EG UK Computer Graphics & Visual Computing (2024)
A. Slingsby and D. Hunter (Editors)

Visualizing Complex Data Decisions: Design Study for Ethical Factors in AI Clinical Decision Support Systems

S. Surodina [1,2], D. Volkova [2], A. Abdul-Rahman [1], and R. Borgo [1]

[1]King's College London, United Kingdom [2]OPORA Health Technologies LTD, United Kingdom

Abstract

Despite the proliferation of Artificial Intelligence (AI) technologies, their uptake in clinical settings has been lacking progress due to complexities of sociotechnical factors and intricacies of decision-making. Fairness and bias of predictive models, ethics and quality of training data, and corresponding compliance requirements become especially pressing while remaining fuzzy and implicit for various stakeholders who make the decisions. We present learnings and future directions from a design study with domain experts and propose a novel approach to encoding and collaborative reasoning on complex requirements for AI-Empowered Clinical Decision Support System (AI-CDSS) design based on Knowledge Graph (KG) representation. The insights will be useful to the community of visualization researchers who work on ethical AI-CDSS design and conduct design studies with clinical partners.

CCS Concepts

• ***Human-centered computing*** → *Visualization design and evaluation methods;* • ***Software creation and management*** → *Requirements analysis;* • ***Applied computing*** → *Decision analysis; Health care information systems;*

1. Introduction

The interest in AI in a clinical context has grown exponentially, fueled, among other factors, by the advances in predictive algorithms and the popularity of tools with user-friendly interfaces to generative AI models. Beyond the hype, the attempts to employ some form of data-driven ML-supported decision-making seem to penetrate every area of health research and care.

Evidence, however, indicates that the AI development and implementation approaches created for other areas are usually hard to translate due to tightly regulated, expertise-reliant, and complex healthcare processes [PLN*22]. Cookie-cutter deployment of generic AI and supporting systems in healthcare are failing and increasingly criticized for lacking what was called a Human-Centred Machine Learning perspective [Cha23]. There is limited consideration for the social, organizational, and environmental factors in AI-CDSS design and deployment, despite such factors being essential in determining AI-CDSS adoption [WZW*23]. Further, as Chomutare et al. [CTS*22] summarize, 'we know very little about the knowledge and beliefs, self-efficacy and other personal attributes of the people involved... These factors represent an important knowledge gap and require further inquiry before AI implementation in healthcare can be more fully understood.'

To solve the challenge of eliciting, abstracting, and structuring such requirements for collaborative analysis, we develop a two-stage visual analytics design approach that adapts the established Design Study frameworks and conduct a series of studies with experts. Based on that, a set of requirements and specifics is formulated, and a novel technique of mapping decisions and requirements as a KG for use requirement collection and collaborative reasoning sessions is identified as a promising tool for improving the design process of AI-CDSS.

2. Related Work

2.1. Collaborative Design Study with Domain Experts

Several methodological frameworks have been proposed for structuring problem-solving processes using visual tools: Nine Stage Framework [SMM12], Design Activity Framework [MMAM14], Nested Model [Mun09], Design by Immersion [HBH*20]. These frameworks account for stages, common methodological errors, sequence and focus of activity. Building on these foundations, we develop a collaborative design process with industry experts to collect and abstract sociotechnical factors through two iterative stages.

2.2. Ethical Factors in AI-CDSS Design

Sociotechnical decisions and requirements are important and diverse in clinical AI implementation [CTS*22]. Current and future compliance needs, cost considerations, limitations of data collection and sharing, and principles of AI-assisted decision-making are increasingly pronounced. A subset of such factors is ethical AI, an umbrella term for requirements that include data privacy, human agency, and transparency of automated decisions, to name a few. They are often vaguely defined and implicit and can be imposed by

