

AI-Centered Grand Challenges in Visual Analytics for Healthcare: Synthesizing the VAHC 2025 Community Experience

Jürgen Bernard*
University of Zürich
Digital Society Initiative

Gabriela Morgenshtern
University of Zürich
Digital Society Initiative

David Gotz†
University of North Carolina
at Chapel Hill

Renata Raidou‡
TU Wiena.arleo@tue.nl

Robert S Laramée‡
University of Nottingham

Silvia Miksch§
TU Wien

Alessio Arleo||
Eindhoven University of
Technology

ABSTRACT

The intersection of AI, healthcare, and visualization is evolving rapidly, posing challenges that cut across disciplinary boundaries and resist easy resolution. The Visual Analytics in Healthcare workshop (VAHC), co-located every other year at the IEEE VIS conference and the AMIA (American Medical Informatics Association) annual conference, has served as a forum to connect the visualization and medical informatics community since 2010. In 2025, to celebrate the 16th edition, we used the workshop as an opportunity to consolidate the community’s collective experience (and expertise) and identify Grand Challenges where the field should prioritize going forward. We combined thematic coding of the 15 accepted VAHC workshop papers with structured group discussions among more than 40 participants, organized around three major themes: “Technical innovation vs. clinical reality”, “Human-centered and scalable VAHC”, and “From foundations to actionable insights”, followed by post-workshop reflexive analysis. Across all three groups, AI emerged as the most consistently recurring concern. In this paper, we report our AI-centered insights from the VAHC 2025 group activity, contextualize them against the broader literature along five Grand Challenges themes, and distill them into five challenge clusters, each concluded with recommendations for future research directions that cross disciplinary boundaries: (1) *trust and bias*, (2) *data and infrastructure*, (3) *explainability and communication*, (4) *human-AI interaction*, and (5) *model reliability and validation*. We share these challenges and their associated research directions as a starting point for discussion and collaboration across the healthcare, AI, and visualization communities. All supplemental materials are available at <https://osf.io/p79uj>.

Index Terms: Visual Analytics, Healthcare Visualization, Challenges and Opportunities, AI in Healthcare, Human-AI Interaction

1 INTRODUCTION

The rapid integration of AI into healthcare is reshaping how clinicians, patients, and researchers interact with health data. The visualization community is uniquely well-positioned to support this transition, yet faces the challenge of identifying where to focus its efforts. The introduction of these technologies requires that researchers address unprecedented obstacles: those of scalability and data quality, resulting from the explosion in consumer-collected data, of those same consumers’ sudden, unprecedented access to health data about themselves, of clinicians’ skepticism towards AI recommendations in clinical decision-making and treatment plan-

ning, and of a lack of patient literacy, which complicates physician-patient relationships.

This is not the first time that the visualization community has challenged healthcare research. In 2015, Caban and Gotz [13] recognized visualization and visual analytics (VA) as disciplines capable of realizing the potential behind the adoption of Electronic Health Records (EHR) on a large scale, rejecting paper records as the standard for recording and sharing clinical data. This “call to action” was written to provoke a reaction in both the visualization and medical informatics communities, pushing the two towards deeper collaboration on research, towards more open, accessible, and personalized care. The Visual Analytics in Healthcare (VAHC) workshop [1] is a forum where the visualization and medical informatics communities can gather to discuss research intersecting visual data analysis and clinical informatics. Topics of interest to the workshop include applications of VA to clinical care, public health data analysis, AI for healthcare data visualization, and patient- and consumer-related methods. In 2024, a special issue of the *Journal of American Medical Informatics Association* (JAMIA) was published on the topic of digital and personal health, following the results of VAHC 2024 [4].

As with the introduction of EHR, the advent of AI methods in healthcare has sparked a radical paradigm shift, impacting various aspects of research and practice (see, e.g., [33, 39, 37]). These include support in patient diagnosis and clinical decisions [37], automated identification of cancerous tissues from X-rays [33], drug research [39], robot-assisted surgical procedures [37]. Within this “revolution”, with opportunities and challenges on the horizon, VAHC 2025 set out to directly involve the community in a collaborative effort to identify the *Grand Challenges* in VAHC, intended to cross boundaries between visualization/VA and medical informatics communities.

In this paper, we report the AI-centered findings from VAHC 2025, co-located with IEEE VIS. Our process is grounded in over a decade of VAHC workshop proceedings, and includes systematic thematic coding of the 15 accepted VAHC 2025 submissions, structured group discussions with more than 40 workshop participants, and post-workshop reflexive analysis. This process reveals five AI-centered challenge clusters (bias and trust, data and infrastructure, explainability and communication, human-AI interaction, and model reliability and validation) that recurred independently across three differently-framed discussion groups, suggesting they reflect community-wide priorities rather than organizer framing. Across all five, a consistent signal emerges: the most persistent AI challenges in healthcare VA are not technical, but rather calibration problems: between trust and evidence, between explainability and clinical utility, between automation and human agency. We offer these clusters and their associated research directions as a synthesis of current community concerns and a step toward a shared research agenda at the intersection of healthcare, AI, and visualization.

