The MoBa GWAS Explorer: Designing Approachable Visualizations of GWAS Data for a Mixed Audience

Hanna Balaka* University of Bergen
Marc Vaudel† University of Bergen
Laura A. Garrison‡ University of Bergen

Figure 1: ***Left panel:*** The *MoBa GWAS Explorer* interface overview: (1) Information bar links to details on the tool and on the MoBa study, (2) Query view, (3) Genomic Heatmap view, (4) Manhattan Plot view, (5) Temporal Forest Plot view, (6) GWAS Explorer Tutorial. ***Right panel:*** The *MoBa GWAS Explorer* Tutorial. (a) and (b) show two steps in the non-expert user onboarding tutorial.

Abstract

Public health studies generate extensive datasets providing important insights into human health. The Norwegian Mother, Father, and Child Cohort Study (MoBa) is a longitudinal cohort study capturing information on pregnancy and early childhood. This information helps uncover the genetic underpinnings of traits or diseases drawing interest from researchers in public health. Non-experts are also attracted to the study, both to understand their contributions as data donors and relevant health determinants. However, the complexity of MoBa data hinders its exploration, analysis, and dissemination. We present a design study exploring the needs and uses of the MoBa dataset in a mixed-user context and introducing the *MoBa GWAS Explorer*, a web-based visual tool for exploration and analysis of MoBa data by a mixed audience. This tool supports experts in exploring and analyzing MoBa data interactively. Though designed primarily for researchers, we explored the potential for onboarding strategies to make this tool more approachable for non-experts. We conducted a qualitative study with both user groups to evaluate their experience with the tool and its usability. Our evaluation indicates that the application, along with the integrated onboarding, has potential to serve both expert and non-expert groups. Supplementary materials for this study are available at https://osf.io/k5bvj/.

Index Terms: Visual analytics, design study, genomics data, onboarding, mixed audience.

*e-mail: hanna.balaka@uib.no
†e-mail: marc.vaudel@uib.no
‡e-mail: laura.garrison@uib.no

1 Introduction

The public health field plays an important role in improving population health by analyzing disease trajectories and health patterns within a population. Genetic epidemiology is an area of public health focused on how human genetic factors influence traits including health and diseases [21]. To gain insights into health and disease patterns, genetic epidemiologists analyze large amounts of data collected through longitudinal cohort studies, such as the Norwegian Mother, Father, and Child Cohort Study (MoBa) [26] that captures important information on pregnancy and early childhood. The study collects participants' genetic information alongside data on observable traits or diseases. Moreover, MoBa includes large-scale gene association studies exploring links between genetic data and physical traits or diseases on a population scale, offering insights into how genetics influence our health.

While providing insights into how genetics shape our health trajectories from pregnancy to early childhood, MoBa study data pose numerous challenges due to their complex nature. The study includes data on a range of traits and diseases across different time points and genetic information of children, fathers, and mothers participating in the study. This results in a dataset which is extensive, multidimensional, and composed of different data types. Furthermore, MoBa data present a special challenge for analyzing the results of **genome-wide association studies (GWAS)**. GWAS results are summary statistics describing associations between genetic information and observable traits and diseases. These statistics are typically organized in tables containing millions of entries. Researchers often rely on various visualization methods to analyze these data that we outline later in this section. While challenging for domain experts, this type of data is even more difficult for people without a specialized background. Understanding how genetics affect specific traits or diseases requires a solid foundation in genetics and statistical analysis, which many non-experts may not have. Here, we define non-experts as individuals interested in learning how genetics shapes human health or how data collected through

cohort studies are being used, and who do not have knowledge of genetics beyond an introductory biology course.

Several visualization techniques have been developed to facilitate the analysis of genetic data [31] and data on observable traits and diseases [34], including specialized toolkits like Gosling [24]. However, these approaches target experts and are ill-suited to non-experts out-of-the-box. Projects like FinnGen [15] and UK Biobank [43] offer web interfaces for exploring genotypic and phenotypic cohort data similar to MoBa. However, they require domain knowledge and are inaccessible to a broader audience, including those who donate data and may want to access their own information. FinnGen's companion scrollytelling website [18] explains genetics and GWAS to non-experts, yet lives detached from the expert interface. We focus on integrating expert and non-expert experiences within a single application through onboarding, investigating the extent to which this strategy can empower non-experts to explore tools designed for experts. This fusion of user experiences can promote greater transparency between researchers and participants by allowing users to directly interact with the tools used by experts. Moreover, this approach has the potential to reduce the development time of such applications and reach a broader audience that ranges along a spectrum from experts to non-experts.

Our primary contribution is a design study that explores the potential for a single visual analytics application to serve two disparate audiences using visualization onboarding, which we evaluate through a between-subjects qualitative study. Through this design study **we also contribute the *MoBa GWAS Explorer***, a web-based visual tool to explore and analyze MoBa data.

2 Related Work

Our research builds on prior work on genomic data visualization, visual science communication, and visual onboarding strategies.

