



City Research Online

City St George's, University of London

Citation: Elshehaly, M., Eddy, L. & Mon-Williams, M. (2023). CLEVER: A Framework for Connecting Lived Experiences with Visualisation of Electronic Records. Paper presented at the IEEE VIS: Visualization & Visual Analytics, 22-27 Oct 2023, Melbourne, Australia. doi: 10.1109/VIS54172.2023.00034

This is the accepted version of the paper.

This version of the publication may differ from the final published version. To cite this item please consult the publisher's version.

Permanent repository link: <https://openaccess.city.ac.uk/id/eprint/31117/>

Link to published version: <https://doi.org/10.1109/VIS54172.2023.00034>

Copyright and Reuse: Copyright and Moral Rights remain with the author(s) and/or copyright holders. Copies of full items can be used for personal research or study, educational, or not-for-profit purposes without prior permission or charge, unless otherwise indicated, provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way. For full details of reuse please refer to [City Research Online policy](#).

CLEVER: A Framework for Connecting Lived Experiences with Visualisation of Electronic Records

Mai Elshehaly*
City, University of London

Lucy Eddy†
University of Bradford

Mark Mon-Williams‡
University of Leeds

ABSTRACT

The disconnect between insights generated from data and real-life practices of decision makers presents a number of open questions for visual analytics (VA). In public service planning, routine data are often perceived as unavailable, biased, incomplete and inconsistent across services. Decision makers often rely on qualitative data - sometimes collected through co-production - to understand the lived experience of communities before formulating a decision. We followed a subjectivist case study approach and immersed ourselves in ongoing co-production activities over the course of one year, to capture how VA can support the dialogue between population health decision-makers and the communities they serve. We present a framework for Connecting Lived Experiences with Visualisation of Electronic Records (CLEVER). The framework regards visualisation as a central component in a complex adaptive decision-making ecosystem and highlights the need to structure domain knowledge across decision contexts in Population Health Management (PHM) at clinical-, service- and district-levels. Our process for developing an initial framework comprised three steps: (i) we elicited decision-making tasks through a series of qualitative data collection activities; (ii) we developed a preliminary domain model to capture data views and a subjective view of the world through human stories; and (iii) we developed a series of visualisation prototypes to instantiate the framework and demonstrated them regularly to stakeholders. In future work, we will conduct ‘deep dives’ to systematically study the role of VA in individual stages of the framework.

Index Terms: Human-centered computing—Visualization—Visualization design and evaluation methods.

1 INTRODUCTION

Routinely collected data, such as Electronic Health Records (EHRs), have immense potential to improve public services and guide efforts to design health interventions [13]. The realisation of this potential requires several challenges to be addressed to bridge the well documented gap between data insights and real-world decision-making [6, 7]. These challenges are exacerbated in Population Health Management (PHM), where different “*departments seek to connect practice to policy for change to happen locally*” [2]. This cross-service approach highlights the need for intelligence to be reliably shared across organisations. PHM further emphasises the role of co-production with local communities in affecting the decision-making process. The diversity of decision-making contexts and the information considered presents several challenges to stages of the visualisation design study methodology, where visualisation users are expected to externalise which data-driven insights they require to make decisions. Conversely, visualisation researchers are often presented with qualitative information, where the link between insight and decision-making is poorly captured.

*e-mail: Mai.Elshehaly@city.ac.uk

†e-mail: l.h.eddy@bradford.ac.uk

‡e-mail: m.mon-williams@leeds.ac.uk

This paper presents a preliminary framework to facilitate the role of visual analytics (VA) researchers and guide their engagement with PHM stakeholders during the *Discover* and *Design* stages of the design study methodology [19]. The framework addresses the need to **C**onnect the **L**ived **E**xperiences of stakeholders with **V**isualisation of **E**lectronic **R**ecords (CLEVER). Through presenting an explicit structure to guide the integration of domain knowledge (e.g., the lived experience of health service users) with data visualisation, the CLEVER framework offers a systematic approach to bridge the gap between stakeholders’ scattered knowledge of the world and data-driven insights. It is based on immersing ourselves in an ongoing case study, over the course of one year, at a population health research facility that maintains a district-wide data repository, linking routine data across education, housing, policing, environment, health, and social care [20]. The study captured the perspectives of policymakers, practitioners, and communities on ways in which data can help improve services for neurodivergent children.