2 METHODS

Longitudinal Literature Grounding. We base our approach on over a decade of VAHC workshop proceedings, ensuring that VAHC 2025 grand challenges activities will be well-grounded. Caban and Gotz [13] presented a structured survey of open problems

*e-mail: {juergen.bernard|gabriela.morgenshtern}@uzh.ch

†e-mail: gotz@unc.edu

‡e-mail: robert.laramee@nottingham.ac.uk

§e-mail: silvia.miksch@tuwien.ac.at

||e-mail: mail@renataraidou.com

||e-mail: a.arleo@tue.nl

in health VA, identifying integration with clinical workflows, scalability, evaluation, and data privacy as persistent gaps, many of which remain unsolved. A subsequent VAHC special issue [4] updated this inventory around six clusters: explainable and trustworthy AI, user-centered design, health communication, EHR accessibility and privacy, equity and inclusion, and interdisciplinary collaboration. External perspectives on evaluation rigor, encoding standards, and the gap between visualization research norms and clinical evidence requirements were additionally drawn from a panel on health visualization challenges at EuroVis 2025 [10]. Together, these sources informed the controlled vocabulary used for submission coding and the three Grand Challenge themes.

Pre-Workshop Submission Coding. We coded all 15 accepted VAHC 2025 papers (seven research papers, four posters and system demonstrations, and four Grand Challenge papers) against a controlled vocabulary of 20 thematic codes, derived inductively from the submissions and anchored in the literature above. [Codes included evaluation methodology, Human-centered AI \(HCAI\), LLM applications, decision-making support, user task characterization, scalable data analysis, wearable data, genomics, reasoning under uncertainty, and visualization fundamentals, among others.](#) This coding provided a transparent, assumption-independent basis for identifying workshop themes and allocating submissions to discussion groups, reducing the risk that theme selection reflected organizer bias rather than the community’s actual research priorities. Coding was performed independently by one VAHC conference organizer (co-author), with all other organizers reviewing assignments; disagreements were resolved through group discussion until consensus was reached.

Theme Derivation and Workshop Execution. Similarity-based clustering of the coded submissions yielded three themes.

- *Theme 1* (Technical Innovation vs. Clinical Reality) clustered around LLM applications, HCAI, and clinical decision-making support, addressing the persistent gap between academic innovation and clinical adoption. The discussion prompt asked: *How can we ensure that the tools we innovate are clinically viable, trusted, and adoptable by end users?*
- *Theme 2* (Human-Centered Scalable VAHC) addressed user task characterization, wearable and genomic data, and personal-to-population scaling. The discussion prompt asked: *How do we design visualization frameworks that serve multiple stakeholders while scaling from individual patient care to population health?*
- *Theme 3* (From Foundations to Actionable Insights) consolidated visualization fundamentals, evaluation methodology, and reasoning under uncertainty. The discussion prompt asked: *What design principles and evaluation standards ensure healthcare visualization tools reliably guide users from exploratory analysis to actionable decisions?*

Formulating each theme as an open prompt rather than a closed definition was a deliberate choice to encourage participants to surface unexpected directions rather than confirm pre-existing assumptions.

The activity portion of the workshop opened with four lightning talks presenting the accepted Grand Challenge papers [44, 16, 6, 43], each framing an open research problem and further operationalizing the tension between technical innovation and clinical adoption (see Table 1). Ultimately, these talks seeded the structured group discussions that followed, in which more than 40 participants engaged in 90-minute sessions per theme, using hand-written notes to identify challenges, sub-themes, and research directions, to lower the barrier to participation. Gotz, Bernard, and Miksch moderated the groups; designated note-takers Arleo, Morgenshtern, and Laramée recorded and summarized outcomes. [Participants spontaneously joined the groups based on their personal interests in the health-related themes. The groups were diverse in healthcare-real-world experience, comprising academic participants at professor, postdoctoral researcher, and PhD student level.](#)

Post-Workshop Analysis and AI Scope Decision. Following the workshop, each moderating pair conducted reflexive thematic analysis of their group’s notes, breaking down discussions into

Table 1: Distribution of AI-related statements in the three workshop participant groups (#) and workshop papers per challenge cluster.