Visualization of genomic data Several visualization techniques have been developed for genomic data analysis [31]. For example, Gosling [24] is a specialized toolkit that offers visualization tools tailored to researchers' datasets and questions. Common visualizations for GWAS data include Manhattan plots, which display SNPs at their genomic positions along the *x-axis* and their corresponding negative logarithms of the *p-values* along the *y-axis* to produce a visual metaphor echoing the Manhattan skyline. Tools like *qqman*[42], *GWAMA* (Genome-Wide Association Meta-Analysis) [25] and *SNPEVG* [45] support creating these plots. However, these packages primarily serve experts with domain knowledge and an understanding of what to visualize. In contrast, our work aims to serve both expert users and individuals without relevant domain background. Other relevant efforts include the FinnGen project [15] and a large-scale biomedical database UK Biobank [43], which allow users to explore genotypic and phenotypic cohort data similar to MoBa. While these projects offer public-facing visual web interfaces, they require domain expertise to be fully usable and are unapproachable for broader audiences.

FinnGen's companion scrollytelling website [18], by contrast, combines illustrations and narrative approaches to explain human genetics and GWAS to non-experts. Another example is a design study conducted by Zhang et al. [46], which explores design strategies for the public communication of GWAS data in MoBa. Both of these works are illustrative of visual science communication as opposed to visual analysis efforts in describing genetics data.

Visual science communication Science communication strategies are essential for bridging the gap between the lab bench and the public sphere [33]. Trumbo [41] emphasizes the challenge of communicating scientific information to individuals without a relevant domain background and their limited access to a specialized vocabulary within a scientific discipline. Visual science communication has become a key practice for enhancing explanation, comprehension, and relatability of scientific topics [16, 13] for a broader audience [8]. Examples range from infographic and illustration work [44, 17] to more technically complex illustrative [35] and narrative visualization projects [36] which blend visual abstraction, storytelling and interactivity to convey messages with sensitivity and clarity. For example, Meuschke et al. [28] explore narrative visualization techniques to effectively communicate scientific findings about diseases to a general audience, including patients and their families. An additional example of narrative visualization, developed outside of academia, is a tool called *The MicrobeScope* [27], which offers four stories on diseases and uses interactive visualization and text to familiarize users with the data.

Visualization onboarding Visualization tools support data understanding, but users still need to learn how to use them, especially those with low visualization literacy or limited familiarity with the underlying data. Effective visualization onboarding [38] can facilitate this learning process and empower users to use visualization tools and better understand the underlying data. Stoiber et al. [37] emphasize the importance of guiding users through complex data representations to improve their comprehension of the representations and support knowledge discovery. Another work by Stoiber et al. [38] examines different visual onboarding methods such as interactive step-by-step guides, scrollytelling, and video tutorials. Their evaluation shows that video tutorials received the most positive feedback, followed by step-by-step guides and, lastly, scrollytelling. Our choice of onboarding strategy is guided in part by these results. Similarly, Kwon et al. [22] examine different methods to help users learn to use visualizations, comparing the following approaches: baseline (i.e., no tutorial), static, video and interactive tutorials. Their results indicate that participants using the interactive tutorial performed better on analytical tasks and reported a more engaging experience. In addition, Tanahashi et al. [40] explore how online guides with exercise questions support users in learning visualizations. Their results suggest that exercise questions in tutorials improve visualization comprehension. Inspired by these studies, we adopt an interactive step-by-step guide with exercises to bridge the gap between expert and non-expert users.

3 Domain Background and Context

Our collaborating domain involves specialized data and terminology that may be unfamiliar. We briefly introduce key concepts relevant to GWAS before discussing our research approach. The MoBa study contains phenotypic and genotypic information. **Phenotypic data** describes phenotypes, which are observable characteristics or traits of an organism, e.g., physical features or disease manifestations. MoBa phenotypic data includes information like physical measurements (e.g., height and weight), medical history, and demographic details (e.g., marital status and income, etc). **Genotypic data** refers to the genetic makeup of an individual and is obtained from an individual's DNA.

DNA, or deoxyribonucleic acid, is a molecule that carries the genetic instructions for an organism [1]. Composed of units called **nucleotides**, each nucleotide consists of three components: a phosphate group, a sugar molecule, and one of four chemical bases: adenine (A), thymine (T), cytosine (C), and guanine (G). These chemicals form specific **base pairs**: A with T, and C with G. The sequence of these DNA base pairs build the complete genetic code, or **genome** [2], which holds the information needed for an organism's development, function, and inheritance. The mammalian genome is organized into structures called **chromosomes**. Humans (typically) have 23 pairs of chromosomes with a total of 46 chromosomes in each cell. Each chromosome in a pair is inherited from one parent.

Genomic variation refers to differences in DNA sequences among individuals or populations [3], and can be caused by differences in base pairs, additions, or deletions. **Single-nucleotide polymorphisms (SNPs)** are the most common type of genomic

variation in humans. Different individuals may have different nucleotides at specific positions in their DNA sequences. These differences, represented as single-letter changes, are SNPs. For example, at a particular position in the genome, one person may have a G, while another person an A. SNPs can affect how genes function and influence specific phenotypes: some SNPs are associated with an increased risk of certain diseases, while others with physical traits such as height or weight. An individual may have zero, one (inherited from one parent), or two (inherited from both parents) copies of a certain SNP, e.g., a G nucleotide, at the same position. In some cases, having two copies of a SNP can have an additive effect on the phenotype, with each copy contributing equally.

