is a central question for future research.

regarded as “means to an end.” Moreover, LLM-generated explanations may merely increase users’ reliance despite not necessarily being accurate reflections of how the answer was produced [10, 19].

We argue that the problem is not an issue of transparency, but a lack of transparency designed for the user. As illustrated in Fig. 1, end-to-end systems often lack well-designed *intermediate artifacts*, which we define as representations of an AI system’s analytical choices that are intentionally created for human review and steering. Unlike computational byproducts, an intermediate artifact is concise, interpretable without requiring technical expertise, and reveals underlying analytical judgments that domain experts can evaluate. This paradigm is not new: Heer [8] described sharing editable structured representations and interactive visualizations between humans and AI to reduce effort in data analysis. We believe this approach is more important than ever in modern data science contexts where AI systems seek to automate large parts of data analysis work, including when the problem formulation is unclear. Therefore, in this paper, we draw from our recent work on AI systems for data science to ask: how does the design of intermediate artifacts turn agentic AI systems from autonomous black boxes into effective tools for thought (TfT)?

2 Case Studies

We ground our discussion of AI data science tools in our experiences building two separate AI-powered systems in the medical domain: HACHI [4] and Tempo [15, 25]. These works are not the only ones

aiming to build more transparent AI workflows for data science [6, 16, 23], but we focus on these due to our personal experiences working on their research teams. We believe learning from the successes and limitations of these tools can help inform an HCI agenda for more autonomous future systems for data science.

2.1 HACHI: Discovering and Annotating Concepts in Clinical Notes

The HACHI system [4] originated as part of a collaboration with a pediatric emergency medicine researcher who was interested in building a decision support model for traumatic brain injury (TBI). Identifying patients who should be tested for TBI is an essential task for emergency medicine providers, but the information that can predict TBI is only available in unstructured clinical notes. The Human+Agent Co-design framework for Healthcare Instruments (HACHI), shown in Fig. 2, was designed to support this data science challenge as well as many other similar tasks. HACHI trains a simple statistical model on *concepts* automatically annotated from clinical text; for example, a simple concept could look like “Has the patient recently had surgery?” The key contribution is that HACHI both identifies concepts and extracts them from notes using an LLM-powered workflow, and it surfaces these concepts to a domain expert user through an interactive interface at intermediate stages of the pipeline. This allows users to inspect the concepts, compare them against their intuition, and reformulate the task.

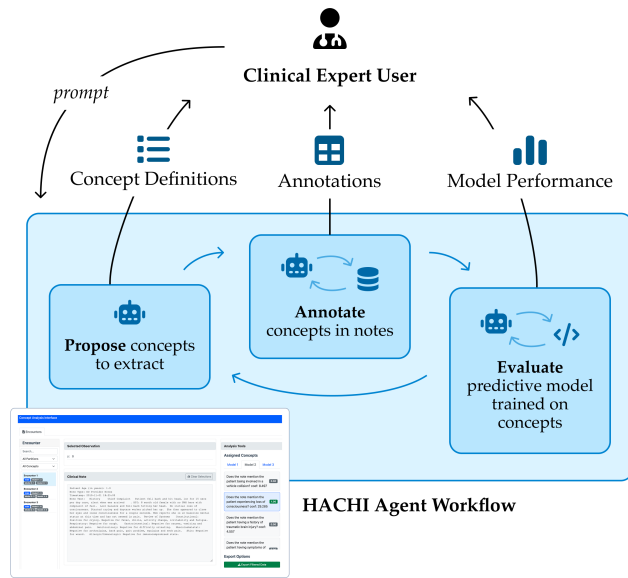


Figure 2: The HACHI workflow trains a predictive model using concepts identified from clinical notes, yielding intermediate artifacts such as the definitions of the discovered concepts, notes and their labels, and the model’s performance.