Challenge Cluster		#	VAHC2025 Papers
	Bias and Trust	12	[28, 20, 43]
	Data and Infrastructure	12	[5, 16, 32, 7, 6]
	Explainability and Comm.	12	[28, 8, 29, 38, 44]
	Human-AI Interaction	23	[5, 6, 20, 23, 21, 30]
	Model Reliability and Validation	10	[43, 38, 44]
Total		69	15

overarching topics. The full author team then collaboratively analyzed all three sets of discussion notes, iteratively identifying convergences across themes. Convergence across independently moderated groups strengthened confidence that the identified topics reflect community-wide priorities rather than the perspective of any single discussion group. [This process resulted in 155 statements being identified across the three themes, with 53, 62, and 40 per group, respectively.](#) AI emerged as a recurring theme across all three groups, as well as across the other VAHC discussion formats (submissions and lightning talks). From the full set of VAHC workshop discussion notes from the group activity, we isolated AI-related statements (69/155) and applied a bottom-up inductive coding approach, deriving codes and sub-codes from the data and resulting in five overarching codes, each with three to four sub-categories (see Table 1). [A detailed breakdown of statements and coding is available as supplemental material.](#) Consistent with reflexive thematic analysis practices, members of the research team acted as active interpreters in this process. The coding process involved discussions to reflect on and refine interpretation by the entire team. For each resulting cluster, we examined its presence across workshop submissions, contextualized findings against the broader literature, and distilled opportunities and future directions to seed cross-disciplinary research agendas.

3 AI-CENTERED GRAND CHALLENGES IN VAHC

We present five challenge clusters that emerged from our analysis . We discuss each challenge cluster through workshop findings, broader literature, and concrete opportunities and directions.

 **Bias and Trust** The workshop discussion on bias and trust (12/69) highlights three core challenges for healthcare. First, multiple forms of *bias* (5/12), model, data, and human (e.g., confirmation bias), raise concerns about reliability and decision quality, emphasizing the need for lifecycle-wide mitigation. This concern aligns with work documenting how biases compound across the AI development pipeline, from data acquisition through deployment, with downstream effects on clinical decision quality and healthcare disparities [15, 14]. Second, *misinformation* (3/12) emerges both from public misconceptions about AI (e.g., equating all AI to LLMs) and from external sources, such as forums, exacerbated during crises such as COVID-19 [3]. Third, *trust calibration* (4/12) is critical: both overtrust (even among experts) and distrust must be managed, promoting informed skepticism. [In visualization research, overtrust was empirically shown to increase with task difficulty, and transparency alone does not reliably recalibrate it \[19\], a finding with direct implications for time-pressured clinical settings.](#) Workshop contributions illustrate these tensions in deployed settings. Nipu et al. [28] report on the post-deployment evaluation of an interactive oncological risk estimator, documenting how clinician feedback on a multi-view AI-backed dashboard shapes revisions to risk communication and cohort-level stratification. Hosseini et al. [20] examine how LLMs adapt visualization rhetoric to population health personas, raising open questions about how persona-driven rhetorical choices may reinforce stakeholder priors and affect calibrated interpretation. In the group activity, VAHC systems were seen as potential tools for communicating and, therefore, strengthening trust between patients, clinicians, and other healthcare providers, as well as a way to reduce or raise awareness about potential

biases. Towards this goal, Sacha et al.'s conceptual framework [36] proposes decomposing uncertainty propagation throughout the VA pipeline and process, providing a basis for explicitly communicating uncertainty at key stages and introducing operator- or system-level interventions that counteract perceptual and cognitive biases. Such an approach, instantiated in healthcare contexts [27], can directly support calibrated trust in clinical VA systems.

Opportunities and Directions. Across sub-themes, biases and trust emerge as situated constructs shaped not only by data and model performance, but by how uncertainty, bias, and provenance are communicated in context. This points to the need for standardized frameworks for trust calibration that account for overtrust and undertrust in deployed clinical AI systems, and that guide how these factors should be operationalized in clinical practice.

 **Data and Infrastructure** The workshop discussion topics on data and infrastructure (12/69) are equally prominent in AI-driven healthcare. Large-scale *data collection* (3/12), from patient journaling to continuous sensing, offers significant potential, but creates analytical bottlenecks and data quality problems, [with most data remaining underutilized or not validated according to quality metrics or medical domain knowledge](#). This highlights the need for effective pipelines that transform raw data into clinically meaningful insights. Workshop contributions also underscore the opportunity to better leverage ubiquitous sensing for the benefit of patient care. Assor et al. [5] present *ParcoursVis*, a progressive VA system operating on EHR-derived event sequences at the scale of millions of patients and billions of events, and report an insight-based evaluation protocol developed with hospital practitioners. Rauscher et al. [32] address the analytical challenges of multi-granular physical activity datasets spanning laboratory and real-life wearable measurements, contributing a linked-view application for hypothesis generation in interdisciplinary intervention studies. At the same time, *governance* concerns (5/12) emerge around data trustworthiness, privacy, and access, alongside calls for standards such as FAIR principles [46]. While shared [data access across a care team can aid effective care, access restrictions in the healthcare system exist \(in part\) to protect patients, and care providers must uphold secure, appropriately restricted access](#). Persistent gaps in data stewardship were found during the COVID-19 pandemic across clinical information systems [31], echoing the presence of this challenge externally from the VAHC research community. Finally, *integration* (4/12) remains a key challenge, requiring interoperability across heterogeneous sources and a clearer data infrastructure to enable seamless, reliable use in clinical and patient-centered contexts. In the group activity, VAHC systems were seen as enablers of integration, exploration, and interpretation of heterogeneous data streams, transforming large-scale raw data into actionable insights while supporting transparency, governance, and interoperability across clinician- and patient-centered systems. Data provenance communication strategies in VA can address issues around accountability and liability present when clinicians begin to engage with patient-provided data, providing information on data sources (patient, clinical, sensor-based vs. qualitative). Systematic review evidence indicates that healthcare professionals and researchers generally view patient-generated health data from wearables and apps positively, while concerns regarding data reliability and integration of heterogeneous data sources into existing clinical workflows continue to constrain clinical adoption [11, 45].