DOI: 10.2312/cgvc.20241227

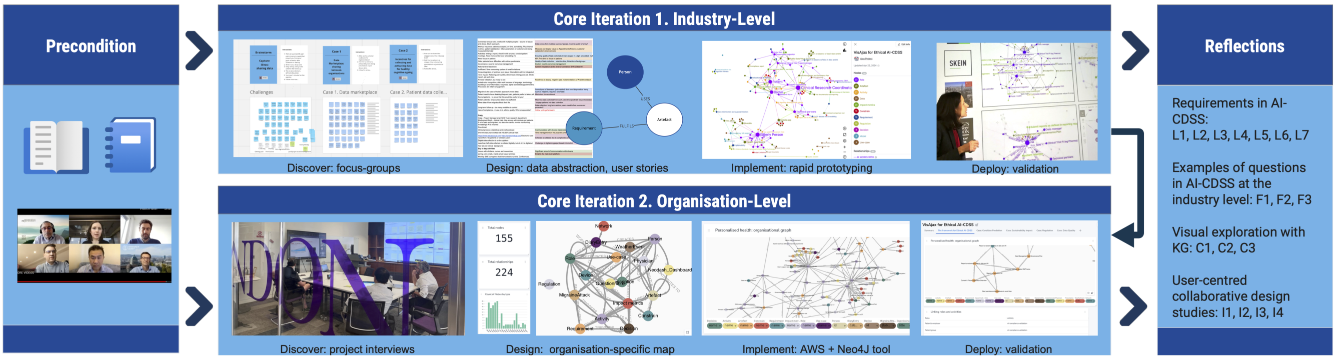

Figure 1: *Structure and milestones of the two-iteration design study framework with industry experts. The study stages' terminology and core structure is based on and extends the Nine Stage Framework [SMM12].*

current and potential regulations, values, and risk assessments. For example, data privacy and the need for large datasets for AI model training call for cross-organizational data sharing and technological solutions such as Federated Learning (FL) [BEG*19]. While FL could resolve some of the constraints, FL architectures complicate the design process further, as they have to involve decision-makers, influencers, and users across organizations where each organization brings its convoluted network of requirements and limitations. Consequently, in the design of AI-CDSS, 'AI' starts long before model development and training at the points of legislative regulation decisions, organizational strategies, and cross-organizational collaborations. There is a gap in understanding how to effectively collect and abstract such requirements, and visualize them for collective decision-making in AI-CDSS.

2.3. Knowledge Graph Visualization of Decision Networks

Knowledge graphs (KG) are a well-established tool for analyzing multi-dimensional data relationships, used in various applications, including health data [PTZY20, SLY*17, BMR21, TPZ*20]. KG has natural advantages in extracting knowledge from multi-source heterogeneous information and can serve as a basis for decision support systems [RML*22, PVGPW17, XWJF19]; however, the current applications are mostly restricted to medical data and do not account for decision networks and links with extended system requirements. Although some aspects of integration of decision context and constraints have been described in other areas, such as Enterprise KG research [Zou20, YL22], KG use for visual reasoning of decision networks has been very limited. KG for decision network analysis has multiple advantages for the critical in healthcare integration of human-in-the-loop reasoning, combining machine-readability, human comprehensibility, and the capacity to integrate into the data structure complex relationship patterns pertinent to human motivations and activities. KG applications hold promise in the emerging field of Neuro-Semantic AI [BFR*24]. We designed a KG-based tool and introduced it in the AI-CDSS design process.

3. Design Study

The design study structure augments the Nine Stage Framework with two stages of Core Iteration in order to capture both industry trends and case-specific stakeholder-driven needs, as well as implements multi-stage collaborative requirement solicitation suitable for the context of AI-CDSS design process such as a large number of diverse, cross-functional, and often implicit decision makers and influencers and differences in terminology between them.

3.1. Precondition phase

The Learn phase started with a literature review on visualization design studies and AI-CDSS, including peer-reviewed papers, non-academic industry reports, case studies, and white papers in parallel with exploratory interviews with experts (n=31).

