
Towards Multi-faceted Human-centered AI

Sajjadur Rahman, Hannah Kim, Dan Zhang, Estevam Hruschka, Eser Kandogan

Megagon Labs

{sajjadur,hannah,dan_z,estevam,eser}@megagon.ai

Abstract

Human-centered AI workflows involve stakeholders with multiple roles interacting with each other and automated agents to accomplish diverse tasks. In this paper, we call for a holistic view when designing support mechanisms, such as interaction paradigms, interfaces, and systems, for these multifaceted workflows.

1 Introduction

Human-centered AI (HCAI) research focuses on various aspects: analyzing practitioners’ work practices along dimensions such as collaboration [6, 29], explainability [21], and trust [14]; conceptualizing specific workflows to inform design guidelines [11, 13, 16, 17, 24]; and developing supporting tools and frameworks [1, 20, 28]. HCAI workflows are multi-faceted, wherein stakeholders with different roles — *e.g.*, product managers, subject matter experts, and data scientists — perform diverse tasks in various phases using different tools. Multiple roles introduce challenges related to collaboration among stakeholders and interaction with tools. The iterative workflows necessitate fine-grained provenance management [20]. The scope and diversity of domains (*e.g.*, NLP, ML, Vision) introduce challenges related to the modeling of domain semantics. To this end, we propose a holistic design that employs plastic interfaces to enable seamless interactions, integrated workflows for ease of task completion, and a graph-based system for managing workflows and provenance.

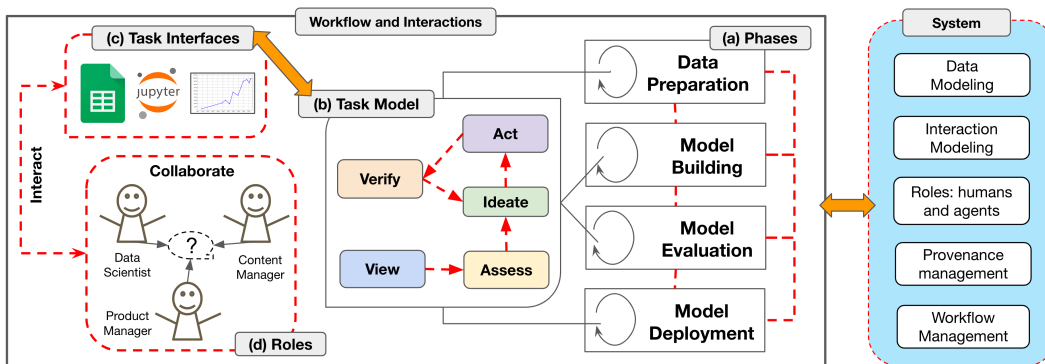


Figure 1: Multifaceted Human-centered AI (M-HCAI) conceptual diagram.

2 M-HCAI Challenges

Figure 1 depicts the conceptual diagram of M-HCAI with various workflow components and system requirements. A workflow is iterative with transitions among *phases* (a). However, a fine-grained characterization, proposed by Rahman et al. [17], reveals the iterative nature of a phase with stakeholders performing diverse *tasks* (c). Stakeholders may interact with different tools and interfaces depending on their task requirement, familiarity, and expertise (d) [20, 27]. Stakeholders with different roles collaborate in various phases within the workflow both synchronously (in-person and virtual meetings) and asynchronously (*e.g.*, Email, Slack, Google Docs) [17, 19, 29].

Such complexity of an MHCAI workflow introduces several challenges. First, *iteration*: iterative workflows often force the users to switch multiple tools tediously, often within a single phase [3, 17, 20]. Second, *interface rigidity*: existing tools lack the plasticity to accommodate diverse roles and expertise as stakeholders ideate and deliberate [5, 14]. Third, *non-standardized collaboration*: as stakeholders collaborate, there may be conflict and misunderstanding [25, 26]. However, there is a lack of standardization around documentation to record these discussions [10, 19, 29]. Finally, *domain diversity*: when we factor in diverse domains within the HCAI setting such as text [4, 20], image, graphs, and tables [15], the lack of formal abstractions and grammar makes it difficult to define and instrument operations in the interfaces systematically.