Opportunities and Directions. Rather than data scarcity, the primary opportunity lies in addressing fragmentation and governance. Future work focuses on developing end-to-end, interoperable pipelines that not only integrate heterogeneous data sources but also make data quality, provenance, and uncertainty explicit throughout. [Specifically, provenance \[47\] uncertainty \[34\] visualization state-of-the-art reports offer a robust conceptual foundation for transparent provenance tracking that support accountability and actionable clinical decision-making](#).

 **Explainability and Communication** Among the statements centered on explainability and communication (12/69), a key

focus is effective *risk communication* (3/12). This includes how to present uncertainty and establish standardized, evidence-based methods for conveying risk to patients. *Visualization* (3/12) plays a critical role in making data provenance, quality, and insights understandable. At the same time, limitations of transparency and tensions between *explainable AI (XAI)* research and clinical practice emerge (6/12), where explainability is often secondary to validated outcomes. Critiques that post-hoc explanations of black-box models provide limited epistemic guarantees for high-stakes clinical decisions motivate the use of inherently interpretable models where feasible [35]. For example, work such as Ghassemi et al.'s [18] argues that current XAI methods are unlikely to deliver the trust, transparency, and bias-mitigation outcomes frequently attributed to them at the patient level, and [recommends](#) rigorous internal and external validation as a more direct path to clinical utility. Broader reflections emphasize a shift toward human-centric [\(explanatory\) AI](#) [25], rather than model-centric notions of explainability, questioning what “explainable” means in context. This includes comparing AI to human decision-making and highlighting the need for context-sensitive, task-oriented explanations. In VAHC 2025, Balaka et al. [8] present the MoBa GWAS Explorer, a design study investigating onboarding strategies that render genome-wide association study data tractable for both expert and non-expert audiences drawn from a longitudinal cohort study. Owen et al. [29] propose an empathy-centred visualization framework for person-centred dementia care, extending explainability from model reasoning to the communication of indirect end-user experience. In our group activity, VAHC systems appear as answers to making AI decisions interpretable by clearly communicating risk, uncertainty, and data provenance through context-sensitive, task-oriented visual interfaces. A key direction emerging here is the need for systematic, context-aware approaches to communicating data, decisions, and interventions for interpretability. During COVID-19, this need became particularly visible. For instance, Baumgartl et al. [9] systematically characterize user groups, task structures, modeling approaches, and data modalities, and analyze the visual encodings employed to communicate data, decisions, and interventions in support of explainability.

Opportunities and Directions. Closing the gap between technical explainability and clinical meaningfulness, where outcomes and usability consistently precede formal model interpretability in practitioners' priorities [12], is critical. This calls for visualization approaches that communicate uncertainty, provenance, and bias in context, evaluated for their measurable impact on decision quality and user confidence.

 **Human-AI Interaction** One third of the statements (23/69) highlight how Human-AI interaction reshapes access, roles, and workflows in healthcare. AI systems, like LLMs, offer *accessibility* (6/23) to communication, enabling openness and supporting patient education, though challenges remain around appropriate abstraction for diverse users. At the same time, *roles* (4/23) are shifting. Clinicians may act as designers or supervisors, raising concerns about de-skilling and the need for clear human oversight and feedback loops. A proposed way to keep agency and accountability with humans is to use visualizations as the integration layer, interfacing AI with clinicians and leveraging an “AI-in-the-loop” paradigm [12]. In terms of *use* (5/23), effective interaction depends on prompt design, multimodal interfaces, and ensuring AI augments rather than replaces expertise. Finally, *workflows* (8/23) are evolving toward collaborative, distributed decision-making, where AI supports both clinicians and patients across care journeys, including onboarding and guidance [40], and community-driven tools. However, [evaluations remain limited in generalizability, motivating future work that explores shared design frameworks for clinical AI workflows](#) [22]. Among the VAHC papers, Lee et al. [23] exemplify this distributed configuration outside the clinic with VAIR, a VA system that couples pose estimation and biomechanical simulation with synchronized views of joint-level risk indicators, supporting expert reasoning for both retrospective injury analysis and proactive intervention planning in sports contexts. VAHC can

support addressing these challenges by enabling multi-purpose, intuitive, collaborative interfaces that enhance communication, clarify roles, coordinate next steps in patient care, and integrate AI into clinical workflows, while keeping humans in the loop.