While this preliminary framework does not offer actionable design principles for visually integrating lived experiences with data, it presents a birdseye view of the soft intelligence required for visualisation to facilitate the dialogue between decision-makers and data; and a roadmap for future work to empirically establish the utility of such dialogue in decision-making. Our contributions in this paper are twofold: (i) We characterise the problem space of PHM through building a preliminary domain model that captures both data and human perspectives. (ii) We propose CLEVER, a socio-technical framework that seeks to answer the questions: (a) what knowledge constructs should visualisation researchers expect stakeholders to have; and (b) what design alternatives can connect this knowledge to data visualisation in PHM? In what follows, we motivate these two questions with relevant literature and a running example.

1.1 The Role of Stakeholders’ Knowledge

Effective visualisations build a mapping between users’ real-world knowledge and their interaction with visual analytics. This mapping aims to bridge gaps in *data*, which hinder the ability to transform them into insights [15], and gaps in users’ *implicit knowledge* of the world, which could hinder their ability to make better-informed decisions. Early work by van Wijk regarded domain knowledge to be a function of VA [21]. A more recent view considers that knowledge priors predate the VA environment, therefore, domain knowledge is expected to act as a “*first-class artifact... that is both an input and an output of VA activities*” [18]. We adopt this view as we examine the input (users’ knowledge of the world) and output (anticipated decision-making tasks) of PHM stakeholders.

Consider, for example, the decision to provision a new service for autistic children within a district. Selecting from a range of possible services requires a quantifiable measure of demand as well as contextual understanding of the anticipated improvements in children’s outcomes and experiences. While EHRs offer a starting point to identify demand, decision-makers acknowledge several barriers (including long waiting lists, inconsistent data recording practices, etc.) that deprive a large proportion of the autistic population from having their care journey adequately represented in EHRs. Therefore, disparate pieces of information need to be collected from individuals and organisations along the service referral pipeline.




Figure 1: Overview of our subjectivist methodology. Qualitative activities including interviews, focus groups and immersion, were conducted in the context of the child neurodiversity case study. White boxes represent steps in our future work to establish the real-world utility of the framework.

In visual analytics, knowledge is often classified into **operational** (informing users’ interaction with a system) and **domain** knowledge (supporting the interpretation of presented information) [5]. Domain knowledge is often viewed from a data-centric lens, in which it determines the users’ ability to interpret data [18], and reason about the “*existence and source*” of error [15]. This data-focused definition assumes a high level of (i) *centralisation*, where individual users are assumed to know enough about the possible mismatches between data and real-world practices, and (ii) *agency* where users are comfortable and confident in mapping insights to real-world situations. For example, in the clinical and epidemiological decision-making context presented by McCurdy et al. [14], individual users knew enough to explain data discrepancies which may impact their interpretation. Similarly, Kerrigan et al.’s taxonomy of domain knowledge elicitation considers machine learning experts as the main stakeholders [12]. A common thread in these studies is that they regard knowledge elicitation from a *human-in-the-loop* perspective, where data and systems take center stage and the role of human expertise is structured around the stages of data pre-processing, model design, etc. This is rarely the case in PHM, where decisions begin and end with human stories [17], while data (if available) play a complementary role to back them with evidence.

1.2 Designing Knowledge-Assisted Visual Analytics

The adapted process model of visualisation by Federico et al. [8] illustrates how visualisation connects between a cognitive (human) space and a computational (machine) space. Our work focuses on considerations for VA design to facilitate a two-way mapping between users’ tacit knowledge K^t and explicit knowledge K^e . Tacit knowledge is defined as knowledge that is not consciously controlled and is often difficult to express or communicate. It is usually acquired through experience or practice and can be transformed into explicit knowledge through a process of codification [16]. In VA research, a user’s tacit knowledge is expected to grow as a result of an *internalisation* process, whereby the user interfaces with knowledge presented in the visualisation [8, 18]. A knowledge *externalisation* process, on the other hand, refers to the elicitation of tacit knowledge from humans to be expressed in the form of rules, models, and computational constructs that can inform data analysis [8], interpretation [14, 21] and guidance for user interaction [4].

