





The MoBa Pregnancy and Child Development Dashboard: A Design Study

R. Ziman¹ , B. Budich^{1,2} , M. Vaude^{1,3}  and L. Garrison¹ 

¹ Department of Informatics, University of Bergen, Norway, ² Department of Simulation and Graphics, University of Magdeburg, Germany, ³ Department of Clinical Science, Mohn Center for Diabetes Precision Medicine, University of Bergen, Norway,

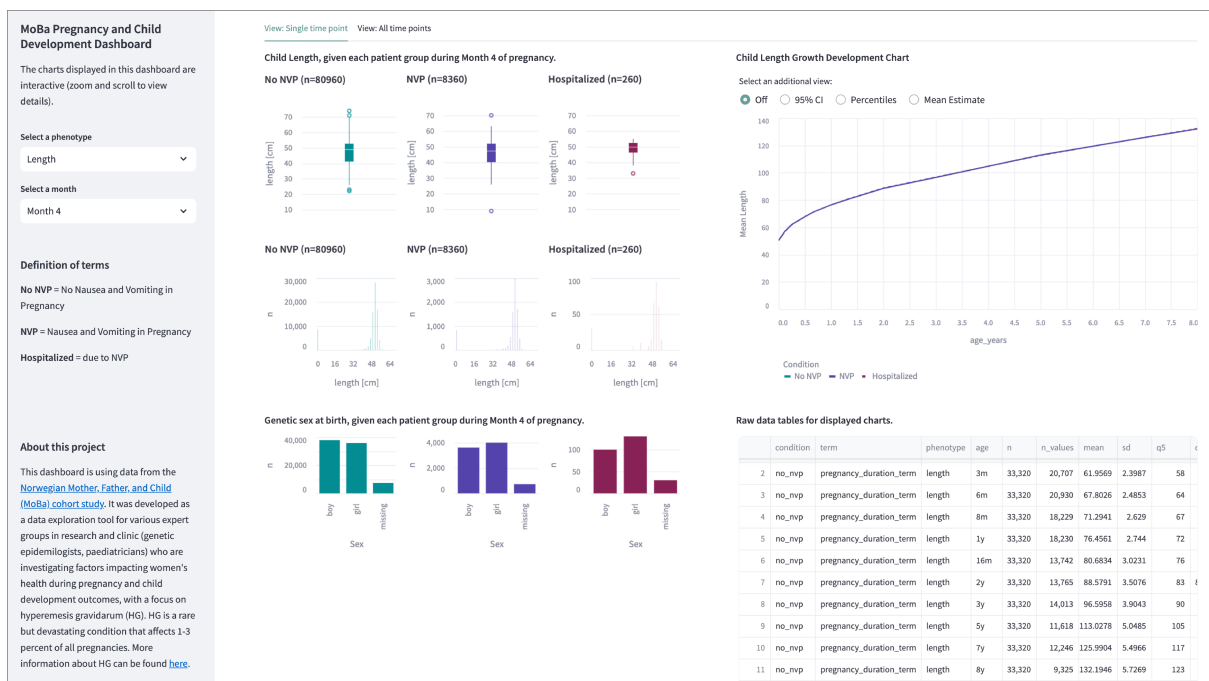


Figure 1: The landing page of the prototype for our web-based MoBa Pregnancy and Child Development Dashboard, designed to aid domain experts investigating hyperemesis gravidarum (HG) and other pregnancy-related health outcomes, using data from the Norwegian Mother, Father, and Child (MoBa) Cohort Study [MBV*16].