Opportunities and Directions. The challenge here is not one of interface design, but of negotiation and human-AI agency through VA support. To this end, VA workflows must be re-examined within AI-enabled analytical environments [42]. Future work prioritizes making responsibilities, expertise boundaries, and accountability between clinicians, patients, and AI systems explicit, inspectable, and adaptable within VAHC systems. This includes mechanisms to expose decision ownership, support bridging between human and AI actors, and enable evaluation of these dynamics.

 **Model Reliability and Validation** Reliable AI in healthcare is not a property established once before deployment; it must be continuously evaluated, monitored, and validated. The workshop discussions (10/69) reveal three complementary challenges spanning the AI model lifecycle. *Evaluation before deployment* (4/10) remains a key challenge, requiring robust benchmarks, cost-benefit considerations, and exploration of automated analysis capabilities. A recent scoping review identifies a predominance of traditional diagnostic metrics (AUROC, sensitivity, specificity) over metrics optimized for post-deployment monitoring [2]. Visualization research offers a direct response: Suh et al. [41] demonstrate that aggregate metrics alone fail to communicate a model's risks and limitations to domain experts, and that visualization-based communication supports more informed assessment of model reliability. *Reliability after deployment* (4/10) raises the problem of *grounding*: anchoring AI outputs in validated knowledge and keeping systems up to date as conditions change. *Dataset shift, whether from population drift, technological change, or behavioral adaptation*, is a failure mode clinicians are often ill-equipped to detect, reinforcing the need for VA applications that help identify distributional change and model degradation to users [17]. Designing reusable visual model interfaces and coordinating their evaluation across multiple patient populations remains methodologically challenging, as many tools are disease-centric and therefore limited in their generalizability across clinical contexts. Communication and design methodologies establishing shared language across clinical and AI experts are key to detecting and addressing dataset shift [24]. Finally, *validation under uncertainty* (2/10) addresses a deeper problem: the standard assumption that model outputs can be verified against ground truth breaks down when expert knowledge itself is incomplete. Interactive VA can support expert validation in clinical settings where objective ground truth is difficult to obtain [26]. In one VAHC 2025 example of single-cell RNA sequencing cell-type classification, authors combine multiple representations (UMAP and dot plots) with SHAP values and model weights to support reasoning under epistemic uncertainty, rather than correctness checking alone [38]. Together, these examples suggest that VAHC systems can contribute to reliability not merely by displaying model outputs, but by supporting lifecycle-oriented evaluation, monitoring, and expert validation.

Opportunities and Directions. A key opportunity here lies in developing more domain-specific, context-aware model validation approaches that account for the complex, multidimensional nature of healthcare and the dynamics of these processes. Future work should establish lifecycle-oriented frameworks that integrate pre-deployment evaluation beyond aggregate metrics, post-deployment monitoring for dataset shift and population drift, and expert validation mechanisms for settings where ground truth is incomplete.

4 CONCLUSION

The VAHC 2025 workshop brought together more than 40 participants to collectively identify core challenges in Visual Analytics for Healthcare, with AI emerging as a recurring theme. Through structured group discussions informed by thematic coding of the workshop's submissions, lightning talks, and open discussion prompts, we identified five AI-centered challenge clusters across

all discussion themes: *bias and trust*, *data and infrastructure*, *explainability and communication*, *human-AI interaction*, and *model reliability and validation*. Across all clusters a consistent signal emerges: many of the most persistent AI challenges in healthcare VA are not primarily technical but also calibration problems, continuously striving for a balance between trust and evidence, explainability and clinical utility, automation and human agency. VAHC solutions are uniquely positioned to simultaneously bridge these tensions, operating at the intersection of data, models, and human reasoning. We offer these challenge clusters as a synthesis of concerns and a step toward a shared research agenda, enabling accountable, reliable, and human-centered AI in VAHC.

ACKNOWLEDGMENTS

This research was partially funded by the Österreichischen Wissenschaftsfonds (FWF, Austrian Science Fund) [10.55776/I6635], the Digital Society Initiative (DSI, Zurich), and the Engineering and Physical Sciences Resource Council (EPSRC UKRI157).