While the role of visualisation is clearly defined in knowledge *internalisation*, it is less clear in the literature if and how such a role exists in supporting knowledge elicitation and *externalisation*. Understanding design considerations that can establish this role is critical if visualisation is to support the collation of evidence from routine data and ‘soft intelligence’ across the PHM ecosystem.

2 THE CLEVER FRAMEWORK

We followed a subjectivist case study approach in developing the CLEVER framework for eliciting implicit knowledge in PHM (Figure 1). Similar to the process of developing a socio-technical evaluation framework for health information systems by Yusof et al. [22], our case study approach helped us gain an in-depth understanding of the contexts surrounding current and prospective use of routinely

linked data in a specific PHM problem; namely, child neurodiversity. A literature review helped us identify the key data analytic questions that support PHM research in this area. It also helped us identify stakeholder groups who interact and interface with the three key contexts in PHM decision-making at **clinical**, **organisational**, and **district** levels [13]. We conducted eleven interviews and three focus group discussions with decision makers at all three levels. We then immersed ourselves in regular research meetings with a PHM research team (of which two members are co-authors) and in co-production sessions hosted by district-level service providers for parents and families of neurodivergent children.

We used a deductive analysis approach to map knowledge constructs that were captured in our qualitative activities to a domain model for PHM. We detail this approach in Section 2.1. Next, we developed an initial framework for human-data collaboration and used this framework to develop a series of visualisation prototypes. The prototypes were shared regularly at PHM research meetings and feedback was used to iterate both the software and the framework. We showcase a snapshot of these iterations in Section 2.2.

2.1 Preliminary Domain Model

What constitutes a useful structure for domain knowledge in PHM? We started with mapping statements captured in our interviews to components of the structural framework for explicit domain knowledge developed by Rind et al. [18]. The framework defines **concepts** as the first modular knowledge construct to be codified, and identifies mappings of concepts to (a) datasets, where knowledge manifests in relevant data items; and (b) VA environment, where knowledge is utilised. Our domain model identified two stages of knowledge construction in these mappings (Figure 2):

Manifestation of a reference population. In PHM, a key concept is the *reference population*, which defines a cohort of interest to a specific public health challenge. This concept can be manifested in data as a mapping $d : \mathbf{R} \rightarrow \mathbf{C}$ where \mathbf{R} is the set of data references defining the cohort (e.g., individuals in a certain locality), and \mathbf{C} are characteristics that map members of the population to relevant features (e.g., ethnicity) or events (e.g., diagnoses) in electronic records. Identifying the reference population requires human expertise for direct labeling and/or indirect characterisation of individuals.

To illustrate this, we consider our autism example from Section 1.1, where two relevant constructs were identified in our analysis: Autistic Spectrum Disorder (ASD) and Developmental Disorder of Motor Function (DDMF). An examination of the Connected Bradford dataset, which contains healthcare records for over 600,000 citizens, revealed that only 154 cases of children with relevant DDMF diagnoses were annotated by clinicians. In comparison, we identified 6,249 Autism cases (despite an estimated prevalence five times lower than DDMF). This can be due to any of the discrepancies listed by McCurdy et al. for public health data [14] (e.g., missingness, temporal discrepancies, etc.). Therefore, diagnostic constructs based on direct labeling alone can lead to the exclusion of large numbers of people, especially from marginalised communities [11]. Expanding **R** to include these high risk cohorts is central to decision-making.

We define population manifestation in data as $d_R : \mathbf{R} \cup \mathbf{R}^e \rightarrow$




Figure 2: Initial framework for structuring soft intelligence and the lived experiences of visualisation stakeholders in PHM.

