Eurographics Workshop on Visual Computing for Biology and Medicine (VCBM) (2026)
Robert Krueger, Eric Mörth, and Katarína Furmanová (Editors)

Conference Paper

Voidlytics for leveraging informative missingness in the study of hyperemesis gravidarum

H. Balaka[1,3], K. E. A. Zhang[1], M. M. Czinkoczky[2], M. Vaudel[2], H. Hauser[1] and L. A. Garrison[1]

[1]Department of Informatics, University of Bergen, Norway
[2]Department of Clinical Science, University of Bergen, Norway
[3]Norwegian Institute of Public Health, Department of Genetics and Bioinformatics, Norway

Abstract

Public health studies commonly use questionnaires to get insights into human health. The Norwegian Mother, Father and Child Cohort Study (MoBa) is a longitudinal study that captures self-reported data on pregnancy symptoms across gestation, in particular also with a potential for an improved understanding of hyperemesis gravidarum (HG), a severe pregnancy complication. Although missing values are prevalent in these data and can hinder the analysis of the collected information, researchers studying MoBa data hypothesize that missingness within observed responses may signal HG. We present a design study exploring the problem space of hypothesis generation regarding relationships between missingness response patterns in the MoBa data and HG. Through this study, we also introduce Voidlytics, a visualization tool that supports hypothesis generation. It facilitates the exploration of missingness response patterns, the explicit formulation of assumptions about how suggestive of HG specific patterns are, as well as testing these assumptions. We provide initial evidence that the tool may support the generation of hypotheses about how indicative of HG missingness response patterns are. Supplementary materials are available at `osf.io/3ycp5`.

Categories and Subject Descriptors (according to ACM CCS): H.5.m [Information interfaces and presentation (e.g., HCI)]: User Interfaces—Miscellaneous

1. Introduction

Studies in public health often use questionnaires to understand population health. Questionnaires collect data on many aspects of health including lifestyle, diet, and existing health conditions, all of which provide valuable insights into the mechanisms underlying complex health conditions. The Norwegian Mother, Father and Child Cohort Study (MoBa) [MBV*16] is one such study that collects longitudinal, self-reported data on pregnancy symptoms at multiple time points across gestation. Bioinformatics researchers working with MoBa data aim to link self-reported pregnancy symptoms to the underlying genetic architecture, surfacing the genetic factors that influence severe pregnancy complications such as hyperemesis gravidarum (HG) [FSS*18, FWZ*24] as well as the broader spectrum of nausea and vomiting in pregnancy.

However, the analysis of questionnaire data is challenged by missingness that can arise from a variety of factors on the participant's end (e.g., participants choosing not to respond or being unable to respond due to health-related issues), or due to a design limitation in a questionnaire (e.g., early MoBa questionnaires recorded only affirmative responses while leaving negative ones undocumented) [KSA*01].

Previous analyses of MoBa data [VSMG13, BHS*09] typically take a complete-case analysis approach that excludes cases with incomplete responses and performs statistical analyses on complete cases. This approach entails a risk of selection bias as it assumes missingness to be statistically random, when missingness in symptom questionnaires is rarely random [VMM22, ZvKC*24]. Women who fail to complete questionnaires due to severe nausea, hospitalization, or pregnancy complications are systematically different from those who return their responses, meaning complete-case analysis disproportionately excludes the most affected individuals. As a result, conclusions drawn from complete-case analysis reflect the symptom experience of healthier, more compliant responders rather than the full population.

The current analysis of pregnancy symptom data in rela-

tion to HG could be improved by methods to explore and analyze missingness in MoBa data. To this end, we present a design study to characterize the problem space of hypothesis generation concerning associations between missingness in the data and HG. Our study also aims to explore visual solutions to support researchers in generating hypotheses about missingness within observed data and HG. This is a collaborative effort between visualization researchers and bioinformatics researchers (domain experts, authors #3 and #4) from the Dept. of Clinical Science who work with MoBa to study HG. Our main contributions include:

- data and task characterization for HG missingness analysis, guided by domain expert goals
- iterative co-design of Voidlytics, a novel tool for HG analysis, that leverages missingness patterns along with the analyst's tacit knowledge about the missing data
- an illustration of how Voidlytics may support generation of hypotheses about patterns of missing responses and HG

Voidlytics is the first tool, in the best of our knowledge, that attempts to visualize and manage the complexity of multidimensional missingness while integrating the analyst's assumptions into missingness analysis.

2. Related work

We review prior research on informative missingness. We describe also the assumptions of missing data in health cohort studies and visual strategies for sensemaking with missing data. The current landscape of HG studies is described in the following section on domain background and context.

2.1. Assumptions of missing data in cohort studies

Missing / incomplete data is inherent in longitudinal health cohort studies that often source data from self-reported questionnaires in combination with electronic health records (EHR) [WCNK13, BHF15, ZvKC*24]. This missingness is seen as a source of bias that compromises the reliability of statistical inferences as well as the generalizability of results for a target patient population. When handling missing data, analysts need to make assumptions about the mechanism behind the missing responses [HAD21, ZvKC*24, Rub76]. While missingness can be assumed to be statistically random–completely at random (MCAR) or at random (MAR)–with consequences only in a loss of statistical power, missingness in cohort studies is often systematic and nonrandom in reality. In the latter, responses can be assumed as missing not at random (MNAR) with a dependency on unobserved data. Studies with self-reported data may deal with selection bias, where patients experiencing adverse health effects are more likely to drop out of a study and are fundamentally different from observed cases [VMM22].

As missingness mechanisms are often not testable using data alone [SG02], the analysis results are dependent on the analysts' assumptions about the missing data. Concerns have been raised about the ambiguous reporting of assumptions [LTC*21] when making methodological decisions about missing data as it can lead to the misuse of common data science approaches for working around missing data, e.g., complete-case analysis or data imputation. To this end, several works call for more explicit ways to articulate and scrutinize assumptions and other domain knowledge for missing data analysis [ZvKC*24, HAD21, LWN19]. Another approach (to investigating missingness mechanisms) considers the missingness itself as informative, analyzing the extent and structure of missing values within a dataset. For example, Tan et al. [TGH*23] examine patterns of missing EHR laboratory test data to support hypothesis generation and inferential analysis of hospital treatment capacity. Our work aims to explore and combine these alternative orientations to missing data analysis, working with (and not around) informative missingness using domain-specific assumptions.