Exploratory interviews. We invited a sample of senior decision-makers from our healthcare industry network. Participants represented roles within a typical Clinical DSS project to capture the variety of explicit and implicit motivations and concerns related to collecting and processing patient health data and were selected based on their experience in implementing AI in healthcare in various roles that would define them as stakeholders-influencers and decision-makers. Personal networks and established relationships were important factors when inviting participants, and therefore, this potentially constitutes a limitation of the study, which mostly represented individuals from the UK. The interviews were conducted via video calls and, where consented by the interviewee, recorded, automatically transcribed, and thematically structured using Otter.ai [Ott22]. Questions were primarily posed in the unstructured format with open-ended questions such as 'Please describe your experience implementing AI-CDSS,' 'What, in your opinion, prevents health organizations from setting up collaborative machine learning partnerships,' and 'What concerns do you have regarding using AI for decision support in your work? '.

Learnings from the Learn phase. As part of this phase, we captured personal experiences, identified regulations, uncertainties, and fuzzy-defined goals that were both for the organizational contexts of the interviewees and expressed opinions about overreaching industry challenges (L1). This validated the initial assumptions on the critical importance of non-technical requirements. Without a direct prompt, many comments addressed dimensions of AI system design that broadly correspond to the theoretical framework of political, economic, socio-cultural, technological, environmental, and legal (PESTEL) factors [JSW09].

Complexity and low clarity of requirements were explicitly mentioned and followed from the answers (L2). Ethics-related questions surfaced as both least understood and least crisply defined, and the absence of practical ethical AI guidances was highlighted (L3). Another marked factor was the multitude of stakeholders at different levels of needs and technical knowledge (L4), aggravated by the team dynamic challenges in the processes of requirements gathering and system design (L5) and low rates of success of AI-CDSS project implementations in the long term (L6). Knowledge gaps and blockers (L7) were also raised, such as fears about costs and complexity, uncertainty, lack of expertise in the AI implementation process, and meeting regulatory requirements. Many decisions depend on the data types they are applied to, for example, privacy sensitivity, cost and time related to the acquisition, and heterogeneity of datasets (L8). These findings allowed us to formulate top-level tasks, narrow the question list, and select participants for the first iteration of the Core phase.

3.2. Core Iteration 1. Industry-level

Discovery Focus Group. The results of the Learn stage served as the basis for the Focus Group's design to cross-validate the findings and further specify goals and abstract tasks with experts. For the co-creation session, 25 industry experts were invited, and 12 eventually took part: four in-person and eight joined the room via video conferencing. Aware and expecting issues with the pitfalls of group dynamics and moderator impact, we introduced a case-study approach modelled on the Harvard Business School case-study method [Ric17, Reb11]: the session design and the corresponding moderator guide was developed around two case presentations by a health data marketplace Chief Executive Officer and by a clinical trial manager. Both introduced a set of real challenges and needs in the context of their respective projects. Online and in-person participants used an interactive brainstorming board designed in Miro [vdDK22]. The presentations were preceded by an open-ended group brainstorming on the requirements and goals and a discussion of each case, suggesting solutions and related problems from each participant's professional experience. The session was video recorded, and later, a thematic analysis was performed using the board and sticky note ideas generated during the session by the participants.

Encoding Decision Network as a Graph. To understand the most salient themes in the participants' responses, we used inductive thematic analysis [BC12] and a hybrid approach [FAAEMC06]. Firstly, we identified use cases. Secondly, dissected goals, performance metrics, tasks, and constraints of users, complex and interdependent relationships, and multi-user workflows. Finally, focus group results were enriched and combined with the information gathered during the Learn phase (L1, L3, L7). We considered the method of user stories [LDWB16] that is traditionally applied to requirements analysis and other ways to account for the needs of collaborative explorations (L2, L4, L5). KG emerged as a suitable option for collaborative analysis and decision network representation. Themes then were manually mapped in a graph with several iterations of ontology construction. We name this version of the graph Mother-KG since it captures top-level industry requirements and expectations. Nodes represented roles and linked activities, goals, and constraints. The focus group highlighted a set of industry-level questions and challenges related to AI-CDSS implementation, including designing systems that cater to all categories of clinical stakeholders in a cost—and effort-efficient manner (F1), considering patient needs when designing systems (F2) and accounting for preventive health measures in illness-focused health systems (F3).