$\mathbf{C} \cup \mathbf{C}^e$, where \mathbf{R}^e extends the original reference set by direct notation of additional references using explicit domain knowledge \overline{K}^e . \mathbf{C}^e represents a manifestation of externalised characteristics that define additional inclusion criteria. Both \mathbf{R}^e and \mathbf{C}^e vary in terms of relevance and definition from one decision-making context to another. At the clinical context, decision-makers have first hand experience of interacting with references (patients) in the data. This context has a high level of agency and knowledge centralisation. Therefore, clinicians can offer domain knowledge constructs that inform the direct labeling of references in \mathbf{R}^e . Experts can contribute knowledge of common practices in selecting clinical codes, and consensus on what constitutes a valid code list (e.g., [9]). Additional characteristics \mathbf{C}^e capture cohorts who risk having a condition but are missing a relevant diagnostic code. These \mathbf{C}^e are described by experts in the form of composite criteria (e.g., a child from an ethnic minority in a deprived area with learning difficulties).

As we step away from clinical contexts, the definition of \mathbf{C}^e becomes less intertwined with data and more reliant on what communities convey to decision-makers about their lived experiences. Factors that contribute to the relevance of each characteristic include the reliability of data sources that could support or counter claims made by communities, and known inequalities that present barriers to the acquisition of such claims (e.g., from marginalised and hard to reach communities). We define two measures for each characteristic in \mathbf{C}^e : (i) a relevance score $\delta_j^e(c_i)$ measures the computational relevance of the i^{th} characteristic to the j^{th} concept; and (ii) the match between concepts and communities' lived experiences, which we denote with with an "imaginary part" symbol $I_j^e(c_i)$ to specify that this measure does not come from data, though it assigns a weight to c_i in specific decision-making contexts.

Composition of care pathways. Decision-makers trace the referral routes of cohorts from one service to another to identify gaps and inform resource provision. This means that for each data reference, a sequence of events $X \rightarrow v_1, \dots, v_q$ can be fetched from data to define service usage pathways. Only a small fraction of these events is typically relevant to service providers. For example, while the EHR for the autism cohort includes hundreds of event types that include GP visits, hospital admissions, prescriptions, etc., the story told by communities typically involves a few milestones (e.g., identification, assessment, etc.) during these service encounters. The concept for each milestone can be manifested in different data sources. Decision-makers are interested in the flow across these milestones to identify inequalities and missed opportunities along the pathway.

Explicit knowledge about care pathways can be captured at two

levels of granularity. A within-service level defines a localised journey that includes clinical history and events typically considered by clinical managers. We define the clinical pathway P_{CI} :

$$P_{CI}(\mathbf{r}, \mathbf{c}) = \{(v_1, \dots, v_s) \in \{v_1, \dots, v_q\} | s \leq q, \delta(v_i) \geq \theta\} \quad (1)$$

to map the reference population \mathbf{r} characterised by \mathbf{c} to an event subsequence $(v_1, \dots, v_s)_j$, where the relevance of each individual event from data $\delta(v_j)$ is compared to a domain-defined threshold θ . A majority of the challenges reported by communities, however, go beyond this localised clinical pathway. Issues of disjointed services and obscure processes create high levels of frustration. PHM decision makers at the district level are interested in those between-service issues as they report and assess service performance indicators. In these contexts, the care pathway becomes a directed graph, in which each node is a subsequence P_{CI} . The set of edges between the subsequences is characterised by several pieces of evidence. Communities' narratives, decision-makers' expectation of referral routes, and performance benchmarks act as human-centred edge attributes that estimate the direction and magnitude of the flow between services, and barriers experienced along the journey.

We define the composition of population pathways as a labeled multidigraph $P_W = (V, E)$, where $V = \{P_{CI_j}\}$ is the set of all possible localised (i.e., within-service) pathways that make up the vertices of the cross-service pathway; E is the set of arcs connecting services. We define Σ_V and Σ_E as finite vocabularies that define possible labels for vertices and edges. Σ_V may codify service types and milestones. Σ_E may codify measures of expected and observed patient flow across services. The ability to tell a story about people's journeys has an affective component that can persuade an audience and inspire change. A careful human-centred definition of Σ_V and Σ_A must capture these affective components and use a language that visualisation viewers can understand and to which they can relate.