Using HACHI to develop a predictive model for TBI led to several surprising findings, particularly around clinicians’ involvement in the process. In contrast to conventional workflows for building models on text data, which might obscure important nuances in code, HACHI enabled clinicians to provide feedback to the AI agent that went well beyond the developers’ expectations. For example, early on the team identified a concept (“brain bleed”) whose high predictive weight turned out to reflect data leakage—information available only after the diagnostic test the model was meant to predict—necessitating removal of contaminated cases and a restart of the pipeline. In the next iteration, they noticed that some concepts captured note-writing style rather than patient characteristics (e.g., whether a note *mentions* the Glasgow Coma Scale vs. whether the patient *has* an abnormal score), leading them to constrain how concept definitions were phrased. And when they discovered starkly different model performance across two hospital campuses (AUC 0.93 vs. 0.71), they reweighted the objective function to ensure equitable performance—a decision driven by values around algorithmic fairness that no purely automated system would have made [4]. Over successive rounds of this feedback, model performance and generalizability improved, with each round requiring only 1–2 hours of team review each time [4]. **HACHI’s user interface was essential to making this feedback loop practical:** it gave collaborators full agency in deciding what needed to be reviewed and in what order, and it greatly simplified navigation over many intermediate outputs.

Nevertheless, the HACHI workflow also revealed many open questions on how best to design human-AI systems. The data leakage issue might have been caught earlier if the manual data extraction step had itself been conducted with AI assistance, ideally

supported by intermediate artifacts that made the timing of the notes visible to the whole team. In addition, there were many other opportunities for the human team to engage with and validate intermediate outputs of the LLM pipeline. By optimizing how often and when the human team engages with the LLM pipeline, the final results from HACHI could be even better. For instance, HACHI currently only lets the user update the top-level prompt rather than directly editing concept definitions, but the latter may be important for providing even more precise feedback.

2.2 Tempo: Helping Domain Experts Work with Temporal Event Data

The Tempo project [15, 25] began after the authors worked with several clinical research teams to build prediction models based on electronic health record (EHR) data, but found it prohibitively difficult to engage domain experts at a critical stage: defining and formatting relevant clinical events for modeling. The goal was to design a tool for data scientists that was expressive enough to extract the EHR data they needed, while allowing domain experts to easily understand and critique how the process worked. The final version of this system is an interactive notebook widget that allows users to write queries in a novel query language, called TempoQL, and visualize the results. To further lower the barrier to entry for people with less data science expertise, Tempo includes an agentic AI Assistant that follows the workflow shown in Fig. 3. In response to a natural-language query, the AI Assistant can automatically search the EHR database for relevant concepts, then construct queries using those concepts. Since TempoQL queries are simpler and much more concise than equivalent SQL code, clinical experts can use them to verify and edit the data extraction and aggregation procedure. After running the query in the interface, the user can send a summary of the results back to the AI agent to correct any errors or improve the query.

Case studies with teams of data scientists and domain experts [25] showed that **the TempoQL language served as an effective intermediate for users to reason about complex data extraction workflows.** For example, a product manager working on a web browsing tool was able to rationalize differences between queries that aggregated events at two different time intervals, and they suggested a new aggregation scheme that could combine the best of both alternatives. We also found that off-the-shelf LLM tools, which were used to implement the AI Assistant, were 2.5x more likely to generate correct TempoQL than SQL for the same task, despite only seeing TempoQL’s syntax at inference time [15]. These results suggest that when designing an autonomous agent for data extraction, using a simple, readable language as the intermediate artifact could have benefits for *both* AI accuracy and user understanding.

Whereas HACHI showed that human involvement can help guide an AI system toward solving the right problem, Tempo illustrates how explicitly designing intermediates around human involvement can lead the AI system to produce better outputs. Yet Tempo’s intermediate query language only supports one part of the data science workflow, extraction of temporal event data. It remains an open question how intermediate artifacts could support other tasks like exploratory analysis and predictive modeling. Given the vastness of modern EHR datasets, Tempo also faces the potential issue of

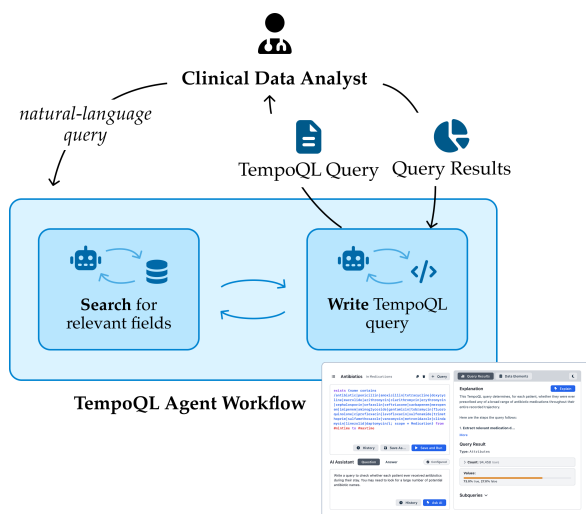


Figure 3: In Tempo, the AI Assistant translates a user’s query into a readable, precise query language to extract temporal data from electronic health records. The TempoQL query and its results serve as intermediate artifacts that help the analyst decide if the data extraction was successful.