Genome-wide association studies (GWAS) rely on SNPs to identify and statistically quantify associations between genomic variations and traits or diseases [9]. Methods for quantification are logistic regression for binary traits (e.g., disease status) and **linear regression for continuous traits** (e.g., height). Our study focuses primarily on quantitative traits.

Different SNPs may have different effects on a trait or disease within a population. These effects are measured by the effect size, commonly represented as the **beta coefficient (β)**. This beta coefficient signifies the deviation from the phenotype mean value caused by the presence of a certain SNP within a genome. Its magnitude reflects the strength of the effect the SNP causes on the phenotype. Its sign implies the character of the effect: if the beta coefficient for a specific SNP is positive, it means the presence of this SNP tends to increase the value of the phenotype and vice versa. For example, a SNP with a positive beta coefficient for BMI implies that this SNP is associated with higher BMI.

The linear regression model for quantitative traits also provides an estimate of the error in beta estimation, called standard error, as well as reporting a p-value reflecting the probability of observing a similar or higher beta and standard error by chance. Domain researchers widely use a **p-value threshold** of $(5e-8)$ to discriminate between **strong and weak associations** [12]. Genetic associations with p-values greater than $(5e-8)$ are considered to have a weak association with a phenotype, while lower p-values are considered genome-wide significant. This corresponds to a threshold of 0.05 with one million independent tests. Genetic associations with p-values higher than $(1e-4)$ are considered spurious. Since p-values of interest are so small, researchers transform them into negative logarithms to ease identification of significant associations, which results in the skyscraper metaphor of the Manhattan plot described previously. These negative logarithms of the p-values are referred to as significance levels. SNPs with significance levels above the threshold $-\log_{10}(5e-8)$ are considered strongly associated with a phenotype, while lower significance levels indicate weak associations. SNPs with significance levels below $-\log_{10}(1e-4)$ are considered spurious. These different thresholds are important to show, with different degrees of priority and salience, for a complete picture of genomic association likelihoods.

4 Data and Tasks

Our approach is guided by Munzner's nested model for visualization design [29]. We begin with a characterization and abstraction of data and tasks.

Data The MoBa study includes data from over 100,000 pregnancies, documenting genomic information and data for a wide range of phenotypes. The data collection has been going on from 1998 (the start of participant recruitment) to the present. As our work is dedicated to visualization techniques rather than an exhaustive analysis of all phenotypes within the MoBa dataset, we focused on three widely available and interpretable phenotypes: *weight*, *length*, and *Body Mass Index (BMI)* of children from birth to eight years of age. These phenotypes serve as primary indicators of growth, development, and overall health of children. MoBa collected measurements for each phenotype at twelve distinct time points to cover the stages of development before puberty. This information enables researchers to examine how associations between genetics and phenotypes change over time. After data collection, domain experts performed GWAS analyses for each phenotype and time point. Moreover, they conducted analyses against three genomes—fetal, maternal, and paternal—allowing researchers to explore whether a trait or health condition is influenced by genetic material inherited from a parent, the genetic material of the fetus, or both. Consequently, the data is stratified along three variable groups: phenotype (categorical), time point (quantitative) and genome (categorical). Following this stratification, the data is organized into multiple subsets, each containing GWAS data for a specific combination of phenotype, time point, and genome, This results in a **multidimensional table with mixed data types**.

Expert tasks Following data characterization and abstraction we proceeded to task analysis, beginning with **expert tasks** that we defined through a requirements analysis. We conducted semi-structured interviews with three domain experts conducting research on MoBa data: a geneticist, a genetic epidemiologist, and a psychiatric genetic epidemiologist. Each researcher participated in a single interview session lasting 45-60 minutes. The interviews included open-ended questions probing participants' tasks with MoBa data, workflows, specific pain points or limitations in current tools or processes. For example, we asked the participants to describe the main features of the MoBa dataset that they were working with or were interested in exploring. For the complete set of interview questions, we refer to supplementary materials. Afterward, the author team met over several sessions to synthesize domain needs into the following tasks and questions:

- **ET1**: Effectively navigate MoBa study variables to **identify and select data subsets** to **explore particular combinations** of phenotypes, time points, and genomes within the study.
- **ET2**: **Compare** the associations between phenotypes and SNPs across fetal, maternal, and paternal genomes.
- **ET3**: **Browse** SNPs for a specific phenotype, time point, and genome combination. They want to **identify** significant SNPs or **look up** and verify a specific SNP.
- **ET4**: **Explore** the dynamics of SNP impact (i.e., beta coefficients) over time.
- **ET5**: **Summarize** the genetic basis of a phenotype.