Abstract

Visual analytics dashboards enable exploration of complex medical and genetic data to uncover underlying patterns and possible relationships between conditions and outcomes. In this interdisciplinary design study, we present a characterization of the domain and expert tasks for the exploratory analysis for a rare maternal disease in the context of the longitudinal Norwegian Mother, Father, and Child (MoBa) Cohort Study. We furthermore present a novel prototype dashboard, developed through an iterative design process and using the Python-based Streamlit App [TTK18] and Vega-Altair [VGH*18] visualization library, to allow domain experts (e.g., bioinformaticians, clinicians, statisticians) to explore possible correlations between women's health during pregnancy and child development outcomes. In conclusion, we reflect on several challenges and research opportunities for not only furthering this approach, but in visualization more broadly for large, complex, and sensitive patient datasets to support clinical research.

CCS Concepts

• Applied computing → Life and medical sciences; • Human-centered computing → Visualization design and evaluation methods; Interaction design process and methods; Interface design prototyping;

1. Introduction

Ongoing clinical research using data from the Norwegian longitudinal Mother, Father, and Child (MoBa) Cohort Study [MBV*16] investigates many health and medical factors that can impact pregnancy health and child outcomes. Of particular interest to our collaborative network is a condition known as *hyperemesis gravidarum* (HG), which is marked by severe, *prolonged nausea and vomiting during pregnancy* (NVP) with at least one related *hospitalization*. This rare but devastating condition affects 0.3–10.8% of pregnant women [FTG*19], and the health consequences for both the mother and child during and after pregnancy are poorly understood. Clinical researchers study longitudinal data and phenotypic traits (the observable physical characteristics of a person's genetic make-up [WT08]) such as height and weight, in relation to health conditions like HG to uncover patterns and possible correlations, and inform treatment development and clinical protocols. Research into the causes and outcomes of HG can be supported by visualizations, which have become central to public health and clinical research as they facilitate the exploration and analysis of large, complex datasets. To this end, requirements of a successful visual support tool in public health include: providing an overview of the data, enabling analysts to integrate expert knowledge (that is, to focus on the dataset and dimensions of interest), using visualizations familiar to the domain expert user, and providing visual support for, among other tasks, association and comparison of the data [PL20]. Understanding the prevalence of a disease and identifying health risks are two relevant public health indicators [ZS17] that a successful visual analytics tool can also help experts achieve.

However, this endeavour faces several challenges, particularly with the MoBa dataset. First, it involves extracting, processing, and analyzing specific data from a complex, heterogeneous dataset containing thousands of variables collected from several disparate patient questionnaires. This can be a cumbersome, often repetitive process of searching and pre-processing data as hypotheses evolve. Second, patient confidentiality must be maintained, therefore individual-level data cannot be reported and must instead be aggregated for the purpose of analysis and dissemination. Third, missing data and data input errors are pervasive throughout the questionnaires, for example, unreported genetic sex of the baby at birth, or incorrect values entered for measurements such as weight and height. Women afflicted with HG or other serious conditions are also more likely to be lost to follow-up in subsequent questionnaires than women who experience otherwise healthy pregnancies, resulting in their under-representation in the study dataset. These last two challenges introduce uncertainty and potential error in the visualization pipeline, which may in turn lead to incorrect interpretation of the data. Such compounding effects make the study of an already rare and poorly understood disease more difficult.

In this paper, we explore how visualization may help to allay some of these issues with the following contributions:

1. Application of the design study methodology [SMM12], the Five Design Sheets (FdS) Method [RHR15], and an iterative design process to develop a public health dashboard.
2. A web-based visual analytics dashboard prototype to investigate possible phenotypic correlations between pregnancy health and child developmental outcomes, with a focus on HG.

2. Related Work

We contextualize our approach through the lens of two core topics that inspired our work: visual analysis of longitudinal data, and visualization dashboards for public health.