2.2. Visual analysis of missing data

Visualizations can help analysts to make sense of missing data, from revealing the extent of missingness within a dataset, to the discovery and interpretation of trends and patterns, to reconciling the uncertainty and assumptions arising from data analysis [SMW*22, SSS*14, EPD05, AFF24]. This is an understudied area of visualization research as missingness mechanisms are challenging both to identify and to translate to data analysis tasks [Fer19]. Within existing literature, a first strategy for missingness analysis is to visualize the structure and distribution of missing data across different data items and dimensions so as to contextualize the missing data [SWB*24]. Several packages for missing data, e.g., Missingno [Bil18], Naniar [TC23], VIM [KT16], offer standard statistical charts such as bar charts, line charts, scatterplots, mosaic plots, and heatmaps to provide an overview of null values as well as to indicate if missing data is concentrated in specific variables or cases. Another strategy is to focus on revealing the relationships and groupings of missingness between data dimensions, such as with interactive set visualizations [RAH22], parallel coordinate plots [Fer19], glyphs [FW21], clustering algorithms [JM22], and coordinated multiple views [AFDD*24, ANI*20, ST17]. A handful of works describe visual tools and encodings for interpreting imputed missing values, e.g., by Song et al. [SS18], but these are not discussed as the scope of this work is on the analysis of missingness without imputation.

Fewer methods focus on helping the analyst understand missingness as it relates to their assumptions and tacit knowledge. This presents a gap to explore, given the important role that an analyst's prior knowledge plays in interpreting and making judgments about graphs with missing values [EPD05, SWB*24]. McCurdy et al. [MGM18] allow experts to record and externalize their uncertainties about observed data using glyphs and text-based annota-

tions, while Franke et al. [FBKW19] allow for a similar task using qualitative labels. Drawing upon these works, Lin et al. [LAML22] characterize a set of conceptual techniques to record an analyst's tacit knowledge or "hunches" about the imperfections of their data, though we have not yet found an example in use in missingness visualization nor for our domain-specific context. We see this as an opportunity to explore toolmaking that can leverage an analyst's tacit knowledge and assumptions in the visual analysis of missing data.

3. Domain background and context

Hyperemesis gravidarum (HG) is a severe pregnancy complication characterized by persistent, intractable nausea and vomiting, frequently requiring hospitalization and affecting approximately 0.3–3% of all pregnancies [MOR*16]. Despite its substantial impact on maternal quality of life and pregnancy outcomes, the underlying mechanisms of HG remain poorly understood, and reliable genetic risk factors have yet to be established. Longitudinal cohort data are essential for the study of HG and the broader spectrum of nausea and vomiting in pregnancy, The Norwegian Mother, Father and Child Cohort Study (MoBa) [MBV*16] is a uniquely rich resource for this purpose, providing a collection of detailed self-reported data on pregnancy symptoms and maternal health across multiple time points during gestation. The typical analysis workflow for MoBa data involves pre-processing and data cleaning followed by exploratory and confirmatory statistical techniques. Factor analysis is one such technique used to examine the relationships between reported symptoms and identify latent variables in the data. These latent factors are subsequently used for downstream interpretation and genetic association analyses. All data exploration and analysis is conducted within HUNT Cloud [HUN26], a secure cloud-based environment for the processing of sensitive research data in Jupyter notebooks.

A key challenge to this workflow is the prevalence and ambiguity of missingness in the MoBa data. As an artifact of early MoBa questionnaire design, a missing response may indicate either the genuine absence of a symptom or that the participant did not complete the question – two different informational states that are indistinguishable without additional context. The study includes metadata on questionnaire return, recording which questionnaires were not returned by the participant. When a questionnaire is not returned, researchers treat the corresponding responses as missing and typically exclude them from analysis. In contrast, we hypothesize that patterns of attrition in the data may be informative of HG. Supporting the exploration and formulation of such hypotheses is a central motivation of this design study.

4. Data, Tasks, and Design Requirements

In our study, we use the design study methodology proposed by Sedlmair et al. [SMM12]. We follow Munzner's nested model for visualization design [Mun09] and begin with the characterization and abstraction of data and tasks.

4.1. Data

The MoBa study includes pregnancy health data from over 100,000 pregnancies across more than 90,000 mothers, recruited between 1999 and 2008. Follow-up data collection from enrolled participants continues to the present. Data are collected via three self-report questionnaires (*Q1-3*) sent out at the gestational weeks 15, 22, and 30. *Q1* includes questions on mental and physical health during the first trimester and prior to conception, *Q2* focuses primarily on diet and nutrition during pregnancy, and *Q3* addresses mental and physical health in the later stages of pregnancy. Our work is dedicated to the visualization of binary questionnaire variables indicative of HG symptoms and outcomes, specifically the prevalence of nausea, vomiting, and the hospitalization of mothers during pregnancy.

The data on these symptoms are represented as a table with the corresponding questions as its attributes and pregnancy instances as its items. *Q1* and *Q2* include questions on nausea and vomiting, but not hospitalization. *Q3* contains questions about nausea, long-term nausea and vomiting, and hospitalization. Each question variable is further split by weeks of pregnancy, e.g, "before week four", "between weeks 5 and 8", or "after week 13". *Q1* probes into the first trimester, while *Q3* the second. However, for the question on hospitalization due to long-term nausea and vomiting, *Q3* considers weeks across the entire pregnancy. The binary questions on nausea and vomiting in *Q2* do not capture the temporal data on these symptoms, but rather whether these symptoms have occurred during the pregnancy.

We extended the questionnaire data with variables from the birth registry, including quantitative measures such as placental weight and pregnancy duration, which our domain collaborators (and co-authors) identified as potentially informative correlates of HG severity. The resulting dataset is a mixed-type multidimensional table combining binary symptom indicators and continuous registry-derived variables.

4.2. Expert tasks

We used contextual inquiry methods [BH99] to elicit and characterize a domain expert's analysis goals for missing data. We started by conducting semi-structured, one-on-one interviews with our domain experts (authors #3 and #4) who work daily with the MoBa questionnaire data. Each interview lasted 60 mins. and was conducted by paired interviewers [AM23]. The researcher was asked to walk through their typical analysis workflow while thinking aloud and describing their actions. The interview also included open-ended questions probing into their analytical practices regarding missing data, such as questions about the treatment of missingness as well as the assumptions made during analysis.

The complete set of interview questions is available in supp. materials at osf.io/3ycp5. Following the interviews, the visualization researchers met in weekly sessions to synthesize the domain goals elicited from the interview findings into a set of high-level expert tasks that can be used to derive design requirements. In tandem, the first author sat weekly with the researchers in their workspace to iteratively refine our characterization of their domain goals.

The outcome of this contextual design process is the co-design of Voidlytics. The tool supports two main domain goals: 1) the domain-specific exploration of missingness patterns in questionnaire data, and 2) hypothesis generation for HG by using the expert's tacit knowledge and assumptions about their missing data. Based on these domain goals, we isolated the high-level expert tasks below following Brehmer and Munzner's multi-level task typology [BM13]:

T1: Discover and explore patterns of missing data. The analyst begins exploratory data analysis (EDA) by examining the relationships between a set of questionnaire variables indicative of HG symptoms and outcomes. The analyst is also interested in exploring the distribution of missing responses for each variable. While they assume HG participants do not fully complete the questionnaires i.e., attrition due to disease, they cannot use current methods to easily explore possible HG cases within the missing data.