“unknown unknowns,” where the AI Assistant might extract only some of the relevant data fields while neglecting others. Future designs could ask the agent to brainstorm alternatives to the solution it presents, helping users identify directions for improvement.

3 Discussion

Amid the proliferation of ever-more-capable AI agents in data science and other domains, it is currently unclear how (and indeed if) it is necessary to bring the human into the loop. After all, wouldn’t a complex black-box workflow that produces the correct answer save the non-data scientist user the most time? Isn’t it sufficient for the AI system developers to validate that the tool uses appropriate processes to derive its results, so end users can focus on the end product? We argue that end-user involvement is most important when the AI needs to function as a TfT: when the problem is unclear or unspecified, when expert intuition should shape the process, and when the credibility of the output depends on the methodology.

The successes we observed in HACHI and TempoQL lead us to call on the HCI community to **design more transparent AI data science tools by intentionally scaffolding workflows with precise and steerable intermediate artifacts**, combining Heer’s framework with the convenience of an end-to-end AI pipeline. For users who have insufficient expertise to think through data science problems independently, well-designed intermediate artifacts can reveal important but subtle considerations that would have required prior experience to foresee. By pointing to ways to reformulate the task, intermediate artifacts can help bridge the “gulf of envisioning,” where users may not know how to specify the task they want to perform until it is at least partially completed [27]. Most importantly, they can empower non-expert users to contribute insights in the areas in which they *do* have expertise. For example, clinicians

can inject their intuition about a disease into the way it is modeled or how concepts are extracted from the notes they write.

The two works discussed in this paper are a starting point for HCI researchers to explore how AI agents’ intermediate artifacts could be more intentionally designed and evaluated. We pose three open questions to guide this research agenda:

When and How Often Should AI Agents Surface Intermediates? In our view, the central question of human-centered AI agent workflow design is at which stages the agent should surface relevant artifacts for critique. While more transparency can always be beneficial, users’ time constraints often make it infeasible for them to review everything that an AI agent does [11], and inexperienced users may not always know which stages of the process are most prone to human-AI misalignment. We suggest that intermediate artifacts should be selected by balancing (a) **how an experienced human might perform the task**, (b) **where choices depend on social values or expert intuition**, and (c) **where different answers might change the problem formulation**. For example, TempoQL queries constitute a recipe for how to retrieve and aggregate the data, which the user can evaluate as they would a sentence in a paper’s methods section. Another potential direction for future work could be to design agentic systems that can themselves decide what user input to solicit.

How Can We Present Intermediate Artifacts? The way intermediate results are presented can directly shape how they support users’ cognition. The works discussed in this paper experimented with three modes of artifact presentation: a **precise, yet readable query language** (TempoQL), **generated natural-language prompts** (HACHI concept extraction prompts), and **input-output pairs** (HACHI concept labels). These interface designs were effective because they did not require data science expertise to interpret and they concisely revealed key choices and misalignments. Future work could envision alternative designs that satisfy these criteria for other stages of the data science pipeline. For instance, a system could draw inspiration from Rewolinski and Yu’s Predictability-Computability-Stability (PCS) framework [22] to encode best practices for data science in an AI workflow, e.g. by presenting variations of an analysis to validate its robustness.

How Can We Evaluate Intermediate Artifacts? Our case studies provided preliminary evidence that engaging domain experts through intermediate artifacts may improve the quality of their data science work. However, quantitatively validating this claim continues to prove challenging [8]. While prior work has examined how novices [18] and scientists [19] use current LLM tools, these studies have either used simple, predefined tasks or evaluated only participants’ self-reported usage. Understanding how AI tools impact people’s analyses is particularly important for high-stakes domains and open-ended problems, yet it is difficult to quantify correctness and robustness for these types of tasks. In some ways, intermediate artifacts could make evaluation easier: for example, more edits to an intermediate output or iterations between stages could represent more effective reasoning support. We invite the community to discuss what rigorous evaluation of AI data science workflows might look like, working towards an effective toolbox for human engagement in the era of AI agents.