A rich body of visualization research exists for visualizing time-oriented data, summarized by Aigner et al. [AMST11]. Of note, we draw inspiration from techniques highlighted in this work for comparing groups at a single time point versus multiple timepoints, and the ability for the user to change flexibly between these two views. Albers et al. investigate various techniques for visualizing aggregated time series data and resultant user performance, and identify box plots as outperforming other visual encodings (e.g., line charts) for aggregate comparison tasks [ACG14]. Further, Javed et al. demonstrate that small multiples are beneficial for time series data to avoid occlusion from overlapping datasets [JME10]. We explore the combination of these techniques in our visualization solution. With the domain expert's focus on statistical analysis in the context of this work, we draw inspiration from validated techniques that integrate both visual analytics and statistical approaches. The Cohort Comparison ("CoCo") visualization tool and associated taxonomy of patient summary metrics (e.g., number of records) and time-based event metrics (e.g., prevalence of an event) provides guidance for comparing two cohorts [MDM*15]. We apply some of these metrics to our clinical area in pregnancy health while expanding the capability of comparing additional cohorts.

The use of visual dashboards to support clinical and public health research is a frequent strategy, with solutions tailored for specific subdomains, research questions, and user goals. While numerous works target strategies at the genome level [NHG19], comparatively few approaches explicitly target phenotype exploration and analysis. PhenoStacks [GGC*16] and PhenoBlocks [GHC*16] are both visual analytics approaches that facilitate expert comparison of observed physical traits. Our approach similarly tackles the comparison of phenotypes, but uniquely enables experts to compare how these traits evolve over time, in this case the course of a child's development. As research in pregnancy health and in particular HG is ongoing, there are currently few visualization methods that support research in this particular topic. A recently developed visualization approach supporting understanding of maternal health is the MoBa Explorer [BVG23], designed for both domain experts and lay audiences that focuses on analysis of genotypic and phenotypic relationships between pregnancy health and child outcomes. In our case, we focus only on a domain expert target user group to support their exploratory goals with these data.

3. Proposed Approach

We followed a design study methodology [SMM12] that incorporated a participatory design approach [JKKS20], using Munzner's nested model for visualization and validation [Mun09]. In close collaboration with the domain expert, a bioinformatician who is a co-author of this work, we engaged in iterative feedback loops throughout the design process to establish the scope and requirements for the dashboard, and for just-in-time validation of design decisions against our identified domain tasks and needs.

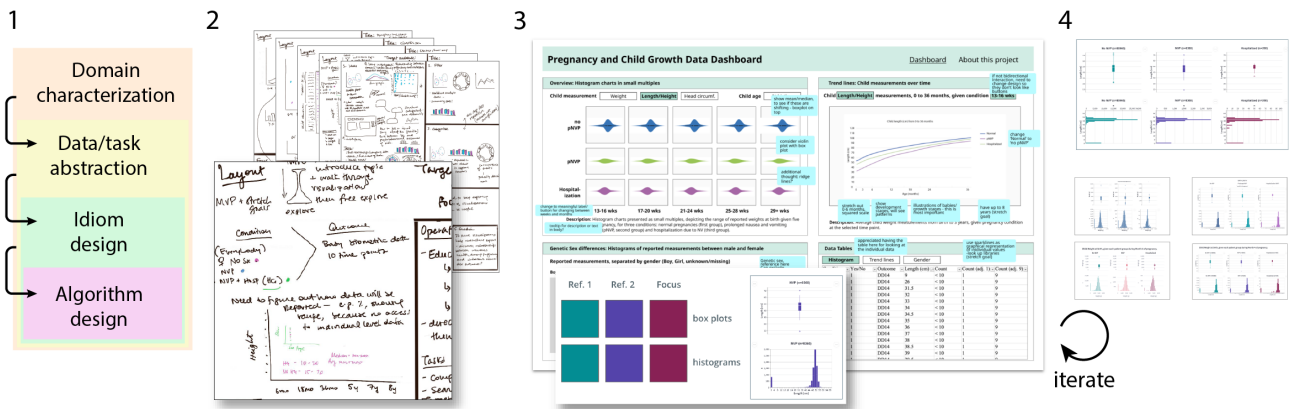


Figure 2: Steps in our design and research process. 1. Nested model for visualization and validation approach [BM13]. 2. FdS method [RHR15]. 3. Mock-up in Figma with domain expert feedback annotated in blue, implemented iteratively. 4. Exploring visualizations and colour schemes through rapid iteration. At and between each step, we engaged in participatory design with the domain expert.