Implementation. A web-based KG tool, Graph Commons (`http://graphcommons.com/`), was chosen for rapid prototyping. To validate the approach, we applied questions F1-F3 as examples to discuss with other industry specialists in interactive sessions. We conducted manual KG exploration and used KG data science analysis methods — centrality, community detection, and pathfinding algorithms [XWJF19].

Learnings from the Core Iteration 1. C1: In the validation sessions, Domain Experts collaboratively contributed to the expansion of nodes and links and the reassignment of element types and links. KG analysis helped to explore deeper the patterns for more effective design and implementation (F1). Applying the clustering algorithm to Mother-KG helped to identify four communities of users and decisions based on shared requirements and limitations, and, therefore, it can be considered prototyping for five groups with common functional and visual characteristics where practical consideration would prevent customization to each of the multiple Roles (Figure 2).

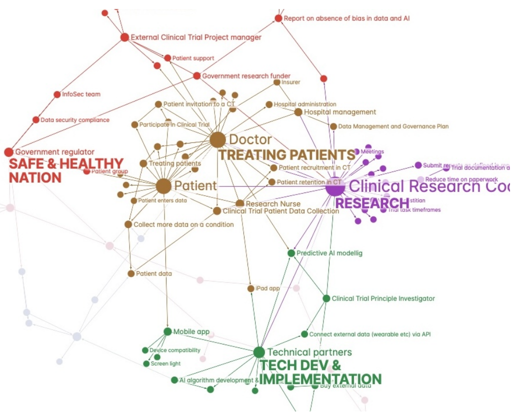

Figure 2: *Result of clustering analysis to assist the task of pattern identification.*

C2: Centrality and path-finding algorithms, as well as manual network exploration of KG, help structuring analysis and interactive discussions (F2) as illustrated in the application of these methods to Patient location in the decision-making process. Despite - in theory - the Patient's centrality to healthcare, in practice, not the Patient but Clinical Administrator is at the centre of the processes, and the Patient's needs are expressed indirectly via other actors. Visual analytics of KG allows to promptly identify such issues. C3: Path analysis similarly assists collaborative exploration. For example, there is a weak indirect link and the long path between society's need to focus on prevention and the core processes in health provision (F3). To resolve this, external stakeholders could be introduced

Proceedings published by Eurographics - The European Association for Computer Graphics.

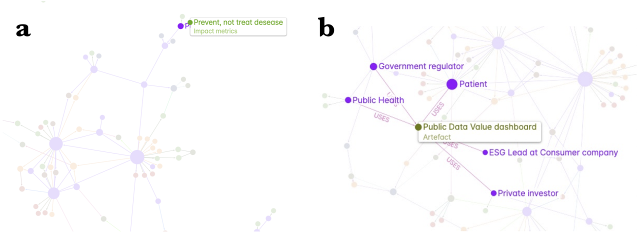

Figure 3: *Representing graph segments and edge editing: a. Despite their importance, prevention-related metrics are not prominent in a typical AI-CDSS. b. Introducing a linking artifact makes it possible to plug in the public interests in prevention directly.*

into the decision-making chain by a connecting artifact – visualization dashboard that communicates preventive impact metrics to the previously excluded actors (Figure 3). Interactive demonstration of adding such linking nodes proved easily understandable and brought up further discussions among experts. Further, the findings validated the evidence that Patient data relates to and fulfils multiple Requirements, including quality control and reporting. Therefore, it could be beneficial to incorporate the patient data schema into a graph linking to the corresponding key decision points.

3.3. Core Iteration 2. Organization-level AI-CDSS System Design

The results of Iteration 1 have informed the implementation and validation in two real-world projects. For both cases, we conducted interviews with key users. After the case interviews and construction of the case-specific KG, the Mother-KG is iteratively adjusted.