REFERENCES

- [1] Visual analytics in healthcare 2025 call for papers. <https://visualanalyticshealthcare.github.io/homepage/2025/call-for-papers.html>. Accessed: 2026-04-15. 1
- [2] E. S. Andersen, J. B. Birk-Korch, R. S. Hansen, L. H. Fly, et al. Monitoring performance of clinical artificial intelligence in health care: a scoping review. *JBIR Evidence Synthesis*, 22(12):2423–2446, 2024. doi: 10.11124/JBIES-24-00042 4
- [3] A. Arleo, R. Borgo, J. Kohlhammer, R. A. Ruddle, H. Scharlach, and X. Yuan. Reflections on the use of dashboards in the COVID-19 pandemic. *IEEE Computer Graphics and Applications*, 45(2):135–142, 2025. doi: 10.1109/MCG.2025.3538257 2
- [4] A. Arleo, A. T. Chen, D. Gotz, S. Kandaswamy, and J. Bernard. Reflections on interactive visualization of electronic health records: past, present, future. *Journal of the American Medical Informatics Assoc.*, 31(11):2423–2428, 10 2024. doi: 10.1093/jamia/ocae249 1, 2
- [5] A. Assor and J.-D. Fekete. Understanding expert exploration in EHR visualization tools: The ParcoursVis use case. In *Workshop on Visual Analytics in Healthcare*, pp. 1–7. IEEE, 2025. doi: 10.1109/VAHC69430.2025.00005 2, 3
- [6] A. Assor, H. Jeon, S. Shin, and J.-D. Fekete. Changing the paradigm from dynamic queries to LLM-generated SQL queries with human intervention. In *Workshop on Visual Analytics in Healthcare*. IEEE, 2025. [Online] <https://visualanalyticshealthcare.github.io/homepage/proceedings/2025/vis25n-sub6788.pdf>. 2
- [7] M. Aupetit and A. Abuthawabeh. InViTAG: A web application for AI-assisted exploration and grouping of health images and data. In *Workshop on Visual Analytics in Healthcare*. IEEE, 2025. [Online] <https://visualanalyticshealthcare.github.io/homepage/proceedings/2025/vis25n-sub6439.pdf>. 2
- [8] H. Balaka, M. Vaudel, and L. A. Garrison. The MoBa GWAS explorer: Designing approachable visualizations of GWAS data for a mixed audience. In *Workshop on Visual Analytics in Healthcare*, pp. 36–42. IEEE, 2025. doi: 10.1109/VAHC69430.2025.00010 2, 3
- [9] T. Baumgartl, M. Ghoniem, T. von Landesberger, G. E. Marai, S. Miksch, S. Mohr, S. Scheithauer, and N. Srivastava. Empowering communities: Tailored pandemic data visualization for varied tasks and users. *IEEE Computer Graphics and Applications*, 45(1):130–138, 2025. doi: 10.1109/MCG.2024.3509293 3
- [10] J. Bernard, D. Borland, R. S. Laramée, B. Preim, R. G. Raidou, and A. Vilanova. Future challenges and unsolved problems in health visualization. In *EuroVis 2025 - Panels and Tutorials*. Eurographics, 2025. doi: 10.2312/evt.20251119 2
- [11] S. Brückner, O. Sadare, S. Fesl, M. Scheibe, C. Lang, and S. Gilbert. Attitudes of healthcare professionals and researchers toward wearable and app derived patient generated health data. *npj Digital Medicine*, 8(1):186, Mar. 2025. doi: 10.1038/s41746-025-01568-4 3
- [12] K. Buhler, T. Holtt, T. Schultz, and P.-P. Vazquez. AI-in-the-Loop: The future of biomedical visual analytics applications in the era of AI. *IEEE Computer Graphics and Applications*, 2025. doi: 10.1109/MCG.2024.3517293 3
- [13] J. J. Caban and D. Gotz. Visual analytics in healthcare - opportunities and research challenges. *J. Am. Medical Informatics Assoc.*, 22(2):260–262, 2015. doi: 10.1093/JAMIA/OCV006 1