Domain problem characterization: To understand what research questions and hypotheses exist in clinical research on HG and child outcomes, we iterated over several sessions with a domain expert to determine the specific variables to be explored from the dataset. Experts working in this space are interested in understanding the impact of HG on child growth and development outcomes, with height and weight as a key focus.

Data/task abstraction: MoBa phenotypic information are heterogeneous (numerical and categorical variables) tabular data, collected from multiple time points (i.e. questionnaires throughout pregnancy and after birth). Our task analysis resulted in three primary questions and tasks within and across three patient groups (*No NVP*, *NVP*, and *Hospitalized*), abstracted according to Brehmer & Munzner’s multilevel task typology [BM13].

T1 Is there a difference between reported NVP at any of five time points and child growth measurements at one time point (e.g., birth weight)? (*Explore and compare attribute distribution, extrema, across groups*).

T2 Is there a difference in child development between boys and girls? (*Explore and compare data attribute similarities across groups*).

T3 Is there a difference between reported NVP at five different time points and child growth measurements trend over time (e.g., data for length/height from birth up to eight years of age)? (*Explore and compare trends*).

Encoding/interaction technique: Our process to develop appropriate encodings and interactions for the visualization was guided by the Five Design Sheets (FdS) Methodology [RHR15] for rapid idea generation and iteration. FdS is a validated method for brainstorming and prototyping new designs and applications in computer science. The first two authors completed this concept generation exercise independently to maximize the number of ideas, including potential layouts and idioms, to explore in the prototyping phase. Given the expert is interested in attribute exploration and comparison across different time granularities, histograms, box plots,

and line charts were deemed the most appropriate visual idioms to represent distribution of the data that use position and length as the main channels for comparison [Mun14]. Box plots are appropriate for summary statistics [ACG14; KA14] and are familiar to the domain. Similarly, histograms were deemed the ideal idiom for displaying distributions of aggregated data in this small multiple view [Mun14]. Violin plots were considered as another potential idiom; however, since we were working with binned data, this visualization could give the appearance of continuous data. Line charts are commonly used to represent child development over time in clinical practice and research [Jül11]. Finally, the visualizations are interactive and linked between different views to facilitate an overview first then details-on-demand [Shn03] exploration process. We developed an initial visual (non-interactive) mock-up using Figma [Far23] to illustrate the dashboard interface, with an exploration of potential idioms as well as design considerations that make use of accessible design recommendations, including colour and layout. The iterative design process involved several rounds of feedback with the domain expert partner as a preliminary validation of design decisions and to ensure user task alignment.

Algorithm design: Our computations to aggregate and display the data are designed for efficient back-end data management and pre-processing, particularly with a large dataset, to minimize computational time and delay for the user on the front-end. All computation times were measured to milliseconds or faster.

4. The MoBa Pregnancy and Child Development Dashboard

Our iterative design and validation approach yielded a web-based application to allow domain experts to explore possible phenotypic correlations between women’s health during pregnancy (*Hospitalized*, *NVP*, and *no NVP* groups) and children’s developmental outcomes, initially height and weight. The dashboard was developed using Vega-Altair [VGH*18] and Streamlit [TTK18]. The tool and additional details regarding its development, is open-sourced and available at [10.17605/OSF.IO/U6KDM](https://doi.org/10.17605/OSF.IO/U6KDM) along with a tutorial video.