T2: Derive an "HG certainty variable" using missingness patterns and tacit knowledge. Analysts often use their tacit knowledge, e.g., assumptions and intuitions about the data and the disease, to "gut check" and inform their analysis. To leverage the analyst's tacit assumptions about missingness patterns **[T1]**, we need to enable the analysts to record and represent their (un)certainty that a response pattern is characteristic of participants with HG.

T3: Compare the derived variable with other variables of interest. Typical EDA involves examining the relationship between different questionnaire variables **[T1]**. By comparing a derived HG variable **[T2]** with other variables of HG outcomes, the analyst can test whether missingness patterns represent possible cases of HG. For example, the analyst can be more confident in their characterization of HG if placental weight is different in missingness patterns defined as "high" vs. "low" HG certainty.

4.3. Design Requirements

We derived a set of design requirements that Voidlytics should fulfill to address the high-level expert tasks above.

R1: Show patterns of missing responses (missingness patterns) and their distribution for each variable of interest (nausea, vomiting, and hospitalization) in *Q1-3* across the duration of pregnancy **[T1]**.

R2: Enable specifying a missingness pattern **[T2]**.

R3: Enable quantifying the user's certainty that a selected missingness pattern is characteristic of HG **[T2]**.

R4: Enable deriving an HG certainty variable based on a set of missingness patterns and HG certainty values specified by the user **[T2]**.

R5: Enable the visual exploration and analysis of the derived HG certainty variable along with other variables of interest from the cohort study or birth registry **[T3]**.

R6: Enable the visual exploration and analysis in the common Python notebook environments.

An important design consideration is that the analyst can only access the questionnaire data in a HUNT Cloud [HUN26], a secured academic cloud environment and digital lab space, which supports interactive computing tools built around Jupyter. Specifically, Jupyter notebooks are used for exploratory data analysis. Accordingly, we needed to account for the requirement to work within this secured environment **[R6]**.

5. Visualization design

Voidlytics comes with three views, **Exploration**, **Definition**, and **Analysis**, that enable the analyst to explore missingness response patterns, generate assumptions on their suggestiveness of HG, and self-quantify and test the assumptions. In addressing the design requirements, we developed this tool as an interactive widget that works across Python notebook environments **[R6]**. Inspired by DimBridge [MAR*24], we used anywidget [MGA24] and combined Python's backend ecosystem with D3 [BOH11] for the frontend. All source code is available in a GitHub repository: github.com/Annchovy/voidlytics.

5.1. Exploration view

Exploring missingness in the participants' responses is a necessary starting point for generating hypotheses on how indicative of HG such missingness is. Illustrated in Fig. 1, we provide the **Exploration view** for examining missingness response patterns in the data. Since the experts are interested in exploring missingness, we disregard response sets with complete data as irrelevant for our missingness-based hypothesis generation. While excluding complete cases may introduce bias, we see this missingness-oriented strategy as complementary to the current analytical approaches for observed (non-missing) data. The remaining subset contains 24,876 responses. To manage this amount of data, we aggregate these responses into 2,459 unique missingness response patterns. Drawing inspiration from a missingness map employed in VIVID [ANI*20], we create an interactive heatmap to display these patterns **[R1]**, as shown in Fig. 1A. The heatmap allows for a compact yet expressive representation of responses across the questions and patterns.

The heatmap's *x*-axis lists the dataset's variables, i.e., the questions in *Q1-3*. We order the questions in three levels: by the questionnaires (*Q1*, *Q2*, *Q3*), by the symptoms in the questions (nausea, vomiting, hospitalization), and by the

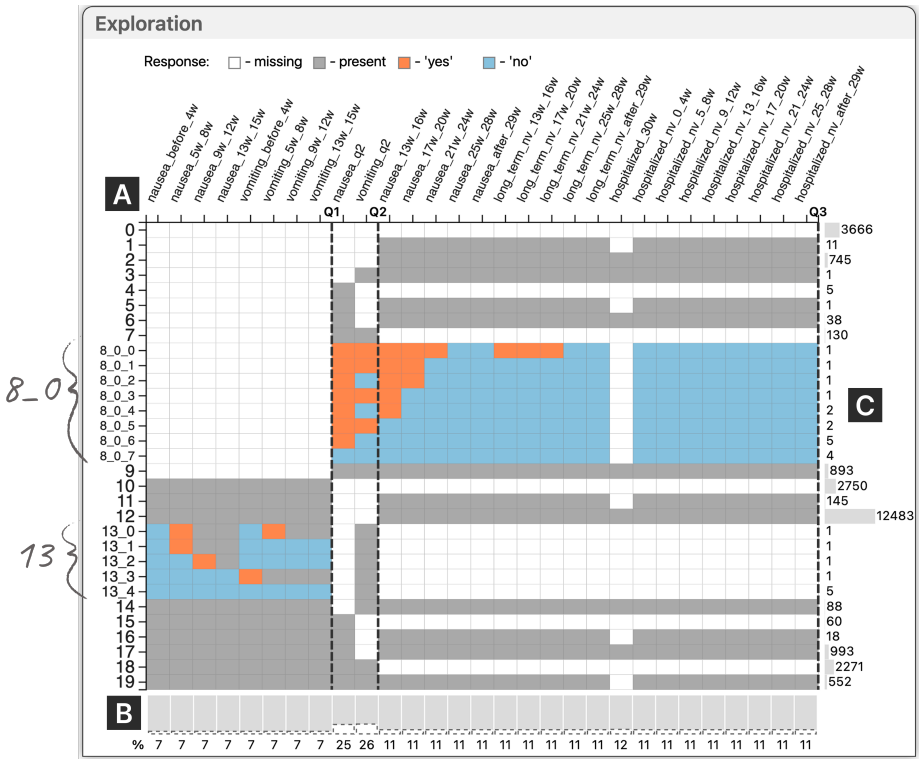

Figure 1: *The Exploration view is an entry point to exploring missing responses. (A) The heatmap displays missingness response patterns. Here, pattern 13 is unfolded to the first level of subpatterns, while pattern 8 is expanded to subpattern 8_0 (the only subpattern in 8) and then further to the second-level subpatterns (8_0_1, 8_0_2, etc.). Affirmative responses (symptom experienced) are encoded in orange, while negative responses (symptom not experienced) in blue. Pattern 13_2 shows the nausea onset between weeks 9 & 12 and no vomiting in Q1. (C) The bottom bar chart shows missingness extent per question.*

symptoms' time points (nausea before week 4, nausea between weeks 5 & 8, etc.). We choose this ordering to match the experts' mental map of the questionnaires' structure. The *y*-axis corresponds to all unique missingness response patterns emerging from the data.

Initially, we group participants into patterns based on whether their responses are missing or observed across all questions, without considering differences in the observed response values. For example, participants with responses to *Q1* but without responses to all questions in *Q3* constitute the same pattern. We encode missing responses as white and observed values as grey. This initial coarse aggregation condenses the data into 20 participant groupings, providing an overview of response patterns in the dataset and laying the foundation for further exploration.