Non-expert tasks We define non-experts as individuals interested in learning how genetics shape human health or how data collected though public health studies are being used, and who do not have deep knowledge of genetics and statistical analysis. We synthesized **non-expert tasks** through literature review [19, 32] and discussions with the same set of domain experts on their experience with public interest in the MoBa study. Following this process we determined the following tasks and questions that non-experts may aim to solve:

- **NET1**: A non-expert may want to **discover** information about genetics and their impact on health to satisfy their curiosity about the topic.
- **NET2**: Non-experts aim to learn how a specific trait is affected by genetics by **looking up** GWAS data for a specific phenotype and time point, and **browsing** the corresponding SNPs' significance levels.
- **NET3**: A non-expert may seek to **discover** more about the methods employed by researchers to examine association between a phenotype and genomic variation.
- **NET4**: MoBa participants wish to learn how their personal information is secured. They may aim to **explore** the entire dataset available to ensure that their data are not revealed.

We used our data and task characterization to develop the *MoBa GWAS Explorer*, an exploratory tool adopting onboarding strategies to help a mixed audience learn about GWAS data from the MoBa

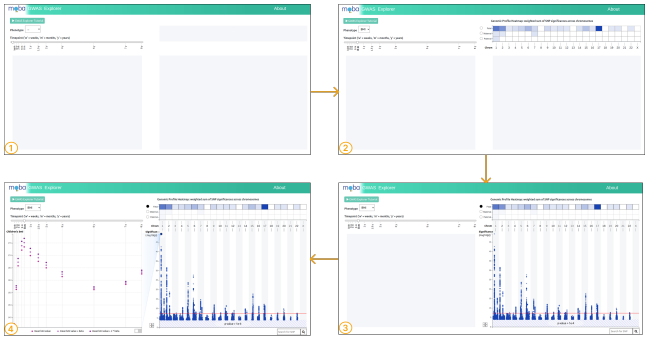

Figure 2: The *MoBa GWAS Explorer* overview. The exploration process: (1) shows the initial state of the dashboard, (2) – the Genomic Heatmap View reveals an overview of the GWAS data for three genomes, (3) – the Manhattan Plot View reveals the GWAS data for a genome of interest, (4) – the Temporal Forest Plot View reveals the dynamics of SNP impact over time.

project while facilitating expert analysis. We developed this tool using *Django* [14] as a backend and *D3.js* [6] to serve the frontend.

5 MoBa GWAS Explorer

Our tool consists of a navigation panel with two pages: *MoBa GWAS Explorer* and *About*. The *About* page contains contextual information about the MoBa study and the interface itself to help researchers and lay audiences understand the context of the *MoBa GWAS Explorer* and, hence, facilitates all the expert and non-expert tasks. Illustrated in Fig. 1, the *GWAS Explorer* page contains our developed tool for GWAS exploration and consists of multiple views: **Query**, **Genomic Heatmap**, **Manhattan Plot** and **Temporal Forest Plot** views. We provide the mapping of the views and the corresponding tasks each of them supports in the supplemental materials. Through these different views, an expert user can, for instance, use the Query and Genomic Heatmap views to find SNPs linked to BMI at 8 months in the maternal genome, then use the Manhattan Plot to identify top associations and the Temporal Forest Plot to explore when they most impact the phenotype. Additionally, the interface features a tutorial designed to onboard non-expert users [NET1, NET3, NET4].

Upon opening the application, users see the interface, as shown in Fig. 2.1, with the Genomic Heatmap, Manhattan Plot, and Temporal Forest Plot initially displayed as placeholders. This approach provides *orienting* support which we adopt from visualization guidance literature [10, 11], an area closely related to onboarding in visual analytics. As the user explores the application, the views gradually reveal the corresponding MoBa data, as shown in Fig. 2, which for non-experts is further facilitated by a tutorial. This gradual unfolding of the views supports incremental exploration, and allows the user to process these multidimensional data in smaller, more manageable subsets [30].

Query view Starting in the top-left corner, the **Query view** (shown in Fig. 1.2) serves as the initial point of interaction with the dashboard and provides users with tools to query the dataset and explore its content. The view includes a drop-down menu for phenotype selection and a timeline slider for specifying a time point of interest [ET1, NET1, NET2, and NET4].

Genomic Heatmap view Adjacent to the Query view, the **Genomic Heatmap view** supports user navigation of the genome variable. The view consists of radio buttons for selecting a genome of interest [ET1, NET1, NET2, and NET4] and a grid heatmap showing the aggregated significance levels of SNPs across chromosomes for each genome [ET2, ET3, ET5, NET1, NET4]. The radio buttons represent the categories of the genome variable (fetal, maternal, and paternal) and enable the user to select a genome to examine. The heatmap offers an overview of GWAS data by displaying the weighted sum of SNP significance levels per genome and chromosome, enabling comparative analysis across genomes and guiding further detailed exploration [ET2]. The heatmap's color scale ranges from zero to the maximum weighted sum across all chromosomes and genomes for the selected phenotype and time point. When a user selects a genome, it is highlighted. Changing the phenotype or time point updates the heatmap to reflect the new query variables.

Manhattan Plot view The **Manhattan Plot view** appears below the Heatmap view. For a given combination of phenotype, time point, and genome, we used a Manhattan plot to represent the corresponding GWAS results [ET1, ET3, ET5, NET1, NET2, NET4]. A Manhattan plot was considered the most suitable representation for this task for several reasons. Firstly, a Manhattan plot provides an overview of genome-wide associations [ET5, NET4]. Secondly, Manhattan plots have a linear layout which facilitates the comparison between data points [7] and, thus, supports the search for the most significant genomic variations [ET3, NET1, NET2]. In addition, the interviews with the experts revealed that Manhattan plots are widely used in scientific publications to communicate GWAS results. As a result, many researchers are likely familiar with this representation, which facilitates its use.