- [14] R. J. Chen, J. J. Wang, D. F. K. Williamson, T. Y. Chen, J. Lipkova, M. Y. Lu, S. Sahai, and F. Mahmood. Algorithmic fairness in artificial intelligence for medicine and healthcare. *Nature Biomedical Engineering*, 7(6):719–742, 2023. doi: 10.1038/s41551-023-01056-8 2
- [15] J. L. Cross, M. A. Choma, and J. A. Onofrey. Bias in medical AI: Implications for clinical decision-making. *PLOS Digital Health*, 3(11):e0000651, 2024. doi: 10.1371/journal.pdig.0000651 2
- [16] M. Ershova and G. Blasilli. Bridging service design, visualizations, and visual analytics in healthcare digital twins: Challenges, gaps, and research opportunities. In *Workshop on Visual Analytics in Healthcare*. IEEE, 2025. [Online] <https://visualanalyticshealthcare.github.io/homepage/proceedings/2025/vis25n-sub6333.pdf>. 2
- [17] S. G. Finlayson, A. Subbaswamy, K. Singh, J. Bowers, A. Kupke, J. Zittrain, I. S. Kohane, and S. Saria. The clinician and dataset shift in artificial intelligence. *New England Journal of Medicine*, 385(3):283–286, 2021. doi: 10.1056/NEJMc2104626 4
- [18] B. A. Ghassemi M, Oakden-Rayner L. The false hope of current approaches to explainable artificial intelligence in health care. *The Lancet Digital Health*, 3(11):e745–e750, 2021. doi: 10.1016/S2589-7500(21)00208-9 3
- [19] S. Ha, S. Monadjemi, and A. Ottley. Guided by ai: Navigating trust, bias, and data exploration in ai-guided visual analytics. *Computer Graphics Forum*, 43(3):e15108, 2024. doi: 10.1111/cgf.15108 2
- [20] A. Hosseini, J. Wood, A. Slingsby, R. Jianu, and M. Elshehaly. What do LLMs prioritise when adapting visualizations to user personas? In *Workshop on Visual Analytics in Healthcare*, pp. 15–21. IEEE, 2025. doi: 10.1109/VAHC69430.2025.00007 2
- [21] A. Kleinau, J. B. Au, J. Ulmer, E. Charbonnier, M. Y. Jabarulla, S. E. Long, D. Hohmann, D. Böthig, H. Brower-Rabinowitsch, A. Horke, M. Avsar, P. Beerbaum, S. Oeltze-Jafra, and T. Uden. A dashboard for visualizing health data of children with aortic valve disease. In *Workshop on Visual Analytics in Healthcare*. IEEE, 2025. [Online] <https://visualanalyticshealthcare.github.io/homepage/proceedings/2025/vis25n-sub8649.pdf>. 2
- [22] V. Lai, C. Chen, A. Smith-Renner, Q. V. Liao, and C. Tan. Towards a science of human-AI decision making: An overview of design space in empirical human-subject studies. In *Fairness, Accountability, and Transparency (FAccT)*, pp. 1369–1385. ACM, 2023. doi: 10.1145/3593013.3594087 3
- [23] C. Lee, U. Gong, T. Lin, S. Zollmann, S. A. Epsley, A. Petway, and H. Pfister. VAIR: Visual analytics for injury risk exploration in sports. In *Workshop on Visual Analytics in Healthcare*, pp. 22–28. IEEE, 2025. doi: 10.1109/VAHC69430.2025.00008 2, 3
- [24] J. K. Lennerz, U. Green, D. F. K. Williamson, and F. Mahmood. A unifying force for the realization of medical AI. *npj Digital Medicine*, 5(1):172, 2022. doi: 10.1038/s41746-022-00721-7 4
- [25] C. Meske, J. Brenne, E. Uenal, S. Oelcer, and A. Doganguen. From explainable to explanatory artificial intelligence: Toward a new paradigm for human-centered explanations through generative ai. *arXiv preprint arXiv:2508.06352*, 2025. doi: 10.48550/arXiv.2508.06352 3
- [26] G. Morgenshtern, A. Verma, S. Tonekaboni, R. Greer, J. Bernard, M. Mazwi, A. Goldenberg, and F. Chevalier. RiskFix: Supporting Expert Validation of Predictive Timeseries Models in High-Intensity Settings. In *EuroVis 2023 (Short)*. Eurographics, 2023. doi: 10.2312/evs.20231036 4
- [27] M. Musleh, R. G. Raidou, and D. Ceneda. Trustme: A context-aware explainability model to promote user trust in guidance. *IEEE Transactions on Visualization and Computer Graphics*, 2025. doi: 10.1109/TVCG.2025.3562929 3
- [28] N. Nipu, L. V. van Dijk, G. Canahuat, C. D. Fuller, and G. E. Marai. Lessons from the development and deployment of an interactive oncological risk estimator. In *Workshop on Visual Analytics in Healthcare*, pp. 29–35. IEEE, 2025. doi: 10.1109/VAHC69430.2025.00009 2
- [29] R. Owen and J. C. Roberts. Embedding empathy into visual analytics: A framework for person-centred dementia care. In *Workshop on Visual Analytics in Healthcare*, pp. 43–49. IEEE, 2025. doi: 10.1109/VAHC69430.2025.00012 2, 3
- [30] P. V. S. Pakianathan, R. Islambouli, H. McGowan, D. Branco, T. Guerreiro, and J. D. Smeddinck. Exploring human-AI interaction with patient-generated health data sensemaking for cardiac risk reduction. In *Workshop on Visual Analytics in Healthcare*. IEEE, 2025. [Online] <https://visualanalyticshealthcare.github.io/homepage/proceedings/2025/vis25n-sub4125.pdf>. 2