Another aspect lacking from the HCAI discourse is the management of complicated workflows, which necessitate robust system design. Zhang et al. [29] identified lack of provenance as a contributing factor in obfuscation and loss of knowledge when data science teams share data. Rahman et al. [17] advocate for instrumenting provenance management mechanisms as built-in features of systems. Kandogan et al. [7] highlight how existing systems fail to connect and exploit context across stakeholders and agents due to a lack of initiative in modeling data, people, and interactions.

3 Multifaceted Human-centered AI: A Holistic View

We posit that the workflow- and system-level challenges are interdependent. Capturing the provenance of interactions and decisions within the workflow requires provenance-aware interfaces with built-in logging capabilities. For the underlying system to effectively persist and catalog the provenance information, the interactions defined on the interfaces need grounding on sound principles. Such principled design requires (a) building abstractions to capture domain semantics (*e.g.*, texts contain words, sentences, POS tags, and opinions), (b) defining operators that represent interactions over those abstractions in an interface (similar to the grammar of interaction graphics [22]), and (c) explicit modeling of stakeholder roles. Finally, managing these human-centered iterative workflows requires developing frameworks that enable the integration of tools and enhance interface scalability [2, 18]. In response, we propose the following design considerations for M-HCAI:

An integrated workflow supporting multiple paradigms. Avoiding the context switching overhead during iteration requires integrating different modalities: programming environments and interactive interfaces (*e.g.*, spreadsheets and visualization tools.) Computational notebooks are suitable for such integration where programming can be complemented by introducing interactive widgets [8, 28]. However, notebooks are not suitable for non-programmers involved in the HCAI process. Moreover, these tools neither manage the fine-grained provenance of user interactions nor support versioning. Therefore, research efforts for bridging these gaps are crucial for the holistic M-HCAI system design.

Plastic interfaces as the basis for interaction. The multi-role inclusion challenge with computational notebooks can be addressed by developing custom interfaces that may act as boundary objects for the stakeholders. Boundary objects are artifacts that are “both plastic enough to adapt to local needs and the constraints of the several parties employing them, yet robust enough to maintain a common identity across site” [9, 23]. Recent work introduces cross-platform capabilities to transform interactive widgets in notebooks into web dashboards that enable shared understanding among stakeholders within an organizational workflow [1]. However, these widgets are not provenance-aware, and the corresponding interactions are domain semantics agnostic. Therefore, further augmentation is required to employ these interfaces to operate in concert with the underlying system.

Systems for supporting M-HCAI. Interactions among automated agents and stakeholders of different expertise and roles may occur at any phase. Therefore, the context (where and when a task or interaction occurred), scope (specific data domain and phase), and abstractions of interactions will vary depending on the scenario. To this end, knowledge graphs [12] (KGs) may serve as the core data model for the support systems managing M-HCAI workflows. KGs can explicitly capture diverse stakeholder roles and automated agents, different interaction types, and the corresponding context [7]. Existing work define grammars to capture operations for text [4, 20], and tabular data analysis [15]. Capturing domain-specific interaction contexts such as operations and data types would require defining such abstractions for M-HCAI.

In this position paper, we draw attention to the multi-faceted nature of HCAI while identifying workflow and system-level challenges. We propose a research agenda that requires multi-disciplinary research efforts spanning databases, AI, visualization, and HCI.