Acknowledgments

Thanks to the HACHI and Tempo research teams (Avni Kothari, Ziyong Ma, Anika Vaishampayan, Richard Boyce, and others) for their contributions to the systems we discussed in this work, and to the clinicians and other domain experts whose usage of our tools and candid feedback revealed opportunities for future design. The authors gratefully acknowledge funding support from the Weill Cancer Hub West.

References

- [1] Joel Becker, Nate Rush, Elizabeth Barnes, and David Rein. 2025. Measuring the Impact of Early-2025 AI on Experienced Open-Source Developer Productivity. *arXiv preprint arXiv:2507.09089* (2025). doi:10.48550/arXiv.2507.09089
- [2] Ian Drosos, Advait Sarkar, Xiaotong Xu, Carina Negreanu, Sean Rintel, and Lev Tankelevitch. 2024. "It's like a rubber duck that talks back": Understanding Generative AI-Assisted Data Analysis Workflows through a Participatory Prompting Study. In *Proceedings of the 3rd Annual Meeting of the Symposium on Human-Computer Interaction for Work (CHIWORK '24)*. Association for Computing Machinery, New York, NY, USA, 1–21. doi:10.1145/3663384.3663389
- [3] Will Epperson, Gagan Bansal, Victor C Dibia, Adam Fournay, Jack Gerrits, Erkang (Eric) Zhu, and Saleema Amershi. 2025. Interactive Debugging and Steering of Multi-Agent AI Systems. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems (CHI '25)*. Association for Computing Machinery, New York, NY, USA, 1–15. doi:10.1145/3706598.3713581
- [4] Jean Feng, Avni Kothari, Patrick Vossler, Andrew Bishara, Lucas Zier, Newton Addo, Aaron Kornblith, Yan Shuo Tan, and Chandan Singh. 2026. Human-AI Co-design for Clinical Prediction Models. doi:10.48550/arXiv:2601.09072 arXiv:2601.09072 [cs].
- [5] Ken Gu, Ruoxi Shang, Tim Althoff, Chenglong Wang, and Steven M. Drucker. 2024. How Do Analysts Understand and Verify AI-Assisted Data Analyses?. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems (CHI '24)*. Association for Computing Machinery, New York, NY, USA, 1–22. doi:10.1145/3613904.3642497
- [6] Jiajing Guo, Vikram Mohanty, Jorge H Piazzentin Ono, Hongtao Hao, Liang Guo, and Liu Ren. 2024. Investigating Interaction Modes and User Agency in Human-LLM Collaboration for Domain-Specific Data Analysis. In *Extended Abstracts of the CHI Conference on Human Factors in Computing Systems (CHI EA '24)*. Association for Computing Machinery, New York, NY, USA, 1–9. doi:10.1145/3613905.3651042
- [7] Siyuan Guo, Cheng Deng, Ying Wen, Hechang Chen, Yi Chang, and Jun Wang. 2024. DS-Agent: Automated Data Science by Empowering Large Language Models with Case-Based Reasoning. doi:10.48550/arXiv.2402.17453 arXiv:2402.17453 [cs].
- [8] Jeffrey Heer. 2019. Agency plus automation: Designing artificial intelligence into interactive systems. *Proceedings of the National Academy of Sciences* 116, 6 (Feb. 2019), 1844–1850. doi:10.1073/pnas.1807184115
- [9] Sirui Hong, Yizhang Lin, Bang Liu, Bangbang Liu, Binhao Wu, Ceyao Zhang, Danyang Li, Jiaqi Chen, Jiayi Zhang, Jinlin Wang, Li Zhang, Lingyao Zhang, Min Yang, Mingchen Zhuge, Taicheng Guo, Tuo Zhou, Wei Tao, Robert Tang, Xiangtao Lu, Xiaowu Zheng, Xinbing Liang, Yaying Fei, Yuheng Cheng, Yongxin Ni, Zhibin Gou, Zongze Xu, Yuyu Luo, and Chenglin Wu. 2025. Data Interpreter: An LLM Agent for Data Science. In *Findings of the Association for Computational Linguistics: ACL 2025*, Wanxiang Che, Joyce Nabende, Ekaterina Shutova, and Mohammad Taher Pilehvar (Eds.). Association for Computational Linguistics, Vienna, Austria, 19796–19821. doi:10.18653/v1/2025.findings-acl.1016