The SNPs in the Manhattan plot are organized by chromosome, and SNPs from different chromosomes are spatially separated. To visually connect the Manhattan plot's chromosome regions with the corresponding columns in the heatmap and distinguish neighboring chromosomes, we use alternating vertical white and grey background shading, shown in Fig. 1.4. Additionally, linking the aggregated data in the heatmap's cells with a more detailed representation of SNPs in the Manhattan plot helps the user identify chromosome regions with dense SNP clusters.

The GWAS dataset contains a significant amount of data, largely due to the extensive SNP arrays that make visualization and analysis especially challenging. However as discussed in Sec. 3, SNPs with significance levels below $-\log_{10}(1e-4)$ are considered spurious and can be mainly ignored. As a result, the Manhattan plot displays only the SNPs with a significance level greater than $-\log_{10}(1e-4)$. To visually represent the SNPs with a significance level below $-\log_{10}(1e-4)$ and avoid the misinterpretation that there are no SNPs below the significant threshold, we included a visual metaphor in the form of a patterned (striped) rectangle, demonstrated in Fig. 1.4. The plot contains a button, shown in Fig. 1.4b, that allows the user to collapse this rectangle. This results in a more focused view and, consequently, facilitates NET1, NET2, and NET4. Following discussions with domain experts, we incorporated a threshold to differentiate between SNPs with strong and weak levels of association within a selected phenotype. This threshold is encoded by a horizontal red line shown in Fig. 1.4 and equals $-\log_{10}(5e-8)$. This encoding helps users quickly identify the most relevant SNPs based on their significance level [ET1, ET3, NET1, NET2, NET4]. To provide details-on-demand [5], hovering over data points displays detailed information about SNPs, such as SNP identifiers, chromosome name, genomic position, and significance (here, the negative logarithm of a p-value) [ET3]. Additionally, we supplemented the Manhattan plot with zooming and panning to help the user focus on specific areas of the visualization.

The user can select a particular SNP by selecting a point or using a search box, shown in Fig. 1.4a [ET1, ET3]. When a SNP is selected, the *MoBa GWAS Explorer* displays blue lines, illustrated in Fig. 1.4, that link the SNP and the next view.

Temporal Forest Plot view The next view is the **Temporal Forest Plot view**. The name refers to a forest plot—a representation

commonly used in medical studies to display study estimates alongside their confidence intervals [23], both of which are available in MoBa data. Initially, the view appears as a placeholder, indicating that more information will be shown once a SNP in the Manhattan Plot is selected. The Temporal Forest Plot displays phenotype mean values alongside the corresponding deviations from the trait means caused by the presence of one or two SNPs in a genome. The view displays several measurements across the time points (hence, our metaphor of a forest over time):

- The mean values of a phenotype selected in the Query view.
- The deviations from the mean values caused by a SNP selected in the Manhattan Plot View when **one** instance of the SNP is present in a genome. For each time point, a deviation is calculated as $mean + \beta$, where *mean* and β are the phenotype mean value and the beta coefficient at this time respectively.
- The deviations from the mean values caused by a SNP selected in the Manhattan Plot View when **two** instances of the SNP are present in a genome. For each time point, a deviation is calculated as $mean + 2 * \beta$.

The y-axis represents phenotype measurements labeled as "Children's *phenotype*", where *phenotype* is the selected phenotype. The slider in the Query View above serves as the *x-axis* for this view. To visually link the two views, dashed lines, shown in Fig. 1.5, extend from the slider's tick labels through the Temporal Forest Plot. We represent the mean values as magenta points, while $mean + \beta$ and $mean + 2 * \beta$—light and dark purple points, respectively. We encode confidence intervals for $mean + \beta$ and $mean + 2 * \beta$ values as semi-transparent light purple lines. Users can zoom along the *y-axis* and pan for closer examination [ET4, NET1, NET2, NET4]. Additionally, a toggle button, shown in Fig. 1.5a, allows switching between two representations: deviations including the phenotype mean values ($mean + \beta$ and $mean + 2 * \beta$) and deviations showing only the beta coefficients β and $2 * \beta$. The latter highlights when a selected SNP has the largest effect on the phenotype, indicated by the highest point [ET4, ET5, NET1, NET2, NET4].

Non-Expert tutorial The *MoBa GWAS Explorer* offers an interactive tutorial shown in Fig. 1.a and Fig. 1.b that supports non-expert tasks NET1, NET2, NET3, and NET4. The tutorial aims to explain MoBa data and their analysis to non-expert users, for example, those who are interested in learning how genetic studies are conducted in MoBa or how genetics may influence their child's health, such as BMI. The tutorial leads the non-expert user through the dashboard's views, describing the analysis pipeline using a particular example: the analysis of GWAS results for BMI at eight months for the fetal genome. The tutorial covers various concepts from basic genetic terms (DNA, SNP) to advanced topics (GWAS, beta coefficients). The tutorial acts as an interactive step-by-step guide integrated in the *MoBa GWAS Explorer* as inspired by prior work [37, 38] that combines textual descriptions and illustrations [44, 35] and learning-by-doing approach [22, 40]. The tutorial follows the martini glass narrative structure [36] to guide the user through the sections of the dashboard. Upon completion, the tutorial closes inviting the user for more independent exploration. We provide a recording of the tutorial in the supplemental materials.