To address the need for data protection and enable better automated reasoning, the prototyped graph database was reiterated in collaboration with industry project partners and deployed on the stack of Neo4j Aura, AWS, React, Node, and NeoDash. Sub-maps are extracted to address the challenge of graph cluttering [HCSMSM00]. Parameter views of nodes, edges, and settings enable the exploration of extensive attributes without visual overload. Table views allow for filtering activities and requirements by roles and other data for more detailed exploration. The visualization tool is available at `http://viz.oporahealth.com/`. As the result of the first Core Iteration 1, we documented and summarized our findings and reflections in this paper. Figure 1 shows the stages of the design study process and the evolution of the visual interfaces and representations.

4. Lessons Learned and Further Work

Successful design and implementation of AI-driven tools in healthcare rest on a multi-dimensional set of requirements and desiderata that are usually spread across diverse roles with often misaligned motivators and needs. Using a visual network analysis can help the design and engage stakeholders. In addition, visualization offers opportunities to address many of the goals and tasks in AI-CDSS, such as exploring features that impact AI predictions. We conducted a design study and distilled the insights for better visualization design in this area discussed in previous sections (L1-L7, C1-C3). Further, the reflections on methodology and suggested future work directions for the visualization community emerged:

I1. The two-stage design study process with industry-wide mapping helps structure and guide organizational research and uncover motivations and factors that front-line users might not be aware of or have difficulty articulating. Requirements mapping and collaborative sens-making using a KG-based visual format show advantages compared to User Stories traditionally used in the industry, however more formal comparative studies will be needed.

I2. The case-study approach was successful in addressing some, but not all, limitations created by group dynamics and other specifics of focus-group format. In future studies, we suggest assembling a smaller group (up to eight participants), allowing more time up to a full day session with corresponding breaks, presenting the cases of completed projects as opposite to planned ones, and aligning the group from the start on the ontology model logic and vocabulary, such as focus on user context and their tasks.

I3. Fuzzy ethics-related requirements can be effectively clarified via collaborative KG ontology construction, which, in turn, pose their own requirements for assisting visualization design that should be explored in further depth: real-time graph construction and analysis workflows, visualization medium (large-screen projection at a room scale), the presence of a moderator, as well as the hybrid (in-person and online) format of interaction.

I4. We identified a set of visualization tasks that aim to help decision-makers to collaboratively decompose and specify high-level organizational factors and effectively share them with colleagues who have deep complementary knowledge but usually have minimal time to dedicate to the process. This calls for considering collective work, soliciting ideas, user engagement, empathizing, and understanding other people's views. In healthcare, especially when working with data on ethical requirements and patients who are not present at the decision table, human-values visual metaphors can be a valuable direction for future research.

I5. Automated analysis assists in prompting insights, improving graph readability by reducing graph cluttering and the occlusion of nodes, and approximating relationship patterns. However, real-time data collection methods force the use of pre-built standard analytical models that can be restrictive. Further KG-based decision support tool improvements can include multi-dimensionality exploration beyond click-and-expand, with node color-coding, size variation, and the use of symbols; introduce multivariate network analysis with attribute-based clustering and other types of analysis.

These findings will be useful to the visualization research community conducting design studies in close collaboration with clinical partners. The next step in our work is to expand and validate the approach in the real-world deployments with the goal to help creating more successful AI systems that deliver healthier, longer lives ethically and equitably.

Acknowledgements

This work was funded by Ageing Research at King's (ARK) and the TARA project - European Union's Horizon Europe research and innovation programme under grant agreement Nº 101057524.