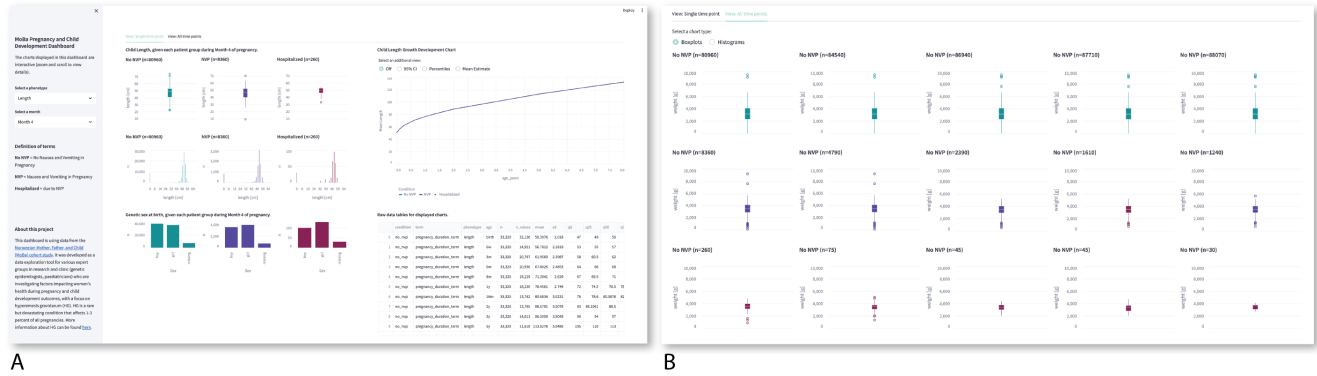


Figure 3: The MoBa prototype: overview of conditions and outcomes data (left) and small multiples view of all time points (right).

The application contains a side panel and two pages of a dashboard, described below via a case study of the expected workflow resulting from a paired session with a domain expert (Fig. 3).

Introduction and variable selection. The side panel displays the title and instruction about the interactions available to the user. Brief information about the data and study are included underneath the variable selection menus, should the user wish to refresh on background information. The user may then select a phenotype and time point of interest, e.g. weight and Month 5. The relevant data will display in the charts in the main dashboard area.

Explore and compare within time point. The user proceeds to the main dashboard which provides an overview and enables deeper investigation into the selected phenotype of children over the three maternal groups of interest. Beginning on the single time point view (Fig. 3A), the user may compare child weights at five months old across the three maternal groups in a summary view (box plots) and investigate extreme values and overall distribution in the histogram view below (T1). The user may also compare between the genetic sexes (T2), with a third ‘unknown’ (unreported) category.

Explore and compare across all time points. The user may also investigate child growth over time in the line chart at the right, separated by maternal condition (T3). Zooming into the line chart shows differences between the three maternal groups (T3). The option to display confidence intervals, percentiles, and mean estimates serving different expert groups was requested by the domain expert to exemplify data uncertainty. An additional tab allows the expert to make visual comparisons of summarized, e.g., weight development within and across maternal groups over all time points (T3) (Fig. 3B) via a 5x3 grid with the time points along the columns and the three patient groups along the rows.

Lastly, the user may conduct a detailed investigation of all summary statistics for all available data via a look-up table at the bottom right of Fig. 3A, a feature that was deemed as useful by the domain expert to capture uncertainty and specific values (T1, T2, T3). The table also serves as an accessibility feature for experts with vision impairment [EBM22].

5. Discussion & Future Work

Our approach incorporates participatory and iterative design with domain experts in the context of a design study methodology. Em-

ploying the FdS method with high fidelity prototyping, we were able to validate our visual and interaction design with a domain expert, culminating in a web-based tool for comparison and analysis of maternal health and child development, with a focus on HG. Next, we plan to validate the tool in a controlled user study to assess the efficacy of our approach in practice. Currently, the prototype includes data for only two child development phenotypes. The dashboard may be expanded with additional phenotypes of interest, including head circumference, BMI, and other measurements such as psychobehavioural outcomes. To support a more integrated research approach in this domain, the dashboard can be linked with other related tools (e.g., the MoBa Explorer [BVG23]) for a more comprehensive ecosystem for longitudinal genotypic and phenotypic cohort studies. Such efforts may also incorporate deep learning models that can run necessary data pre-processing and clean-up steps, as well as generate and support multiple-view visualizations, similar to work conducted by Wu et al. 2021 [WWZ*21].