Synthesising causal pathways and simulating interventions. To complement the human-centred service pathway narrative above, evidence synthesis that informs decision-making typically requires causal analysis to examine factor-outcome interactions along pathways that have already been defined by communities, while uncovering new unexpected relationships through data-driven analyses. Jin et al. define a Hawkes process based approach to modeling these causal relationships for event sequences [10]. In our domain model, we define a causal pathway P_C which augments the real-world narratives captured in P_W in one of three ways: (a) add new vertices to V or modify their labels Σ_V ; (b) modify the sequence (v_1, \dots, v_s) that make up the subsequences in P_{CI_j} ; and/or (c) add new links to A with their corresponding Σ_A .

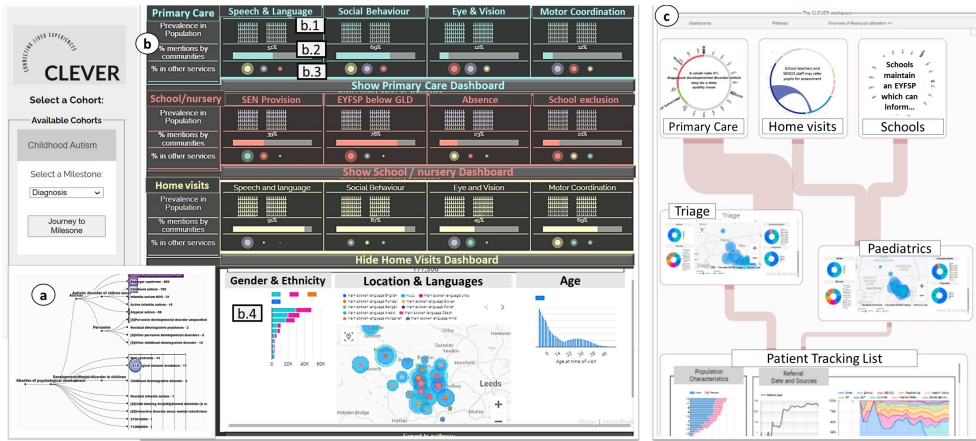


Figure 3: Examples of implementing parts of the CLEVER interface to include reference population selection (a,b) and pathway creation list views (c).

2.2 Framework for Visualisation Design

The domain model in Section 2.1 sets a framework for structuring knowledge constructs that can support decision-making in PHM. Here, we explore ways in which this structure can guide interface design and support the externalisation, sharing, and enhancement of some of these knowledge constructs. This initial framework is not exhaustive. Each cell in Figure 2 can be explored as a standalone design study and we acknowledge that more work is needed to identify the most appropriate visual design alternatives for individual cells and for transitions across cells. However, in giving a few examples, we seek to propose one way in which the CLEVER framework can inform design ideas that motivate users to externalise and share their tacit understanding of people (the reference population) and situations (care pathways). Further work that underpins the externalisation of cause-effect relationships (e.g., causal pathways and what-if scenarios) is left for future work.

2.2.1 Selecting the reference population

Clinicians and researchers work together to identify the reference set R^e by considering their domain knowledge of disease coding ontologies and practices, and national guidelines that govern these practices. They use this information to either select records directly or define rules for inclusion of clinical codes. To support the externalisation of this information, visualisation should display semantic relationships between clinical codes in light of their manifestation in data (e.g., Figure 3(a)). This enables a dialogue between clinicians and service providers on matters regarding data quality and access to diagnosis. Visualisation researchers should then seek providers' input on the relevance of specific characteristics $\delta_j^e(c_i)$. For example, if speech and language concerns are known to be discussed in a specific service's context (e.g., with health visitors), then $\delta_j^e(c_i)$ is increased for data manifestations of these concerns c_i wrt to the autism diagnosis concept in this service-level context.