References

[BC12] Braun V., Clarke V.: Thematic analysis. In *APA handbook of research methods in psychology, Vol 2: Research designs: Quantitative, qualitative, neuropsychological, and biological.*, APA handbooks in psychology®. American Psychological Association, Washington, DC, US, 2012, pp. 57–71. doi:10.1037/13620-004. 3

[BEG*19] Bonawitz K., Eichner H., Grieskamp W., Huba D., Ingerman A., Ivanov V., Kiddon C., Konečný J., Mazzocchi S., McMahan H. B., Van Overveldt T., Petrou D., Ramage D., Roselander J.: Towards Federated Learning at Scale: System Design. *Proceedings of the 2nd SysML Conference, Palo Alto, CA, USA* (2 2019). URL: http://arxiv.org/abs/1902.01046. 2

[BFR*24] Belle V., Fisher M., Russo A., Komendantskaya E., Nottle A.: Neuro-Symbolic AI + Agent Systems: A First Reflection on Trends, Opportunities and Challenges. In *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)* (2024), vol. 14456 LNAI, Springer Science and Business Media Deutschland GmbH, pp. 180–200. doi:10.1007/978-3-031-56255-6{_}10. 2

[BMR21] Biswas S., Mitra P., Rao K. S.: Relation Prediction of Co-Morbid Diseases Using Knowledge Graph Completion. *IEEE/ACM Transactions on Computational Biology and Bioinformatics 18*, 2 (2021), 708–717. doi:10.1109/TCBB.2019.2927310. 2

[Cha23] Chancellor S.: Toward Practices for Human-Centered Machine Learning. *Communications of the ACM 66*, 3 (2 2023), 78–85. doi:10.1145/3530987. 1

[CTS*22] Chomutare T., Tejedor M., Svenning T. O., Marco-Ruiz L., Tayefi M., Lind K., Godtliebsen F., Moen A., Ismail L., Makhlysheva A., Ngo P. D.: Artificial Intelligence Implementation in Healthcare: A Theory-Based Scoping Review of Barriers and Facilitators. *International Journal of Environmental Research and Public Health 19*, 23 (2022). URL: https://www.mdpi.com/1660-4601/19/23/16359, doi:10.3390/ijerph192316359. 1

[FAAEMC06] Fereday J., Adelaide N., Australia S., Eimear Muir-Cochrane A.: Demonstrating Rigor Using Thematic Analysis: A Hybrid Approach of Inductive and Deductive Coding and Theme Development. *International Journal of Qualitative Methods* (2006). 3

[HBH*20] Hall K. W., Bradley A. J., Hinrichs U., Huron S., Wood J., Collins C., Carpendale S.: Design by Immersion: A Transdisciplinary Approach to Problem-Driven Visualizations. *IEEE Transactions on Visualization and Computer Graphics 26*, 1 (1 2020), 109–118. doi:10.1109/TVCG.2019.2934790. 1

[HCSMSM00] Herman I., Computer Society I., Melanceon G., Scott Marshall M.: Graph Visualization and Navigation in Information Visualization: A Survey. *IEEE Transactions on Visualization and Computer Graphics 6* (2000), 24–43. URL: http://www.cs.brown.edu/people/rt/gd.html., doi:10.1109/2945.841119. 4

[JSW09] Johnson G., Scholes K., Whittington R.: *Fundamentals of strategy*. Pearson Education, 2009. 3

[LDWB16] Lucassen G., Dalpiaz F., Werf J. M. E. M. v. d., Brinkkemper S.: The Use and Effectiveness of User Stories in Practice. In *Requirements Engineering: Foundation for Software Quality* (Cham, 2016), Daneva M., Pastor O., (Eds.), Springer International Publishing, pp. 205–222. 3

[MMAM14] McKenna S., Mazur D., Agutter J., Meyer M.: Design Activity Framework for Visualization Design. *IEEE Transactions on Visualization and Computer Graphics 20* (1 2014), 2191–2200. doi:10.1109/TVCG.2014.2346331. 1

[Mun09] Munzner T.: A Nested Model for Visualization Design and Validation. *IEEE Transactions on Visualization and Computer Graphics 15* (2009), 921–928. doi:10.1109/TVCG.2009.111. 1