The required data aggregation step posed a challenge for processing and visualizing non-zero values of less than 10. Decisions about how to assign values—and how to indicate this step was performed—can cause downstream effects in the (mis)representation and (mis)interpretation of these data [BH84]. Further investigation into the consequences of such data processing changes is necessary. Related to this, raw data from the MoBa cohort study needs to be more reliable, as this issue is ongoing in most clinical datasets. Future work to incorporate explicit depiction of missing data, drawing inspiration from approaches such as The Missingness Glyphs (MissiG) [FW21] may enable deeper understanding of the “patchiness” and subsequent reliability and representation of certain patient profiles in the data. This is a significant challenge in our context because study participants in the patient group of interest (i.e. women diagnosed with HG) are known to have higher attrition rates—later questionnaires are less likely to be filled out due to urgent health issues other barriers to participation—which can introduce further bias and uncertainty in the data. This hinders research and educational efforts about the disease for clinicians and patients who could benefit from such resources. Future visualization research to surface such issues would be invaluable to the medical research community to better identify and account for such biases in the data. Continued research into the questions we explore in this work is an opportunity to engage in

visualization as a method to enhance clinical research for maternal health, and promote patient education and advocacy as a result.

Acknowledgments: The Norwegian Mother, Father and Child Cohort Study is supported by the Norwegian Ministry of Health and Care Services and Ministry of Education and Research. We are grateful to all the participating families in Norway who take part in this ongoing cohort study. This work is funded by Univ. of Bergen and Trond Mohn Foundation (#813558, Visualizing Data Science for Large Scale Hypothesis Management in Imaging Biomarker Discovery). The authors wish to thank the reviewers for their valuable feedback on this manuscript.