Figure 3(b) shows a matrix view in which different characteristics (e.g., speech and language, social behaviour, etc.) are displayed for individual service contexts. For each characteristic c_i , two pieces of information are displayed: (i) the prevalence of c_i 's values in this service's cohort (b.1) as well as in other services (b.3) are shown as indicators of $\delta_j^e(c_i)$; (ii) a qualitative measure that shows the percentage of participants who have mentioned this characteristic is shown as a measure of $I_j^e(c_i)$ (b.2). The lived experiences of patient communities can inform further measures of I^e . For example, parents of autistic children have expressed that for years their concerns about social behaviour were dismissed by clinicians and therefore not recorded in their children's EHR. Our suggested design

ensures that this c_i is weighted positively in **b.2**, even when lower data manifestations are displayed in **b.1** and **b.3**.

Additional information that can supplement a users' understanding of characteristics contributed by services and assess their value to the definition of the reference population may include demographic information that are known to capture inequalities or uncertainties. For example, if demographic information (e.g., locality, ethnicity, etc.) is known to introduce biases, summary information or measures of bias can be presented for each. We take the former approach by introducing the idea of embedded dashboards in our prototype, where dashboards are overlaid on the page as pop-up screens and dynamically hidden (Figure 3(b.4)). Alternatively, Borland et al. include measures of bias directly in the visualisation [3].

2.2.2 Creating service pathways

Figure 3(c) presents one possible way to show service pathways and align population characteristics with wider cross-service information. Visualisation techniques were chosen to match the type of information needed within each localised pathway P_{Cl_j} . The vertical position of each P_{Cl_j} is determined based on qualitative narratives that were captured in communities' stories. The top row includes three pathways from primary care, home visits and schools. Participants told us that these are the most common services they encountered in the early identification stage for Autism. We used Bremer's "Using data storytelling with chord diagram" [1] to visualise the flow between different traits discussed within each of these instances of P_{Cl_j} . Subsequent stages in the journey included a triage process, referrals to see a paediatrician, and later being put on a waiting list. Each of these instances of P_{Cl_j} comes with its own set of questions about what happens to people in the corresponding stage. We used embedded dashboards to address these questions. Finally, pipe-shaped links are used to link the pathways into a higher-level P_W .

3 CONCLUSION AND FUTURE WORK

We presented a framework that provides structure to domain knowledge coming from different stakeholders in the PHM decision-making process, with an aim to inspire design studies that address visual analytic needs in this domain. Our preliminary domain model defines four knowledge constructs that marry both data and human expertise. The CLEVER framework presents opportunities to design novel interactive visualisations that facilitate cross-context integration and communication. We intend to continue iterating the framework through a series of "deep dives" into each of the contextual cells to further identify the most critical knowledge constructs contributed by stakeholders and the role that visualisation can play in this externalisation process.