- [10] Sunnie S. Y. Kim, Jennifer Wortman Vaughan, Q. Vera Liao, Tania Lombrozo, and Olga Russakovsky. 2025. Fostering Appropriate Reliance on Large Language Models: The Role of Explanations, Sources, and Inconsistencies. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems (CHI '25)*. Association for Computing Machinery, New York, NY, USA, 1–19. doi:10.1145/3706598.3714020
- [11] Avni Kothari, Patrick Vossler, Jean Digitale, Mohammad Forouzannia, Elise Rosenberg, Michele Lee, Jennee Bryant, Melanie Molina, James Marks, Lucas Zier, and Jean Feng. 2025. When the Domain Expert Has No Time and the LLM Developer Has No Clinical Expertise: Real-World Lessons from LLM Co-Design in a Safety-Net Hospital. doi:10.48550/arXiv.2508.08504 arXiv:2508.08504 [cs].
- [12] Harry Li, Gabriel Appleby, Kenneth Alperin, Steven R Gomez, and Ashley Suh. 2025. What is the Agent Doing? Visualizing Agentic AI Querying Workflows. In *1st Workshop on Logo GenAI, Agents, and the Future of VIS*. Vienna, Austria. <https://visxgenai.github.io/subs-2025/2173/2173-doc.pdf>
- [13] Yuan Liao, Jiang Bian, Yuhui Yun, Shuo Wang, Yubo Zhang, Jiaming Chu, Tao Wang, Yuchen Li, Xuhong Li, Shilei Ji, and Haoyi Xiong. 2025. SageCopilot: an LLM-empowered Autonomous Agent for Data Science as a Service. *IEEE Transactions on Services Computing* (2025), 1–18. doi:10.1109/TSC.2025.3635384
- [14] Zachary C. Lipton. 2017. The Mythos of Model Interpretability. doi:10.48550/arXiv.1606.03490 arXiv:1606.03490 [cs].
- [15] Ziyong Ma, Richard D. Boyce, Adam Perer, and Venkatesh Sivaraman. 2025. TempoQL: A Readable, Precise, and Portable Query System for Electronic Health Record Data. doi:10.48550/arXiv.2511.09337 arXiv:2511.09337 [cs].
- [16] Angela Mastrianni, Hope Twede, Aleksandra Sarcevic, Jeremiah Wander, Christina Austin-Tse, Scott Saponas, Heidi Rehm, Ashley Mae Conard, and Amanda K. Hall. 2025. AI-Enhanced Sensemaking: Exploring the Design of a Generative AI-Based Assistant to Support Genetic Professionals. *ACM Transactions on Interactive Intelligent Systems* 15, 4 (Dec. 2025), 1–30. doi:10.1145/3756326
- [17] Ludovico Mitchener, Angela Yiu, Benjamin Chang, Mathieu Bourdenx, Tyler Nadolski, Arvis Sulovari, Eric C. Landsness, Daniel L. Barabasi, Siddharth Narayanan, Nicky Evans, Shriya Reddy, Martha Foiani, Aizad Kamal, Leah P. Shriver, Fang Cao, Asmamaw T. Wassie, Jon M. Laurent, Edwin Melville-Green, Mayk Caldas, Albert Bou, Kaleigh F. Roberts, Sadjana Zagorac, Timothy C. Orr, Miranda E. Orr, Kevin J. Zwezdaryk, Ali E. Ghareeb, Laurie McCoy, Bruna Gomes, Euan A. Ashley, Karen E. Duff, Tonio Buonassisi, Tom Rainforth, Randall J. Bateman, Michael Skarlinski, Samuel G. Rodrigues, Michaela M. Hinks, and Andrew D. White. 2025. Kosmos: An AI Scientist for Autonomous Discovery. arXiv:2511.02824 [cs.AI] <https://arxiv.org/abs/2511.02824>
- [18] Sydney Nguyen, Hannah McLean Babe, Yangtian Zi, Arjun Guha, Carolyn Jane Anderson, and Molly Q Feldman. 2024. How Beginning Programmers and Code LLMs (Mis)read Each Other. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems (CHI '24)*. Association for Computing Machinery, New York, NY, USA, 1–26. doi:10.1145/3613904.3642706