[Ott22] Otter.ai I.: Otter. ai, 2022. URL: https://otter.ai/. 2

[PLN*22] Petersson L., Larsson I., Nygren J. M., Nilsen P., Neher M., Reed J. E., Tyskbo D., Svedberg P.: Challenges to implementing artificial intelligence in healthcare: a qualitative interview study with healthcare leaders in Sweden. *BMC Health Services Research 22*, 1 (12 2022). doi:10.1186/s12913-022-08215-8. 1

[PTZY20] Pham T., Tao X., Zhang J., Yong J.: Constructing a knowledge-based heterogeneous information graph for medical health status classification. *Health Information Science and Systems 8*, 1 (2020), 1–14. URL: https://doi.org/10.1007/s13755-020-0100-6, doi:10.1007/s13755-020-0100-6. 2

[PVGPW17] Pan J. Z., Vetere G., Gómez-Pérez J. M., Wu H.: *Exploiting linked data and knowledge graphs in large organizations*. Springer, Cham, Switzerland, 2017. 2

[Reb11] Rebeiz K. S.: An insider perspective on implementing the Harvard Case Study Method in business teaching. *Online Submission* (2011). 3

[Ric17] Richards L. G.: Special session: Learning design thinking using engineering case studies. In *2017 IEEE Frontiers in Education Conference (FIE)* (2017), pp. 1–3. doi:10.1109/FIE.2017.8190560. 3

[RML*22] Rolf B., Mebarki N., Lang S., Reggelin T., Cardin O., Mouchère H., Dolgui A.: Using Knowledge Graphs and Human-Centric Artificial Intelligence for Reconfigurable Supply Chains: A Research Framework. *IFAC-PapersOnLine 55*, 10 (2022), 1693–1698. doi:10.1016/j.ifacol.2022.09.641. 2

[SLY*17] Shi L., Li S., Yang X., Qi J., Pan G., Zhou B.: Semantic Integration of Heterogeneous Medical Knowledge and Services. *Research Article Semantic Health Knowledge Graph 2017* (2017), 8–10. URL: https://doi.org/10.1155/2017/2858423. 2

[SMM12] Sedlmair M., Meyer M., Munzner T.: Design Study Methodology: Reflections from the Trenches and the Stacks. *EEE Transactions on Visualization and Computer Graphics* (2012). 1, 2

[TPZ*20] Tao X., Pham T., Zhang J., Yong J., Goh W. P., Zhang W., Cai Y.: Mining health knowledge graph for health risk prediction. *World Wide Web 23*, 4 (2020), 2341–2362. doi:10.1007/s11280-020-00810-1. 2

[vdDK22] van den Driesche C., Kerklaan S.: The value of visual co-analysis models for an inclusive citizen science approach. Inspired by co-creation methods from design thinking. *fteval JOURNAL for Research and Technology Policy Evaluation 53*, 54 (9 2022), 51–60. URL: http://repository.fteval.at/id/eprint/631, doi:10.22163/fteval.2022.571. 3

[WZW*23] Wang L., Zhang Z., Wang D., Cao W., Zhou X., Zhang P., Liu J., Fan X., Tian F.: Human-centered design and evaluation of AI-empowered clinical decision support systems: a systematic review, 2023. doi:10.3389/fcomp.2023.1187299. 1

[XWJF19] Xiang X., Wang Z., Jia Y., Fang B.: Knowledge graph-based clinical decision support system reasoning: A survey. In *Proceedings - 2019 IEEE 4th International Conference on Data Science in Cyberspace, DSC 2019* (6 2019), Institute of Electrical and Electronics Engineers Inc., pp. 373–380. doi:10.1109/DSC.2019.00063. 2, 3

[YL22] Yang B., Liao Y. m.: Research on enterprise risk knowledge graph based on multi-source data fusion. *Neural Computing and Applications 34*, 4 (2 2022), 2569–2582. doi:10.1007/s00521-021-05985-w. 2

[Zou20] Zou X.: A Survey on Application of Knowledge Graph. *Journal of Physics: Conference Series 1487*, 1 (4 2020). doi:10.1088/1742-6596/1487/1/012016. 2

Proceedings published by Eurographics - The European Association for Computer Graphics.