The user launches the tutorial via the button shown in Fig. 1.6. Upon activation, the dashboard fades and becomes inactive, while a modal window appears in the center of the screen showing the onboarding content. This window contains textual descriptions, some of which are supplemented with references to external sources and illustrations to provide a second channel for information processing [44, 35]. The textual descriptions are split into steps, with each step focusing on a specific term or concept related to genetics and MoBa study. The user uses buttons to navigate through each step, and may exit the tutorial at any time. To support certain textual descriptions, the tutorial reveals parts of the dashboard. When these parts are revealed, the user can not interact with them unless instructed to do so by the tutorial. While guiding the user through different elements of a tool, the tutorial displays or highlights these elements on their mention in the onboarding window. For example, when describing SNPs with strong associations, the tutorial brings up the relevant SNPs by reducing the opacity of the other SNPs [20]. Similarly, when the window explains SNPs with weak significance, the corresponding SNPs are brought into focus. Following the learning-by-doing approach [22, 40], certain steps prompt the user to take specific actions. For instance, when explaining Temporal Forest Plot view and effect size estimate, the tutorial instructs the user to toggle the deviation view, as shown in Fig. 1.a. Advancing the the next step is not possible until the user has completed the necessary action (Fig. 1.a and Fig. 1.b). The tutorial concludes, and invites the user to explore the tool independently.

User ID	User type	Expertise	Study phase
E01	Expert	Genetic epidemiology	Intermediate
E02	Expert	Genetic epidemiology	Intermediate
E03	Expert	Genetics	Summative
NE01	Non-Expert	Digital culture	Summative
NE02	Non-Expert	Digital culture	Summative
NE03	Non-Expert	Comparative politics	Summative
NE04	Non-Expert	Psychology	Summative
NE05	Non-Expert	Carpentry	Summative

Table 1: Overview of participants for intermediate and summative evaluations.

6 Evaluation

Our primary aim with this study was to understand the feasibility of the *MoBa GWAS Explorer* for use by both expert and non-expert groups, supported by visualization onboarding. In the early phases of our study we included checkpoints for discussion and feedback with two domain experts. We did not include intermediate evaluation with non-expert users due to resource and time constraints. We evaluated the completed application in a pilot study with detailed 1:1 interviews with a domain expert and five non-experts. This section outlines the evaluation process and feedback from both groups.

Expert evaluation During the development phase, we conducted an **intermediate evaluation** with two **domain experts** [E01, E02] through our collaboration with the MoBa research community. We presented early prototypes to expert E01 and asked for feedback on their clarity and usefulness for the expert's work. Expert E02 helped evaluate a later version of the tool through open-ended exploration while thinking aloud. Afterward, we had an informal discussion on the dashboard's strengths and weaknesses.

Following feedback integration, we held a summative evaluation of the tool in a 1:1 semi-structured interview with the geneticist who had not been involved in the intermediate interviews [E03]. To simulate a real-world use case, we asked the expert to explore the *MoBa GWAS Explorer* freely. We observed and screen-recorded the session to identify moments of hesitation or confusion. Afterward, we asked the geneticist usability-focused questions drawing on the principles of usability testing [4]. For instance, we asked the expert to describe how easily they could navigate through the dashboard. The goal was to assess end-to-end user experience, gather overall user impression and identify opportunities for improvement. When asked to rate their satisfaction with the overall usability of the dashboard on a scale of 0–3 (0: *not satisfied* and 3: *fully satisfied*), the expert gave a score of 1.5 referring to the navigation as a weak point, especially after selecting a phenotype, where the next steps were unclear. Despite this, they considered the dashboard moder-

ately useful and highly likely to be used in their professional work. Additionally, the geneticist noted positive aspects of the dashboard and provided recommendations for additional features and possible improvements, such as a tutorial for experts to learn about the offered interactions and navigation within the application.

Non-expert evaluation We evaluated the dashboard through 1:1 interviews with five **non-expert** participants. We recruited the participants using convenience and purposive sampling [39], targeting individuals without knowledge of genetics and statistical analysis. Each interview lasted approximately one hour and consisted of three stages. In the **first stage**, the participants were asked pretesting questions covering their background and demographics, motivation to engage with genomic data, and familiarity with science, health, genetics, as well as their data literacy. For example, we asked how often they encountered data visualizations in general, specifically in health-related contexts. We also probed their privacy concerns regarding the sharing of personal information—an essential part of the MoBa study. This is particularly relevant in the Norwegian context where many people are study participants and may be concerned about how their data are being used. The **second stage** involved direct interaction with the *MoBa GWAS Explorer*. We asked the participant to interact with the dashboard freely, simulating a real-world scenario when the user explores the application independently. We observed their interaction and recorded the time required to complete the tutorial. The **final stage** included posttesting questions on their experience with the tutorial, their understanding of its content, and interpretation of the application views.