We order patterns in the heatmap by missingness across the question sequence, e.g., responses missing (*Q1*, *Q2*, *Q3*) appear first, followed by those missing (*Q1*, *Q2*), (*Q1*, *Q3*), and (*Q1*). This ordering improves the heatmap's structure and readability, and surfaces potential cases of severe early-pregnancy symptoms (i.e., participants felt too sick to return *Q1*), a typical manifestation of HG.

We also summarize the extent of missingness across all questions, as shown in Fig. 1B. We calculate missingness percentages using all (complete and incomplete) responses to situate the missingness patterns in the context of the full dataset. As shown in Fig. 1C, we additionally supplement the heatmap with a bar chart displaying the counts of response sets distributed across the patterns **[R1]**.

For the detailed exploration of missingness patterns, the heatmap offers drill-down interactions **[R1]**. Inspired by (un)foldable data visualization techniques [BDBT25], our tool allows the analyst to click on a pattern and unfold it into corresponding subpatterns of higher granularity.

There are two levels of subpatterns that can be unfolded. At the first level, a clicked pattern unfolds into subpatterns based on nausea and vomiting onsets reported in *Q1*. Here, the onsets are the time points when nausea and vomiting first occurred in the response sequence. For example, responses with nausea onset before week four and those with the onset between weeks 13 and 15 represent different subpatterns.

There are 174 such subpatterns in the data. For each subpattern, we encode responses *"Yes"* (marking the onset) in orange and *"No"* (before the onset) in blue. Fig. 1 shows pattern "13_2" with the nausea onset between weeks nine and 12, whereas for pattern "13_3" nausea did not appear in the first trimester, but vomiting began before week four. Such splitting uses a key characteristic of HG, namely the occurrence of nausea and vomiting in early pregnancy. Disaggregating a pattern based on the symptoms' onsets enables immediate identification of response groupings with no nausea and vomiting in early pregnancy, as well as response groups consistent with the timing of HG-related symptoms.

The analyst can click on an unfolded subpattern and expand it to the second (and lowest) level of subpatterns revealing finer distinction between responses, as shown in Fig. 1 (pattern "8_0"). At this level, the aggregated data is broken down into the unique, individual patterns that expose the participants' exact responses across all questions. When the analyst clicks a lowest-level subpattern, the subpatterns are folded back into the highest-level pattern.

Within each top-level pattern, subpatterns are ordered by severity scores. We calculate severity scores as weighted sums of responses. For each pattern, we construct a binary vector with *"Yes"* as 1 (a participant experienced the symptom) and *"No"* or missing response as 0 across all questions. For each vector, we compute a linear combination of its elements with corresponding symptom weights. When creating the widget, the analyst can specify the symptom weights and, for example, up-weigh hospitalization questions. Since persistent nausea and vomiting are common among women with HG, ordering by such severity scores brings up patterns with persistent or / and serious symptoms. Fig. 1 shows subpatterns of "8_0" ordered by severity scores.

The bar chart showing the distribution of response sets follows (un)folding interactions and adjusts to match the revealed patterns. As the heatmap exposes more patterns, its cells' height decreases and vice versa. If the cell height drops below a readable threshold, the heatmap expands vertically to preserve patterns' legibility. The tool supports interactive panning for the exploration of patterns beyond the visible area. Additionally, Voidlytics offer zooming interactions and hovering over cells to reveal their values.

We also considered visualizing missingness using multiple stacked bar charts, where each chart showed the proportion of missing and observed responses for a single pregnancy symptom across this symptom's time points. This approach would provide a compact summary of missingness in the data and allow the analyst to drill down into specific subsets of participants by brushing and linking. However, this method would primarily support the assessment of missingness proportions, whereas the domain goals are focused on exploring missingness patterns. Revealing such patterns would require the analyst to iteratively select subsets of responses across multiple stacked bar charts. As the number of symptoms and time points grows, the interaction effort increases and the exploratory process becomes more complex. Moreover, revisiting previously identified patterns would require the analyst to repeat the same sequence of interactions or for the approach to provide a mechanism for retracing previous interaction steps. In contrast, the heatmap directly visualizes the missingness response patterns, and requires comparatively little interaction to move from their overview to detail. While the heatmap introduces its own challenges, such as displaying many patterns simultaneously which potentially increases the visual search effort, it provides a more immediate and comprehensive overview of the patterns and their distribution present in the dataset.

Exploration of missing responses allows the analyst to examine their assumptions about missingness response patterns as indicative of HG. To externalize their assumptions, the analyst selects a pattern of interest by clicking on its label in the heatmap and proceeds to the next view **[R2]**.

5.2. Definition view

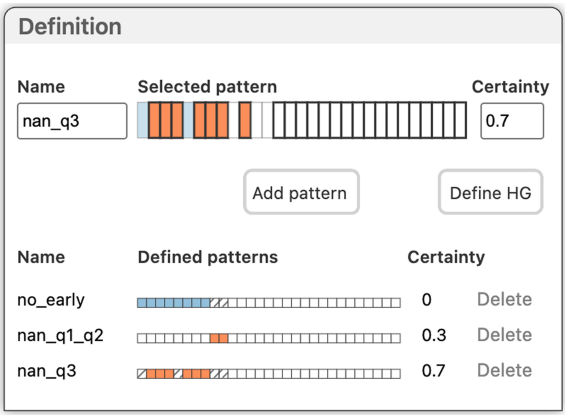

Figure 2: *The Definition view allows the analyst to specify how indicative of HG a missingness pattern is. A selected pattern is automatically imported from the Exploration view into the Definition view, where the analyst can modify it and specify its name and certainty.*

Complementary to the **Exploration view**, the **Definition view** allows the analyst to define a missingness response pattern and quantify their certainty on how indicative of HG this pattern is **[R3]**.

Fig. 2 illustrates the **Definition view** with an example of a pattern selected in the **Exploration view**. Through the inputs "Name" and "Certainty" in the view, the analyst can assign a name and a certainty to a pattern in focus **[R3]**. Additionally, the analyst can modify a selected pattern. The analyst may judge specific segments of the pattern suggestive of HG, while the rest uninformative. By clicking on the cells of a selected pattern, the analyst can include / exclude responses.

For example, in Fig. 2, the questions on nausea and vomiting before week four are deselected, and essentially labeled as "any value or missing". The pattern modification is supported by revealing a cell's question and response when hovered over. By default, the whole pattern is selected.

Once the analyst defined a pattern and its HG certainty, they can press "Add pattern" to include their assumption to the list of defined patterns, shown in Fig. 2. The analyst can revise the defined patterns by deleting and creating new ones. Based on specified patterns, the analyst can press "Define HG" and derive a new HG variable with the defined certainties assigned to matching response sets in the data **[R4]**. For instance, response sets consistent with pattern "nan_q3" in Fig. 2 receive HG certainty 0.7. By default, all response sets have a HG certainty of zero. We also assign a certainty of zero to complete response sets as they provide no missingness-based evidence for HG cases.

The derived HG variable can be accessed through the widget instance and seamlessly attached to the original dataset for use in downstream statistical analyses alongside other variables of interest. This procedure turns assumptions about missingness within observed data into a feature itself.