- [31] N. Queralt-Rosinach, R. Kaliyaperumal, C. H. Bernabé, Q. Long, S. A. Joosten, H. J. van der Wijk, E. L. Flikkenschild, K. Burger, A. Jacobsen, B. Mons, et al. Applying the FAIR principles to data in a hospital: challenges and opportunities in a pandemic. *Biomedical Semantics*, 13(1):12, 2022. doi: 10.1186/s13326-022-00263-7 3
- [32] J. Rauscher, M. Schmidt, U. Schlegel, M. Bieleke, J. Schüller, and D. A. Keim. MotiVator: Analyzing physical activity study data in lab and life through visual analytics. In *Workshop on Visual Analytics in Healthcare*, pp. 8–14. IEEE, 2025. doi: 10.1109/VAHC69430.2025.00006 2, 3
- [33] S. Reardon. Rise of Robot Radiologists. *Nature*, 576(7787):S54–S58, Dec. 2019. doi: 10.1038/d41586-019-03847-z 1
- [34] G. Ristovski, T. Preusser, H. K. Hahn, and L. Linsen. Uncertainty in medical visualization: Towards a taxonomy. *Computers & Graphics*, 39:60–73, 2014. doi: 10.1016/j.cag.2013.10.015 3
- [35] C. Rudin. Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. *Nature Machine Intelligence*, 1(5):206–215, 2019. doi: 10.1038/s42256-019-0048-x 3
- [36] D. Sacha, H. Senaratne, B. C. Kwon, G. Ellis, and D. A. Keim. The role of uncertainty, awareness, and trust in visual analytics. *IEEE Transactions on Visualization and Computer Graphics*, 22(1):240–249, 2016. doi: 10.1109/TVCG.2015.2467591 3
- [37] Z. Shang, V. Chauhan, K. Devi, and S. Patil. Artificial Intelligence, the Digital Surgeon: Unravelling Its Emerging Footprint in Healthcare – The Narrative Review. *Journal of Multidisciplinary Healthcare*, Volume 17:4011–4022, Aug. 2024. doi: 10.2147/JMDH.S482757 1
- [38] J. Shell, A. K. Davis, J. Bai, and J. Aurisano. Explaining model predictions when users aren’t sure: Xai for scRNA-seq cell-type classification. In *Workshop on Visual Analytics in Healthcare*. IEEE, 2025. [Online] <https://visualanalyticshealthcare.github.io/homepage/proceedings/2025/vis25n-sub3665.pdf>. 2, 4
- [39] S. Singh, R. Kumar, S. Payra, and S. K. Singh. Artificial Intelligence and Machine Learning in Pharmacological Research: Bridging the Gap Between Data and Drug Discovery. *Cureus*, Aug. 2023. doi: 10.7759/cureus.44359 1
- [40] C. Stoiber, D. Ceneda, M. Wagner, V. Schetinger, T. Gschwandtner, M. Streit, S. Miksch, and W. Aigner. Perspectives of visualization onboarding and guidance in va. *Visual Informatics*, 6(1):68–83, 2022. doi: 10.1016/j.visinf.2022.02.005 3
- [41] A. Suh, G. Appleby, E. W. Anderson, L. Finelli, R. Chang, and D. Cashman. Are metrics enough? guidelines for communicating and visualizing predictive models to subject matter experts. *IEEE Transactions on Visualization and Computer Graphics*, 30(7):4137–4153, 2024. doi: 10.1109/TVCG.2023.3259341 4
- [42] S. Vink, K. Byungmoo, C. Brumar, M. Yang, K. Potter, and R. Chang. Position Paper: The Opportunity for Visual Analytics in the Age of Generative AI. In A. Villanova and D. Archambault, eds., *EuroVA 2026: EuroVis Workshop on Visual Analytics*. Eurographics, 2026. doi: 10.2312/eurova.20261014 4
- [43] V. von Wyl and J. Bernard. Grand challenge: The urgent need for methods and tools to mediate between the research cultures of health sciences and visual analytics. In *Workshop on Visual Analytics in Healthcare*. IEEE, 2025. [Online] <https://visualanalyticshealthcare.github.io/homepage/proceedings/2025/vis25n-sub9654.pdf>. 2
- [44] A. Z. Wang, D. Borland, and D. Gotz. Visual analytics for causal reasoning from real-world health data. In *Workshop on Visual Analytics in Healthcare*. IEEE, 2025. [Online] <https://visualanalyticshealthcare.github.io/homepage/proceedings/2025/vis25n-sub1773.pdf>. 2
- [45] Q. Wang and R. Laramee. Ehr star: The state-of-the-art in interactive ehr visualization. *Computer Graphics Forum*, 41(1):69–105, 2022. doi: 10.1111/cgf.14424 3
- [46] M. D. Wilkinson, M. Dumontier, I. J. Aalbersberg, G. Appleton, M. Axton, A. Baak, N. Blomberg, J.-W. Boiten, L. B. da Silva Santos, P. E. Bourne, et al. The FAIR guiding principles for scientific data management and stewardship. *Scientific Data*, 3:160018, 2016. doi: 10.1038/sdata.2016.18 3
- [47] K. Xu, A. Ottley, C. Walchshofer, M. Streit, R. Chang, and J. Wenskovich. Survey on the analysis of user interactions and visualization provenance. *Computer Graphics Forum*, 39(3):757–783, 2020. doi: 10.1111/cgf.14035 3