REFERENCES

- [1] Using data storytelling with chord diagram. <https://www.visualcinnamon.com/2014/12/using-data-storytelling-with-chord/>, Oct 2020. [Online; accessed July-2023].
- [2] What is population health? <https://www.cdc.gov/pophealthtraining/whatis.html>, Oct 2020. [Online; accessed June-2022].
- [3] D. Borland, J. Zhang, S. Kaul, and D. Gotz. Selection-bias-corrected visualization via dynamic reweighting. *IEEE Transactions on Visualization and Computer Graphics*, 27(2):1481–1491, 2021. doi: 10.1109/TVCG.2020.3030455
- [4] D. Ceneda, T. Gschwandtner, T. May, S. Miksch, H.-J. Schulz, M. Streit, and C. Tominski. Characterizing guidance in visual analytics. *IEEE transactions on visualization and computer graphics*, 23(1):111–120, 2016.
- [5] C. Chen. Top 10 unsolved information visualization problems. *IEEE computer graphics and applications*, 25(4):12–16, 2005.
- [6] E. Dimara and J. Stasko. A critical reflection on visualization research: Where do decision making tasks hide? *IEEE Transactions on Visualization and Computer Graphics*, 28(1):1128–1138, 2021.
- [7] M. Elshehaly, K. Sohal, T. Lawton, M. Bryant, and M. Mon-Williams. Creative visualisation opportunities workshops: A case study in population health. In *2022 IEEE Evaluation and Beyond - Methodological Approaches for Visualization (BELIV)*, pp. 11–19. IEEE Computer Society, Los Alamitos, CA, USA, oct 2022. doi: 10.1109/BELIV57783.2022.00006
- [8] P. Federico, M. Wagner, A. Rind, A. Amor-Amorós, S. Miksch, and W. Aigner. The role of explicit knowledge: A conceptual model of knowledge-assisted visual analytics. In *2017 IEEE conference on visual analytics science and technology (vast)*, pp. 92–103. IEEE, 2017.
- [9] M. Fonseca, B. MacKenna, A. Mehrkar, O. Collaborative, C. E. Walters, G. Hickman, L. Fisher, J. Pearson, P. Inglesby, S. Bacon, et al. Primary care coding activity related to the use of online consultation systems or remote consulting: an analysis of 53 million peoples’ health records using opensafely. *medRxiv*, pp. 2023–01, 2023.
- [10] Z. Jin, S. Guo, N. Chen, D. Weiskopf, D. Gotz, and N. Cao. Visual causality analysis of event sequence data. *IEEE transactions on visualization and computer graphics*, 27(2):1343–1352, 2020.
- [11] B. Kelly, D. Mason, E. S. Petherick, J. Wright, M. A. Mohammed, and C. Bates. Maternal health inequalities and gp provision: investigating variation in consultation rates for women in the born in bradford cohort. *Journal of Public Health*, 39(2):e48–e55, 2017.
- [12] D. Kerrigan, J. Hullman, and E. Bertini. A survey of domain knowledge elicitation in applied machine learning. *Multimodal Technologies and Interaction*, 5(12), 2021. doi: 10.3390/mti5120073
- [13] T. Lippeveld. Routine health information systems: the glue of a unified health system. In *Keynote address at the Workshop on Issues and Innovation in Routine Health Information in Developing Countries, Potomac, March*, pp. 14–16, 2001.
- [14] N. McCurdy, J. Gerdes, and M. Meyer. A framework for externalizing implicit error using visualization. *IEEE transactions on visualization and computer graphics*, 25(1):925–935, 2018.
- [15] N. Mccurdy, J. Gerdes, and M. Meyer. A framework for externalizing implicit error using visualization. *IEEE Transactions on Visualization and Computer Graphics*, 25(1):925–935, 2019. doi: 10.1109/TVCG.2018.2864913
- [16] N. R. Milton. *Knowledge acquisition in practice: a step-by-step guide*. Springer Science & Business Media, 2007.
- [17] L. Patricio, D. Sangiorgi, D. Mahr, M. Čaić, S. Kalantari, and S. Sundar. Leveraging service design for healthcare transformation: Toward people-centered, integrated, and technology-enabled healthcare systems. *Journal of Service Management*, 31(5):889–909, 2020.
- [18] A. Rind, M. Wagner, and W. Aigner. Towards a structural framework for explicit domain knowledge in visual analytics. In *2019 IEEE Workshop on Visual Analytics in Healthcare (VAHC)*, pp. 33–40. IEEE, 2019.
- [19] M. Sedlmair, M. Meyer, and T. Munzner. Design study methodology: Reflections from the trenches and the stacks. *IEEE Transactions on Visualization and Computer Graphics*, 18(12):2431–2440, 2012. doi: 10.1109/TVCG.2012.213
- [20] K. Sohal, D. Mason, J. Birkinshaw, J. West, R. R. McEachan, M. Elshehaly, D. Cooper, R. Shore, M. McCooe, T. Lawton, et al. Connected bradford: a whole system data linkage accelerator. *Wellcome Open Research*, 7, 2022.
- [21] J. J. Van Wijk. The value of visualization. In *VIS 05. IEEE Visualization, 2005.*, pp. 79–86. IEEE, 2005.
- [22] M. M. Yusof, J. Kuljis, A. Papazafeiropoulou, and L. K. Stergioulas. An evaluation framework for health information systems: human, organization and technology-fit factors (hot-fit). *International Journal of Medical Informatics*, 77(6):386–398, 2008. doi: 10.1016/j.ijmedinf.2007.08.011