References

- [1] A. Bäuerle, Á. A. Cabrera, F. Hohman, M. Maher, D. Koski, X. Suau, T. Barik, and D. Moritz. Symphony: Composing interactive interfaces for machine learning. In *CHI Conference on Human Factors in Computing Systems*, pages 1–14, 2022.
- [2] M. Bendre, T. Wattanawaroon, S. Rahman, K. Mack, Y. Liu, S. Zhu, Y. Lu, P.-J. Yang, X. Zhou, K. C.-C. Chang, et al. Faster, higher, stronger: Redesigning spreadsheets for scale. In *2019 IEEE 35th International Conference on Data Engineering (ICDE)*, pages 1972–1975. IEEE, 2019.
- [3] S. Chattopadhyay, I. Prasad, A. Z. Henley, A. Sarma, and T. Barik. What’s wrong with computational notebooks? pain points, needs, and design opportunities. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, pages 1–12, 2020.
- [4] L. Chiticariu, R. Krishnamurthy, Y. Li, S. Raghavan, F. Reiss, and S. Vaithyanathan. SystemT: An algebraic approach to declarative information extraction. In *Proceedings of the 48th Annual Meeting of the Association for Computational Linguistics*, pages 128–137, Uppsala, Sweden, July 2010. Association for Computational Linguistics. URL <https://aclanthology.org/P10-1014>.
- [5] S. R. Hong, J. Hullman, and E. Bertini. Human factors in model interpretability: Industry practices, challenges, and needs. *Proceedings of the ACM on Human-Computer Interaction*, 4 (CSCW1):1–26, 2020.
- [6] E. Kandogan, A. Balakrishnan, E. M. Haber, and J. S. Pierce. From data to insight: Work practices of analysts in the enterprise. *IEEE Computer Graphics and Applications*, 34(5):42–50, 2014. doi: 10.1109/MCG.2014.62.
- [7] E. Kandogan, M. Roth, I. Terrizzano, E. Seabolt, P. Schwarz, H. Krishnareddy, and A. Agarwal. Context analytics: Vision, architecture, opportunity. In *2018 IEEE 34th International Conference on Data Engineering Workshops (ICDEW)*, pages 1–8. IEEE, 2018.
- [8] M. B. Kery, D. Ren, F. Hohman, D. Moritz, K. Wongsuphasawat, and K. Patel. mage: Fluid moves between code and graphical work in computational notebooks. In *Proceedings of the 33rd Annual ACM Symposium on User Interface Software and Technology*, pages 140–151, 2020.
- [9] S. Leigh Star. This is not a boundary object: Reflections on the origin of a concept. *Science, Technology, & Human Values*, 35(5):601–617, 2010.
- [10] M. Mitchell, S. Wu, A. Zaldivar, P. Barnes, L. Vasserman, B. Hutchinson, E. Spitzer, I. D. Raji, and T. Gebru. Model cards for model reporting. In *Proceedings of the conference on fairness, accountability, and transparency*, pages 220–229, New York, NY, USA, 2019. ACM.
- [11] M. Muller, I. Lange, D. Wang, D. Piorkowski, J. Tsay, Q. V. Liao, C. Dugan, and T. Erickson. How data science workers work with data: Discovery, capture, curation, design, creation. In *Proceedings of the 2019 CHI conference on human factors in computing systems*, pages 1–15, New York, NY, USA, 2019. ACM.
- [12] N. Noy, Y. Gao, A. Jain, A. Narayanan, A. Patterson, and J. Taylor. Industry-scale knowledge graphs: Lessons and challenges. *Commun. ACM*, 62(8):36–43, jul 2019. ISSN 0001-0782. doi: 10.1145/3331166. URL <https://doi.org/10.1145/3331166>.
- [13] S. Passi and S. Jackson. Data vision: Learning to see through algorithmic abstraction. In *Proceedings of the 2017 ACM conference on computer supported cooperative work and social computing*, pages 2436–2447, New York, NY, USA, 2017. ACM.
- [14] S. Passi and S. J. Jackson. Trust in data science: Collaboration, translation, and accountability in corporate data science projects. *Proceedings of the ACM on Human-Computer Interaction*, 2 (CSCW):1–28, 2018.