Participants came from different backgrounds unrelated to genetics. After being introduced to MoBa study, the participants expressed interest in learning more about the study, how genetics relate to diseases and how researchers uncover this information. Most participants rated their familiarity with science and genetics as low to average [NE01, NE02, NE03, NE05], except for one participant who rated their knowledge slightly above average [NE04]. In our data literacy pre-assessment, most participants reported limited exposure to data visualizations, especially in health contexts. To assess privacy concerns, we used an informal 0-–3 scale to measure participants' comfort levels with sharing personal information, where 0 indicated *not comfortable*, and 3 was *fully comfortable*. Three out of five participants [NE03, NE04, NE05] would be comfortable sharing their personal information, giving a score of 3. The remaining participants [NE01, NE02] rated it 2, expressing concerns about sharing biological data. The results suggest that while some users may be open to contributing their data, data sharing remains a concern, underscoring the importance of clear communication about data use in the *MoBa GWAS Explorer*.

In the **second stage**, on introduction to the tool two participants [NE02, NE04] immediately launched the tutorial. The remaining participants [NE01, NE03, NE05] started interacting with the views instead, explaining that that they *"never read instructions or tutorials, usually skip them"* [NE03] and *"felt confident with the new tool"* [NE05]. These participants expressed confusion about the views and eventually launched the tutorial. Overall, participants smoothly followed the tutorial steps and successfully performed the instructed actions. Only one required clarification for how to complete tasks in the tutorial, thinking they had to close the onboarding window before interacting with the views. All participants progressed quickly through most sections explaining phenotypes, SNPs, GWAS, genomes, and the Manhattan plot. However, they spent most of their time on the Temporal Forest Plot and beta coefficients. In total, the tutorial took between 9–24 minutes with an average time of 16 minutes. After finishing the tutorial, all participants were interested in continuing to interact with the dashboard by exploring different variables. When asked about the tutorial's content, participants demonstrated an understanding of the term "phenotype", but most struggled to describe "SNP" and "GWAS." All participants found the explanations of "beta coefficients" unclear, although they grasped the link to deviations in phenotypes. "Confidence intervals" remained unclear to most participants.

Participants were also asked questions to evaluate their understanding of the dashboard's views after tutorial completion. Their answers pointed at the Temporal Forest Plot view being the most difficult to comprehend. In the end, the participants were asked about what they found interesting in the tutorial. Several participants noted that they enjoyed interacting with the application and found parts of the content interesting. For example, *"It was interesting to see the changes in the Temporal Forest Plot"* [NE03], and *"I liked the interactivity"* [NE01]. However, they also criticized the complexity of the content and its length, noting *"some terms were challenging"* [NE05] and *"too complicated sometimes"* [NE04].

7 Reflections and Future Directions

We contributed a design study that explored the potential for a single visual analytics application to serve two disparate audiences through the intervention of visualization onboarding. To support our investigation we developed the *MoBa GWAS Explorer*, a web-based visual tool for exploring and analyzing GWAS data presented by the MoBa study for a mixed audience. Experts may use familiar idioms to visually explore new questions about the data, while non-experts have the same feature access in addition to an interactive onboarding tutorial. This tutorial explains how MoBa data are used and analyzed, walking users through the analysis process and providing explanations of essential genetic concepts.

While the tool shows potential for such mixed-audience user, our evaluation highlighted several limitations and areas for improvement. One key issue was the confusion expressed by experts when interacting with some parts of the application. This suggests that developing a tutorial for expert users could be a useful direction. A possible solution may be to readapt the existing non-expert tutorial to explain the offered visualizations and interactions to researchers, while omitting explanations of basic genomic terms and concepts that domain experts are familiar with. Another issue concerns the non-expert participants who did not launch the tutorial—due either to an expressed confidence in their ability to independently learn the application, or a wish to save time. These attitudes may stem from seeing the use of a tutorial as admitting a lack of knowledge on a certain topic, which may feel uncomfortable for some. Additionally, reluctance to engage with and invest time the tutorial may come from a lack of interest in the data—users may not immediately see personal relevance or value of the study. Engagement with such a tutorial may change with increased relevance or personal interest in the data. While the tool shows promise, some participants struggled to understand certain terms and concepts, for example, "confidence intervals." This raises the question of how we could redesign or reconsider the complexity of the text and tutorial to better facilitate understanding.

Designing visual analytics tools to support the different needs and interests of truly mixed audiences remains a relatively open research area. While our work investigated lightweight guidance and an onboarding tutorial strategy to reach non-expert users, many open and exciting questions remain to explore.

Supplemental Materials

Supplemental materials available at https://osf.io/k5bvj/.

Acknowledgments

We thank the researchers from the Norwegian Institute of Public Health and our study participants for their time and insights. We also thank the Univ. of Bergen VisGroup and our reviewers for their valuable suggestions.