5.3. Analysis view

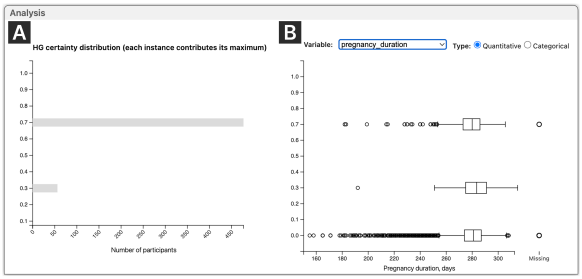

Figure 3: *The Analysis view facilitates integration of the self-quantified HG variable into exploratory data analysis. (A) The histogram displays the distribution of defined HG certainties. (B) The analyst can relate the defined HG variable to another variable of interest, e.g., pregnancy duration.*

Once the analyst has derived an HG variable, the **Analysis view**, shown in Fig. 3, offers visual exploration and testing of the resulting variable. The view consists of two components: a histogram displaying the HG certainty distribution and a plot relating the defined HG variable to a registry variable of interest, e.g., pregnancy duration. The histogram, shown in Fig. 3A, appears immediately after the analyst presses "Define HG" in the Definition view, and displays only certainty values greater than zero. We exclude cases with HG certainty of zero due to their significantly higher frequency – by default, all response sets receive a certainty of zero. The inclusion of such cases would skew the x-axis and reduce the visibility of non-zero certainty values, which are the prime focus of the analysis. We could also use a logarithmic scale to mitigate the effect of zero-certainty cases, but it would complicate the direct comparison of results. It is important to note that one response set can match multiple patterns, and thus multiple certainties. To simplify the visual analysis of the derived HG variable, we consider only the maximum certainty value per response set, bringing up high-confidence HG cases which are of the greatest interest to the experts.

In the right panel, shown in Fig. 3B, the analyst can select an additional variable of interest. For a selected quantitative variable, the view displays a boxplot of the variable for each certainty value. Here, we consider all response sets, including those with HG certainty of zero. Similarly to the HG histogram, for each response set, we reduce the array of certainties to its maximum value. These boxplots enable comparison of the variable's distribution across HG certainty values **[R5]**. This way, the analyst can test their hypotheses on how indicative of HG the defined patterns are. Such testing can lead to rejection of the current hypotheses and generation of new ones, refinement of the specified assumptions or their acceptance and further incorporation in downstream statistical analyses. For a demonstration of the tool we refer to the video included in supp. materials at `osf.io/3ycp5`.

6. Demonstration

We explore the potential of Voidlytics in context of a case based on MoBa data, conducted in close collaboration with our domain experts (authors #3 and #4). Factoring in demand characteristics that appear in close collaboration [BRS11], we chose not to report the experts' subjective assessments of the tool but rather strive to provide a factual account of the analysis sessions conducted with Voidlytics. We describe the tool's capabilities for hypothesis generation and discuss its differences with the experts' usual workflow.

6.1. Protocol

We conducted an individual think-aloud session with each of the domain experts to observe how they used Voidlytics. Our goal with these sessions was to explore the potential utility of Voidlytics for hypothesis generation based on missing data. We kept an open session structure lasting between 60–90 minutes; these were conducted by paired interviewers. Each expert was first introduced to the case displayed, the three views of the tool, and the interactions available per view. The experts were then free to explore and use the tool while thinking aloud and describing their analysis actions. We observed and screen-recorded each session. The written protocol is available in supp. materials at `osf.io/3ycp5`.

6.2. Hypothesis generation

Session #1. One domain expert (author #4) started off by examining the **Exploration view**. Upon initial examination,

they expressed their prior assumption about missingness in the data: *"these [participants] who did not answer the first questionnaire, but answered the second one [. . .] would be the most interesting for us."* As as result, they unfolded one of the patterns with missing *Q1* and returned *Q2* as well as *Q3* to get a more detailed look at response patterns with the noted group of interest. The expert continued exploring the heatmap through iterative folding and unfolding of the patterns. Along the way, they noted that *"if they [participants] dropped Q3 we don't know if they were hospitalized or not"*, stating that hospitalization is deemed a strong outcome of HG. The expert then began selecting patterns as shown in Fig. 4. They first selected a pattern based on their assumption that *"[cases with] dropped Q2 and Q3 and no symptoms [according to Q1] – we can be pretty sure that they are healthy."*; they quantified this pattern's HG certainty as 0.1 in the **Definition view**. Another pattern was selected on the assumption that patients *"who did not answer the first [questionnaire], but answered later [in Q1 and Q2], the chances are higher [that they had HG]."*; they assigned this pattern a higher certainty of 0.5. After several iterations, the expert settled on four patterns.

Next, they pressed "Define HG" and proceeded to the **Analysis view** to test their assumptions. At the time of conducting this specific session, the view's right panel was using a scatterplot to relate the defined HG variable to a registry variable, e.g., placental weight or pregnancy duration. Given that MoBa data is population-based, the scatterplot was not an efficient representation of the relationships between the variables due to significant overplotting of data points. The view did not allow the expert to test their assumptions. Following the session, we revised the **Analysis view** by changing the right panel to display a set of boxplots, which we reasoned as a more efficient visual idiom for representing population-level summary statistics, and one familiar to the domain experts in their usual workflow. This improvement was carried into the second session.

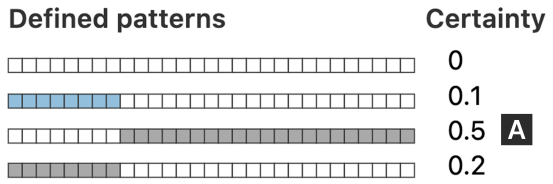

Figure 4: *Author #4 defined four patterns in total. (A) They assigned an HG certainty of 0.5 to response sets without Q1, but with Q2 and Q3.*

Session #2. The second expert (author #3) also began their interaction with Voidlytics by using the **Exploration view**. Through (un)folding interactions, they explored the heatmap and familiarized themselves with the patterns in the data. When asked to select and quantify patterns with different HG certainties, the expert noted that missingness analysis went against their day-to-day intuition with this data, stating that they *"keep thinking of the [observed] answers and not missingness."* Over the session, they became more comfortable with the missingness heatmap and were hypothesizing about missingness patterns as an indicator of HG. While exploring the heatmap, the expert explained some of their assumptions about the data. For instance, they *"[didn't] want to choose people who didn't have any data in the first trimester"* as *"the first trimester is the most likely for everyone to have nausea and vomiting, even people without HG"*. At the same time, the expert acknowledged that this reason for excluding such patterns may also be the reason to examine this response subset as it could include HG cases with severe symptoms that led to missing data in *Q1*.