References

- [ACG14] ALBERS, DANIELLE, CORRELL, MICHAEL, and GLEICHER, MICHAEL. “Task-driven evaluation of aggregation in time series visualization”. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. 2014, 551–560. DOI: [10.1145/2556288.2557200](https://doi.org/10.1145/2556288.2557200) 2, 3.
- [AMST11] AIGNER, WOLFGANG, MIKSCH, SILVIA, SCHUMANN, HEIDRUN, and TOMINSKI, CHRISTIAN. *Visualization of Time-Oriented Data*. Vol. 4. Springer, 2011. DOI: [10.1007/978-0-85729-079-3_2](https://doi.org/10.1007/978-0-85729-079-3_2).
- [BH84] BABAD, YAIR M and HOFFER, JEFFREY A. “Even no data has a value”. *Communications of the ACM* 27.8 (1984), 748–756. DOI: [10.1145/358198.358204](https://doi.org/10.1145/358198.358204).
- [BM13] BREHMER, MATTHEW and MUNZNER, TAMARA. “A Multi-Level Typology of Abstract Visualization Tasks”. *IEEE Transactions on Visualization and Computer Graphics* 19.12 (2013), 2376–2385. DOI: [10.1109/TVCG.2013.1243](https://doi.org/10.1109/TVCG.2013.1243).
- [BVG23] BALAKA, H, VALEN, R, GARRISON, LA, and VAUDEL, M. “MoBa Explorer: Enabling the Navigation of Data from the Norwegian Mother, Father and Child Cohort Study (MoBa)”. (2023) 2, 4.
- [EBM22] ELAVSKY, FRANK, BENNETT, CYNTHIA, and MORITZ, DOMINIK. “How accessible is my visualization? Evaluating visualization accessibility with Chartability”. *Computer Graphics Forum* 41.3 (2022), 57–70. DOI: [10.1111/cgf.14522](https://doi.org/10.1111/cgf.14522) 4.
- [Far23] FARIA, PEDRO. *figma: Web Client/Wrapper to the 'Figma API'*. <https://github.com/pedropark99/figma>, <https://pedropark99.github.io/figma/>. 2023 3.
- [FTG*19] FEJZO, MARLENA S, TROVIK, JONE, GROOTEN, IRIS J, et al. “Nausea and vomiting of pregnancy and hyperemesis gravidarum”. *Nature reviews Disease primers* 5.1 (2019), 62. DOI: [10.3389/fmed.2021.809270](https://doi.org/10.3389/fmed.2021.809270).
- [FW21] FERNSTAD, SARA JOHANSSON and WESTBERG, JIMMY JOHANSSON. “To Explore What Isn’t There—Glyph-Based Visualization for Analysis of Missing Values”. *IEEE Transactions on Visualization and Computer Graphics* 28.10 (2021), 3513–3529. DOI: [10.1109/TVCG.2021.3097407](https://doi.org/10.1109/TVCG.2021.3097407) 4.
- [GGC*16] GLUECK, MICHAEL, GVOZDIK, ALINA, CHEVALIER, FANNY, et al. “PhenoStacks: cross-sectional cohort phenotype comparison visualizations”. *IEEE transactions on visualization and computer graphics* 23.1 (2016), 191–200. DOI: [10.1109/TVCG.2016.2598469](https://doi.org/10.1109/TVCG.2016.2598469) 2.
- [GHC*16] GLUECK, MICHAEL, HAMILTON, PETER, CHEVALIER, FANNY, et al. “PhenoBlocks: Phenotype Comparison Visualizations”. *IEEE Transactions on Visualization and Computer Graphics* 22.1 (2016), 101–110. DOI: [10.1109/TVCG.2015.2467733](https://doi.org/10.1109/TVCG.2015.2467733) 2.
- [JKKS20] JÄNICKE, STEFAN, KAUR, PAWANDEEP, KUZMICKI, PAWEL, and SCHMIDT, JOHANNA. “Participatory Visualization Design as an Approach to Minimize the Gap between Research and Application.” *Vis-Gap@ Eurographics/EuroVis*. 2020, 35–42. DOI: [10.2312/visgap.20201108](https://doi.org/10.2312/visgap.20201108) 2.
- [JME10] JAVED, WAQAS, MCDONNELL, BRYAN, and ELMQVIST, NIKLAS. “Graphical perception of multiple time series”. *IEEE transactions on visualization and computer graphics* 16.6 (2010), 927–934. DOI: [10.1109/TVCG.2010.1622](https://doi.org/10.1109/TVCG.2010.1622).
- [Jú11] JÚLIÚSSON, PETUR. *About the Bergen Growth Study*. 2011. URL: <https://www.vekststudien.no/research-project/about-the-bergen-growth-study/> 3.
- [KA14] KRZYWINSKI, MARTIN and ALTMAN, NAOMI. “Points of significance: visualizing samples with box plots.” *Nature methods* 11.2 (2014). DOI: [10.1038/nmeth.2813](https://doi.org/10.1038/nmeth.2813) 3.