- [15] D. Petersohn, S. Macke, D. Xin, W. Ma, D. Lee, X. Mo, J. E. Gonzalez, J. M. Hellerstein, A. D. Joseph, and A. Parameswaran. Towards scalable dataframe systems. *Proceedings of the VLDB Endowment*, 13(11).
- [16] K. H. Pine and M. Liboiron. The politics of measurement and action. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*, pages 3147–3156, New York, NY, USA, 2015. ACM.
- [17] S. Rahman and E. Kandogan. Characterizing practices, limitations, and opportunities related to text information extraction workflows: A human-in-the-loop perspective. CHI '22, New York, NY, USA, 2022. Association for Computing Machinery. ISBN 9781450391573. doi: 10.1145/3491102.3502068. URL <https://doi.org/10.1145/3491102.3502068>.
- [18] S. Rahman, K. Mack, M. Bendre, R. Zhang, K. Karahalios, and A. Parameswaran. Benchmarking spreadsheet systems. In *Proceedings of the 2020 ACM SIGMOD International Conference on Management of Data*, pages 1589–1599, 2020.
- [19] S. Rahman, P. Siangliulue, and A. Marcus. Mixtape: Mixed-initiative team action plan creation through semi-structured notes, automatic task generation, and task classification. *Proceedings of the ACM on Human-Computer Interaction*, 4(CSCW2):1–26, 2020.
- [20] S. Rahman, P. Griggs, and Ç. Demiralp. Leam: An interactive system for in-situ visual text analysis. In *CIDR*, pages 1–7, cidrdb.org, 2021. CIDRDB.
- [21] A. Rule, A. Tabard, and J. D. Hollan. Exploration and explanation in computational notebooks. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, pages 1–12, New York, NY, USA, 2018. ACM.
- [22] A. Satyanarayan, D. Moritz, K. Wongsuphasawat, and J. Heer. Vega-lite: A grammar of interactive graphics. *IEEE transactions on visualization and computer graphics*, 23(1):341–350, 2016.
- [23] S. L. Star. The structure of ill-structured solutions: Boundary objects and heterogeneous distributed problem solving. In *Distributed artificial intelligence*, pages 37–54. Elsevier, 1989.
- [24] D. Wang, J. D. Weisz, M. Muller, P. Ram, W. Geyer, C. Dugan, Y. Tausczik, H. Samulowitz, and A. Gray. Human-ai collaboration in data science: Exploring data scientists’ perceptions of automated ai. *Proceedings of the ACM on Human-Computer Interaction*, 3(CSCW):1–24, 2019.
- [25] M. E. Whiting, A. Blaising, C. Barreau, L. Fiuza, N. Marda, M. Valentine, and M. S. Bernstein. Did it have to end this way? understanding the consistency of team fracture. *Proceedings of the ACM on Human-Computer Interaction*, 3(CSCW):1–23, 2019.
- [26] M. E. Whiting, I. Gao, M. Xing, N. J. Diarrassouba, T. Nguyen, and M. S. Bernstein. Parallel worlds: Repeated initializations of the same team to improve team viability. *proceedings of the ACM on Human-Computer Interaction*, 4(CSCW1):1–22, 2020.
- [27] K. Wongsuphasawat, Y. Liu, and J. Heer. Goals, process, and challenges of exploratory data analysis: An interview study. *arXiv preprint arXiv:1911.00568*, 2019.
- [28] Y. Wu, J. M. Hellerstein, and A. Satyanarayan. B2: Bridging code and interactive visualization in computational notebooks. In *Proceedings of the 33rd Annual ACM Symposium on User Interface Software and Technology*, pages 152–165, 2020.
- [29] A. X. Zhang, M. Muller, and D. Wang. How do data science workers collaborate? roles, workflows, and tools. *Proceedings of the ACM on Human-Computer Interaction*, 4(CSCW1): 1–23, 2020.