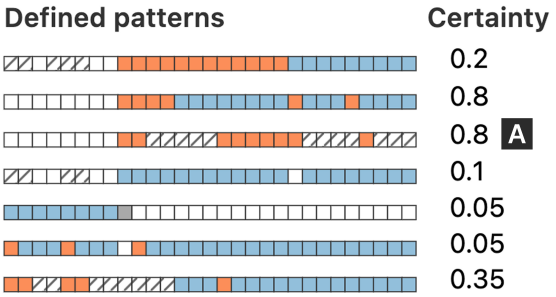

Figure 5: *Author #3 defined seven patterns in total. (A) They assigned an HG certainty of 0.8 to a pattern with missing Q1, reported nausea and vomiting in Q2 and reported long-term nausea and vomiting throughout the 2nd trimester, as well as hospitalization between weeks 17 & 20 in Q3.*

After identifying a pattern of interest, the expert selected this pattern of interest to formulate a hypothesis in the **Definition view**. The pattern in Fig. 5A demonstrates one of the assumptions formulated by the expert: response sets with missing *Q1*, followed by nausea and vomiting reported in *Q2*, and long-term nausea and vomiting in *Q3*, are classified as cases with a high HG certainty (0.8). After several iterations of exploration and pattern specification, the expert defined the HG variable based on the patterns shown in Fig. 5.

However, the **Analysis view** did not show significant differences in placental weights nor pregnancy durations across groups with high vs. low HG certainty. Consequently, the expert noted that their defined assumptions seem incorrect. They discussed reviewing and reiterating on their defined patterns and HG certainty values.

6.3. Results

The sessions suggest that the experts were able to explore missingness response patterns and hypothesize how indicative of HG these were. The **Exploration view** allowed the experts to (re)formulate their assumptions about missingness response patterns in the data. Through interaction with the **Exploration** and **Definition** views, they were able to

self-quantify their certainty about a pattern as indicative of HG, thereby making explicit their tacit assumptions regarding these patterns and HG. In the **Analysis view**, the expert could assess their assumptions by looking at the distributions of registry variables across different HG certainties. When there were no significant differences in distributions, the expert concluded their assumptions were likely incorrect and expressed a need to revise the patterns and certainties.

Though more validation efforts are required, we see this process of building, validating, and reiterating on assumptions as initial evidence of Voidlytics' potential to support hypothesis generation regarding missingness response patterns in relation to HG. One expert (author #3) notes that resulting hypotheses can be carried into downstream analysis: *"These visualizations are not the end product. It's kind of an intermediary. And then we would take it [HG variable] forward typically to check whether our genetic signals are associated with any of these patterns."*. However, both experts commented that it was unusual for them to analyze missingness response patterns. As author #3 notes, *"it takes a bit of time to get used to it because we are not used to comparing missingness patterns"*, while also mentioning that they *"haven't thought about missingness as a trait."* Over the session, this expert adjusted their intuition about missingness patterns as a result of interacting with Voidlytics.

7. Discussion, limitations, and future directions

Missingness-interrogative way of thinking. Voidlytics is designed to support the analyst in inferring, externalizing and testing their tacit assumptions about missingness patterns in MoBa data in relation to HG. Our study shows that the domain experts found it difficult initially to think about their data in terms of both observed and missing responses, yet they were able to shift their perspective in the process of using the tool. This suggests Voidlytics may encourage analysts to take on a missingness-interrogative approach to their work which supports identifying and incorporating informative missingness in data analysis as a feature (rather than an error). While this approach introduces initial cognitive load, provoking deliberate considerations of missingness may yield benefits of "desirable difficulty", e.g., inducing more active and attentive processing of information and generation of inferences to fill in knowledge gaps [HAS11].

Limitations in exploration and analysis. The current approach focuses on missingness response patterns and disregards complete responses. While this may introduce bias, we see our approach as complementary to the current exploratory and analytical methods used by the experts. Future work could explore the joint analysis of missing and complete response patterns for hypothesis generation of HG. Another constraint to comprehensive analysis lies in the **Analysis view**. Only the maximum certainty value for each response set is displayed, which may conceal variations in certainty. To avoid a possible upward bias, future work should explore the representation of all certainties assigned to each response set, e.g., using uncertainty encodings [SS18].

Limitations in usability and hypothesis generation. Currently, the **Definition view** does not give analysts the flexibility to customize a pattern of interest from scratch; this limited the missingness patterns that one expert (author #4) wished to test. Future designs should better support pattern-building during the analysis. Additionally, this view requires improved support for interpreting displayed patterns. Future work can explore on-demand visualizations to provide a detailed overview of a selected/defined pattern. Another area of improvement lies in reformulating assumptions: we note that analysts wish to iteratively refine their patterns until they identify ones that are indicative of HG. This raises the question on how we can better support hypothesis generation by providing visual indicators of meaningful relationships between missingness patterns and registry variables that distinguish these relationships from chance findings.

Scalability. Voidlytics displays only a set of binary questionnaire variables related to nausea, vomiting, and hospitalization. Yet MoBa data include ordinal, quantitative and categorical non-binary variables, some of which have hierarchal structure. With a growing number of variables, it becomes difficult to effectively display and pan across the response patterns in the heatmap. Future work should explore techniques to handle this complexity, such as by aggregating and filtering patterns based on analyst-defined conditions. For patterns, alternative solutions could also include more frequency-aware approaches, such as adjusting a pattern's height based on the number of its response sets.

Evaluation. Aside from think-aloud demonstrations, we considered other evaluation methods, e.g., user satisfaction and usability testing, to assess the value of Voidlytics. These methods were deemed unsuitable due to A) biases from co-design between visualization researchers and domain experts, and B) test criteria not related to our intended contribution (hypothesis generation of HG via missingness patterns). Given our limited access to two domain experts, future work should further assess the validity of Voidlytics through an independent evaluation with additional domain experts.

8. Conclusion

We present a design study providing initial evidence that Voidlytics may support hypothesis generation about the links between HG and missingness response patterns in the MoBa cohort study data. We show how our approach enables the consideration of missingness patterns and the analyst's tacit knowledge as an integral part of formulating hypotheses regarding indication variables for HG. Our study also suggests that Voidlytics may encourage an alternative, missingness-oriented way of thinking when exploring questionnaire data and can have broader applicability to other contexts that intrinsically involve missing data.

Supplementary materials

The source code for Voidlytics is available at github.com/Annchovy/voidlytics. All other supplemental materials are available at osf.io/3ycp5, released under a CC BY 4.0 license. The repository includes the interview questions for eliciting domain goals, protocol for the case study session, and manuscript figures.

Acknowledgments

We thank the Norwegian Institute of Public Health (NIPH) for gathering and maintaining the Norwegian Mother, Father and Child Cohort Study (MoBa) dataset. All analyses in the MoBa cohort were performed using digital labs in the HUNT Cloud at the Norwegian Univ. of Science and Techn., Trondheim, Norway.

Ethics declarations. Informed consent was obtained from all MoBa study participants. The administrative board of MoBa led by NIPH approved the study protocol. The establishment of MoBa and the initial data collection was based on a license from the Norwegian Data Protection Agency and approval from The Regional Committee for Medical and Health Research Ethics (REC). The MoBa cohort is currently regulated by the Norwegian Health Registry Act. The study was approved by The Regional Committee for Medical Research Ethics (REC 744956).