- [MBV*16] MAGNUS, PER, BIRKE, CHARLOTTE, VEJRUP, KRISTINE, et al. “Cohort profile update: the Norwegian mother and child cohort study (MoBa)”. *International journal of epidemiology* 45.2 (2016), 382–388. DOI: [10.1093/ije/dyw029](https://doi.org/10.1093/ije/dyw029) 1, 2.
- [MDM*15] MALIK, SANA, DU, FAN, MONROE, MEGAN, et al. “Cohort comparison of event sequences with balanced integration of visual analytics and statistics”. *Proceedings of the 20th International Conference on Intelligent User Interfaces*. 2015, 38–49. DOI: [10.1145/2678025.2701407](https://doi.org/10.1145/2678025.2701407) 2.
- [Mun09] MUNZNER, TAMARA. “A nested model for visualization design and validation”. *IEEE transactions on visualization and computer graphics* 15.6 (2009), 921–928. DOI: [10.1109/TVCG.2009.1112](https://doi.org/10.1109/TVCG.2009.1112).
- [Mun14] MUNZNER, TAMARA. *Visualization analysis and design*. CRC Press, 2014. DOI: [10.1201/b175113](https://doi.org/10.1201/b175113) 3.
- [NHG19] NUSRAT, SABRINA, HARBIG, THERESA, and GEHLENBORG, NILS. “Tasks, techniques, and tools for genomic data visualization”. *Computer Graphics Forum*. Vol. 38. 3. Wiley Online Library, 2019, 781–805. DOI: [10.1111/cgf.13727](https://doi.org/10.1111/cgf.13727) 2.
- [PL20] PREIM, BERNHARD and LAWONN, KAI. “A survey of visual analytics for public health”. *Computer Graphics Forum*. Vol. 39. 1. Wiley Online Library, 2020, 543–580. DOI: [10.1111/cgf.13891](https://doi.org/10.1111/cgf.13891) 2.
- [RHR15] ROBERTS, JONATHAN C, HEADLEAND, CHRIS, and RITSOS, PANAGIOTIS D. “Sketching designs using the five design-sheet methodology”. *IEEE transactions on visualization and computer graphics* 22.1 (2015), 419–428. DOI: [10.1109/TVCG.2015.2467271](https://doi.org/10.1109/TVCG.2015.2467271) 2, 3.
- [Shn03] SHNEIDERMAN, BEN. “The eyes have it: A task by data type taxonomy for information visualizations”. *The craft of information visualization*. Elsevier, 2003, 364–371. DOI: [10.1016/B978-155860915-0/50046-9](https://doi.org/10.1016/B978-155860915-0/50046-9) 3.
- [SMM12] SEDLMAIR, MICHAEL, MEYER, MIRIAH, and MUNZNER, TAMARA. “Design study methodology: Reflections from the trenches and the stacks”. *IEEE transactions on visualization and computer graphics* 18.12 (2012), 2431–2440. DOI: [10.1109/TVCG.2012.2132](https://doi.org/10.1109/TVCG.2012.2132) 2.
- [TTK18] TREUILLE, ADRIEN, TEIXEIRA, THIAGO, and KELLY, AMANDA. *Streamlit App*. <https://www.streamlit.io/>. Accessed: (2024-03-15). 2018 1, 3.
- [VGH*18] VANDERPLAS, JACOB, GRANGER, BRIAN, HEER, JEFFREY, et al. “Altair: Interactive statistical visualizations for python”. *Journal of open source software* 3.32 (2018), 1057. DOI: [10.21105/joss.01057](https://doi.org/10.21105/joss.01057) 1, 3.
- [WT08] WOJCZYNSKI, MARY K and TIWARI, HEMANT K. “Definition of phenotype”. *Advances in genetics* 60 (2008), 75–105. DOI: [10.1016/S0065-2660\(07\)00404-X](https://doi.org/10.1016/S0065-2660(07)00404-X) 2.
- [WWZ*21] WU, AOYU, WANG, YUN, ZHOU, MENGYU, et al. “Multi-Vision: Designing analytical dashboards with deep learning based recommendation”. *IEEE Transactions on Visualization and Computer Graphics* 28.1 (2021), 162–172. DOI: [10.1109/TVCG.2021.3114826](https://doi.org/10.1109/TVCG.2021.3114826) 4.
- [ZS17] ZAKKAR, MOUTASEM and SEDIG, KAMRAN. “Interactive visualization of public health indicators to support policymaking: An exploratory study”. *Online journal of public health informatics* 9.2 (2017). DOI: [10.5210/objphi.v9i2.8000](https://doi.org/10.5210/objphi.v9i2.8000) 2.