- [19] Gabrielle O'Brien. 2025. How Scientists Use Large Language Models to Program. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems (CHI '25)*. Association for Computing Machinery, New York, NY, USA, 1–16. doi:10.1145/3706598.3713668
- [20] Mizanur Rahman, Amran Bhuiyan, Mohammed Saidul Islam, Md Tahmid Rahman Laskar, Ridwan Mahbub, Ahmed Masry, Shafiq Joty, and Enamul Hoque. 2025. LLM-Based Data Science Agents: A Survey of Capabilities, Challenges, and Future Directions. doi:10.48550/arXiv.2510.04023 arXiv:2510.04023 [cs].
- [21] Shuo Ren, Can Xie, Pu Jian, Zhenjiang Ren, Chunlin Leng, and Jiajun Zhang. 2026. Towards Scientific Intelligence: A Survey of LLM-based Scientific Agents. doi:10.48550/arXiv.2503.24047 arXiv:2503.24047 [cs].
- [22] Zachary T. Rawolinski and Bin Yu. 2025. PCS Workflow for Veridical Data Science in the Age of AI. doi:10.1098/rsta.2024.0605 arXiv:2508.00835 [cs].
- [23] Sarah Schömbis, Yan Zhang, Jorge Goncalves, and Wafa Johal. 2026. From Conversation to Orchestration: HCI Challenges and Opportunities in Interactive Multi-Agent Systems. In *Proceedings of the 13th International Conference on Human-Agent Interaction (HAI '25)*. Association for Computing Machinery, New York, NY, USA, 158–168. doi:10.1145/3765766.3765795
- [24] Wenqi Shi, Ran Xu, Yuchen Zhuang, Yue Yu, Jieyu Zhang, Hang Wu, Yuanda Zhu, Joyce Ho, Carl Yang, and May D. Wang. 2024. EHRAgent: Code Empowers Large Language Models for Few-shot Complex Tabular Reasoning on Electronic Health Records. doi:10.48550/arXiv.2401.07128 arXiv:2401.07128 [cs].
- [25] Venkatesh Sivaraman, Anika Vaishampayan, Xiaotong Li, Brian R Buck, Ziyong Ma, Richard D Boyce, and Adam Perer. 2025. Tempo: Helping Data Scientists and Domain Experts Collaboratively Specify Predictive Modeling Tasks. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems (CHI '25)*. Association for Computing Machinery, New York, NY, USA, 1–18. doi:10.1145/3706598.3713664
- [26] Zhande Song, Jieyu Lu, Yuanqi Du, Botao Yu, Thomas M. Pruyn, Yue Huang, Kehan Guo, Xiuzhe Luo, Yuanhao Qu, Yi Qu, Yinkai Wang, Haorui Wang, Jeff Guo, Jingru Gan, Parshin Shojaei, Di Luo, Andres M. Bran, Gen Li, Qiyuan Zhao, Shao-Xiong Lennon Luo, Yuxuan Zhang, Xiang Zou, Wanru Zhao, Yifan F. Zhang, Wucheng Zhang, Shunan Zheng, Saiyang Zhang, Sartaa Takrim Khan, Mahyar Rajabi-Kochi, Samantha Paradi-Maropakis, Tony Baltoiu, Fengyu Xie, Tianyang Chen, Kexin Huang, Weiliang Luo, Meijing Fang, Xin Yang, Lixue Cheng, Jiajun He, Soha Hassoun, Xiangliang Zhang, Wei Wang, Chandan K. Reddy, Chao Zhang, Zhiling Zheng, Mengdi Wang, Le Cong, Carla P. Gomes, Chang-Yu Hsieh, Aditya Nandy, Philippe Schwaller, Heather J. Kulik, Haojun Jia, Huan Sun, Seyed Mohamad Moosavi, and Chenru Duan. 2025. Evaluating Large Language Models in Scientific Discovery. doi:10.48550/arXiv.2512.15567 arXiv:2512.15567 [cs].
- [27] Hari Subramonyam, Roy Pea, Christopher Pondoc, Maneesh Agrawala, and Colleen Seifert. 2024. Bridging the Gulf of Envisioning: Cognitive Challenges in Prompt Based Interactions with LLMs. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems (CHI '24)*. Association for Computing Machinery, New York, NY, USA, 1–19. doi:10.1145/3613904.3642754