Data availability. MoBa data are managed by the NIPH and can be made available to researchers, provided approval from REC, compliance with the EU General Data Protection Regulation (GDPR) and approval from the data owners. The consent given by the participants does not open for storage of data on an individual level in repositories or journals. Researchers who want access to data sets for replication should apply through helsedata.no. Access to data sets requires approval from REC in Norway and an agreement with MoBa.

References

[AFDD*24] ALSUFYANI S., FORSHAW M., DEL DIN S., YARNALL A., ROCHESTER L., FERNSTAD S. J.: Multi-level visualization for exploration of structures in missing data. In *EG UK Computer Graphics & Visual Computing* (2024). doi:10.2312/cgvc.20241212. 2

[AFF24] ALSUFYANI S., FORSHAW M., FERNSTAD S. J.: Visualization of missing data: A state-of-the-art survey. *arXiv preprint arXiv:2410.03712* (2024). doi:10/rhz3. 2

[AM23] AKBABA D., MEYER M.: "Two heads are better than one": Pair-interviews for visualization. In *2023 IEEE Visualization and Visual Analytics (VIS)* (2023), IEEE, pp. 206–210. doi:10.1109/VIS54172.2023.00050. 3

[ANI*20] ALEMZADEH S., NIEMANN U., ITTERMANN T., VÖLZKE H., SCHNEIDER D., SPILIOPOULOU M., BÜHLER K., PREIM B.: Visual analysis of missing values in longitudinal cohort study data. *Comput Graph Forum 39*, 1 (2020), 63–75. doi:10.1111/cgf.13662. 2, 4

[BDBT25] BLUDAU M.-J., DÖRK M., BRUCKNER S., TOMINSKI C.: Fluidly revealing information: A survey of un/foldable data visualizations. *Comput Graph Forum 44*, 3 (2025), e70152. doi:10.1111/cgf.70152. 5

[BH99] BEYER H., HOLTZBLATT K.: Contextual design. *Interactions 6*, 1 (1999), 32–42. doi:10/cj4277. 3

[BHF15] BIERING K., HJOLLUND N. H., FRYDENBERG M.: Using multiple imputation to deal with missing data and attrition in longitudinal studies with repeated measures of patient-reported outcomes. *Clin Epidemiol* (2015), 91–106. doi:10.2147/CLEP.S72247. 2

[BHS*09] BRANTSÆTER A. L., HAUGEN M., SAMUELSEN S. O., TORJUSEN H., TROGSTAD L., ALEXANDER J., MAGNUS P., MELTZER H. M.: A dietary pattern characterized by high intake of vegetables, fruits, and vegetable oils is associated with reduced risk of preeclampsia in nulliparous pregnant Norwegian women. *J Nutr 139*, 6 (2009), 1162–1168. doi:10.3945/ajcn.110.001248. 1

[Bil18] BILOGUR A.: Missingno: A missing data visualization suite. *J Open Source Softw 3*, 22 (2018), 547. doi:10/rh2h. 2

[BM13] BREHMER M., MUNZNER T.: A multi-level typology of abstract visualization tasks. *IEEE transactions on visualization and computer graphics 19*, 12 (2013), 2376–2385. doi:10.1109/TVCG.2013.124. 4

[BOH11] BOSTOCK M., OGIEVETSKY V., HEER J.: D³ data-driven documents. *IEEE Trans Vis Comput Graph 17*, 12 (Dec. 2011). doi:10.1109/TVCG.2011.185. 4

[BRS11] BROWN B., REEVES S., SHERWOOD S.: Into the wild: Challenges and opportunities for field trial methods. In *Proc CHI* (2011), ACM, pp. 1657–1666. doi:10/d89p8q. 7

[EPD05] EATON C., PLAISANT C., DRIZD T.: Visualizing missing data: Graph interpretation user study. In *IFIP Conference on Human-Computer Interaction* (2005), Springer, pp. 861–872. doi:10.1007/11555261_68. 2

[FBKW19] FRANKE M., BARCZOK R., KOCH S., WELTECKE D.: Confidence as first-class attribute in digital humanities data. In *Proc 4th Workshop on Visualization for the Digital Humanities (VIS4DH)* (2019), IEEE, pp. 1–5. 3

[Fer19] FERNSTAD S. J.: To identify what is not there: A definition of missingness patterns and evaluation of missing value visualization. *Inf Vis 18*, 2 (2019), 230–250. doi:10/ggjbs6. 2

[FSS*18] FEJZO M. S., SAZONOVA O. V., SATHIRAPONGSASUTI J. F., HALLGRÍMSDÓTTIR I. B., VACIC V., MACGIBBON K. W., SCHOENBERG F. P., MANCUSO N., SLAMON D. J., MULLIN P. M., ET AL.: Placenta and appetite genes GDF15 and IGFBP7 are associated with hyperemesis gravidarum. *Nat Commun 9*, 1 (2018), 1178. doi:10.1038/s41467-018-03258-0. 1

[FW21] FERNSTAD S. J., WESTBERG J. J.: To explore what isn't there–glyph-based visualization for analysis of missing values. *IEEE Trans Vis Comput Graph 28*, 10 (2021), 3513–3529. doi:10.1109/TVCG.2021.3065124. 2

[FWZ*24] FEJZO M., WANG X., ZÖLLNER J., PUJOL-GUALDO N., LAISK T., FINER S., VAN HEEL D. A., BRUMPTON B., BHATTA L., HVEEM K., ET AL.: Multi-ancestry GWAS of severe pregnancy nausea and vomiting identifies risk loci associated with appetite, insulin signaling, and brain plasticity. *MedRxiv* (2024). doi:10/rhz2. 1

[HAD21] HANEUSE S., ARTERBURN D., DANIELS M. J.: Assessing missing data assumptions in EHR-based studies: A com-

plex and underappreciated task. *JAMA Network Open 4*, 2 (2021), e210184. doi:gj8vs8. 2

[HAS11] HULLMAN J., ADAR E., SHAH P.: Benefitting infovis with visual difficulties. *IEEE Trans Vis Comput Graph 17*, 12 (2011), 2213–2222. doi:10.1109/TVCG.2011.175. 9

[HUN26] HUNT CLOUD: HUNT Cloud documentation. https://docs.hdc.ntnu.no/, 2026. Accessed: May 01 2026. 3, 4

[JM22] JIMÉNEZ E., MACÍAS R.: Graphical tools for visualization of missing data in large longitudinal phenomena. *Comput Graph Forum 41*, 1 (2022), 438–452. doi:10/rhzz. 2

[KSA*01] KORKEILA K., SUOMINEN S., AHVENAINEN J., OJANLATVA A., RAUTAVA P., HELENIUS H., KOSKENVUO M.: Non-response and related factors in a nation-wide health survey. *Eur J Epidemiol 17*, 11 (2001), 991–999. doi:b88dvb. 1

[KT16] KOWARIK A., TEMPL M.: Imputation with the R package VIM. *J Stat Softw 74* (2016), 1–16. doi:10/gf4wzf. 2

[LAML22] LIN H., AKBABA D., MEYER M., LEX A.: Data hunches: Incorporating personal knowledge into visualizations. *IEEE Trans Vis Comput Graph 29*, 1 (2022), 504–514. doi:10.1109/TVCG.2022.3209451. 3

[LTC*21] LEE K. J., TILLING K. M., CORNISH R. P., LITTLE R. J., BELL M. L., GOETGHEBEUR E., HOGAN J. W., CARPENTER J. R., ET AL.: Framework for the treatment and reporting of missing data in observational studies: The treatment and reporting of missing data in observational studies framework. *J Clin Epidemiol 134* (2021), 79–88. doi:10/gh55hn. 2

[LWN19] LUPATTELLI A., WOOD M. E., NORDENG H.: Analyzing missing data in perinatal pharmacoepidemiology research: Methodological considerations to limit the risk of bias. *Clin Ther 41*, 12 (2019), 2477–2487. doi:10/gq6kr5. 2

[MAR*24] MONTAMBAULT B., APPLEBY G., ROGERS J., BRUMAR C. D., LI M., CHANG R.: Dimbridge: Interactive explanation of visual patterns in dimensionality reductions with predicate logic. *IEEE Trans Vis Comput Graph 31*, 1 (2024), 207–217. doi:10.1109/TVCG.2024.3456391. 4

[MBV*16] MAGNUS P., BIRKE C., VEJRUP K., HAUGAN A., ALSAKER E., DALTVEIT A. K., HANDAL M., HAUGEN M., HØISETH G., KNUDSEN G. P., ET AL.: Cohort profile update: The Norwegian mother and child cohort study (MoBa). *Int J Epidemiol 45*, 2 (Apr. 2016). doi:10/f8pdwv. 1, 3

[MGA24] MANZ T., GEHLENBORG N., ABDENNUR N.: Any notebook served: Authoring and sharing reusable interactive widgets. In *Proc 23rd Python in Science Conference* (Tacoma, Washington, July 2024). doi:10.25080/NRPV2311. 4

[MGM18] MCCURDY N., GERDES J., MEYER M.: A framework for externalizing implicit error using visualization. *IEEE Trans Vis Comput Graph 25*, 1 (2018), 925–935. doi:10.1109/TVCG.2018.2864913. 2

[MOR*16] MCPARLIN C., O'DONNELL A., ROBSON S. C., BEYER F., MOLONEY E., BRYANT A., BRADLEY J., MUIRHEAD C. R., NELSON-PIERCY C., NEWBURY-BIRCH D., ET AL.: Treatments for hyperemesis gravidarum and nausea and vomiting in pregnancy: A systematic review. *JAMA 316*, 13 (2016), 1392–1401. doi:10.1001/jama.2016.14337. 3

[Mun09] MUNZNER T.: A nested model for visualization design and validation. *IEEE Trans Vis Comput Graph 15*, 6 (Nov. 2009), 921–928. doi:10/dzqbc2. 3

[RAH22] RUDDLE R. A., ADNAN M., HALL M.: Using set visualisation to find and explain patterns of missing values: A case study with nhs hospital episode statistics data. *BMJ Open 12*, 11 (2022), e064887. doi:10/rh2g. 2

[Rub76] RUBIN D. B.: Inference and missing data. *Biometrika 63*, 3 (1976), 581–592. doi:10/g82vnq. 2

[SG02] SCHAFER J. L., GRAHAM J. W.: Missing data: Our view of the state of the art. *Psychol Methods 7*, 2 (2002), 147. doi:dsvqnz. 2

[SMM12] SEDLMAIR M., MEYER M., MUNZNER T.: 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.2012.213. 3

[SMW*22] SUN M., MA Y., WANG Y., LI T., ZHAO J., LIU Y., ZHONG P.-S.: Toward systematic considerations of missingness in visual analytics. In *2022 IEEE Visualization and Visual Analytics (VIS)* (2022), IEEE, pp. 110–114. doi:10/rh2m. 2

[SS18] SONG H., SZAFIR D. A.: Where's my data? evaluating visualizations with missing data. *IEEE Trans Vis Comput Graph 25*, 1 (2018), 914–924. doi:10/gqnj3m. 2, 9

[SSS*14] SACHA D., STOFFEL A., STOFFEL F., KWON B. C., ELLIS G., KEIM D. A.: Knowledge generation model for visual analytics. *IEEE Trans Vis Comput Graph 20*, 12 (2014), 1604–1613. doi:10.1109/TVCG.2014.2346481. 2

[ST17] SJÖBERGH J., TANAKA Y.: Visualizing missing values. In *2017 21st International Conference Information Visualisation (IV)* (2017), IEEE, pp. 242–249. doi:10/rhzw. 2

[SWB*24] SUN M., WANG Y., BOLTON C., MA Y., LI T., ZHAO J.: Investigating user estimation of missing data in visual analysis. In *Proc 50th Graphics Interface Conference* (2024), pp. 1–13. doi:10.1145/3670947.3670977. 2

[TC23] TIERNEY N., COOK D.: Expanding tidy data principles to facilitate missing data exploration, visualization and assessment of imputations. *J Stat Softw 105*, 7 (2023), 1–31. doi:10.18637/jss.v105.i07. 2

[TGH*23] TAN A. L., GETZEN E. J., HUTCH M. R., STRASSER Z. H., GUTIÉRREZ-SACRISTÁN A., LE T. T., DAGLIATI A., MORRIS M., HANAUER D. A., MOAL B., ET AL.: Informative missingness: What can we learn from patterns in missing laboratory data in the electronic health record? *J Biomed Inform 139* (2023), 104306. doi:10.1016/j.jbi.2023.104306. 2

[VMM22] VEJRUP K., MAGNUS P., MAGNUS M.: Lost to follow-up in the Norwegian mother, father and child cohort study. *Paediatr Perinat Epidemiol 36*, 2 (2022), 300–309. doi:10.1111/ppe.12821. 1, 2

[VSMG13] VIKANES Å. V., STØER N. C., MAGNUS P., GRJIBOVSKI A. M.: Hyperemesis gravidarum and pregnancy outcomes in the Norwegian mother and child cohort–a cohort study. *BMC Pregnancy Childbirth 13*, 1 (2013), 169. doi:10/gb8xz4. 1

[WCNK13] WELLS B. J., CHAGIN K. M., NOWACKI A. S., KATTAN M. W.: Strategies for handling missing data in electronic health record derived data. *EGEMS (Wash DC) 1*, 3 (2013), 1035. doi:10.13063/2327-9214.1035. 2

[ZvKC*24] ZAMANIAN A., VON KLEIST H., CIORA O.-A., PIPERNO M., LANCHO G., AHMIDI N.: Analysis of missingness scenarios for observational health data. *J Pers Med 14*, 5 (2024), 514. doi:10.3390/jpm14050514. 1, 2