REFERENCES

[1] Deoxyribonucleic acid (dna). https://www.genome.gov/genetics-glossary/Deoxyribonucleic-Acid. 2

[2] Genome. https://www.genome.gov/genetics-glossary/Genome. 2

[3] Genomic Variation — genome.gov. https://www.genome.gov/genetics-glossary/Genomic-Variation. 2

[4] How To Write Good Usability Testing Questions [2025 Guide]. https://careerfoundry.com/en/blog/ux-design/how-to-write-usability-testing-questions/, Dec 2022. 5

[5] C. Ahlberg and E. Wistrand. IVEE: An information visualization and exploration environment. In *Proc IEEE VIS*, pp. 66–73. IEEE Computer Society, USA, 1995. doi: 10/cw9dcn 4

[6] M. Bostock, V. Ogievetsky, and J. Heer. D3 data-driven documents. *IEEE Trans Vis Comput Graph*, 17(12), Dec. 2011. doi: 10.1109/TVCG.2011.185 4

[7] M. Burch and D. Weiskopf. On the benefits and drawbacks of radial diagrams. In *Handbook of human centric visualization*, pp. 429–451. Springer, 2013. doi: 10/ptv8 4

[8] A. Burns, C. Lee, R. Chawla, E. Peck, and N. Mahyar. Who do we mean when we talk about visualization novices? In *Proc. CHI*, pp. 1–16. Association for Computing Machinery, New York, NY, USA, 2023. doi: 10/gsfpjx 2

[9] W. S. Bush and J. H. Moore. Chapter 11: Genome-wide association studies. *PLoS Comput Biol*, 8(12):e1002822, 2012. doi: 10/gfr9pz 3

[10] D. Ceneda, T. Gschwandtner, T. May, S. Miksch, H.-J. Schulz, M. Streit, and C. Tominski. Characterizing guidance in visual analytics. *IEEE Trans Vis Comput Graph*, 23(1):111–120, Jan. 2017. doi: 10/f92d6c 4

[11] D. Ceneda, T. Gschwandtner, S. Miksch, and C. Tominski. Guided visual exploration of cyclical patterns in time-series. In *Proc VDS*, 2018. 4

[12] Z. Chen, M. Boehnke, X. Wen, and B. Mukherjee. Revisiting the genome-wide significance threshold for common variant gwas. *G3*, 11(2):jkaa056, Jan. 2021. doi: 10/ptwt 3

[13] J. Christiansen. *Building science graphics: an illustrated guide to communicating science through diagrams and visualizations*. AK Peters/CRC Press, 2022. doi: 10/gtmvqj 2

[14] Django Software Foundation. Djanjo project. https://www.djangoproject.com/, 2025. 4

[15] FinnGen. FinnGen results — r7.finngen.fi. https://r7.finngen.fi/, 2017. 2

[16] F. Frankel and A. H. DePace. *Visual strategies: A practical guide to graphics for scientists & engineers*. Yale University Press, 2012. 2

[17] L. A. Garrison, M. Meuschke, J. E. Fairman, N. N. Smit, B. Preim, and S. Bruckner. An exploration of practice and preferences for the visual communication of biomedical processes. In *Proc. VCBM*, pp. 1–12. Eurographics Association, France, 2021. doi: 10/n94j 2

[18] GeneViz. What is finngen? https://geneviz.aalto.fi/what-is-finngen/, 2024. 2

[19] S. B. Haga, W. T. Barry, R. Mills, G. S. Ginsburg, L. Svetkey, J. Sullivan, and H. F. Willard. Public knowledge of and attitudes toward genetics and genetic testing. *Genetic testing and molecular biomarkers*, 17(4):327–335, Apr. 2013. doi: 10/gnhpfm 3

[20] H. Hauser. Generalizing focus+ context visualization. In *Scientific visualization: The visual extraction of knowledge from data*, pp. 305–327. Springer, 2006. doi: 10/dh5vbj 5

[21] M. J. Khoury, T. H. Beaty, and B. H. Cohen. *Fundamentals of genetic epidemiology*, vol. 22. Monographs in Epidemiology and Biostatistics, 1993. 1

[22] B. C. Kwon and B. Lee. A comparative evaluation on online learning approaches using parallel coordinate visualization. In *Proc. CHI*, pp. 993–997, 2016. doi: 10/ghppzx 2, 5

[23] A. G. Lalkhen and A. McCluskey. Statistics v: Introduction to clinical trials and systematic reviews. *Continuing Education in Anaesthesia, Critical Care & Pain*, 8(4):143–146, Aug. 2008. doi: 10/dnmcbd 5

[24] S. LYi, Q. Wang, F. Lekschas, and N. Gehlenborg. Gosling: A grammar-based toolkit for scalable and interactive genomics data visualization. *IEEE Trans Vis Comput Graph*, 28(1):140–150, Jan. 2021. doi: 10/gr2d8n 2

[25] R. Mägi and A. P. Morris. GWAMA: software for genome-wide association meta-analysis. *BMC Bioinformatics*, 11(288):1–6, May 2010. doi: 10/c2v2vv 2

[26] P. Magnus, C. Birke, K. Vejrup, A. Haugan, E. Alsaker, A. K. Daltveit, M. Handal, M. Haugen, G. Høiseth, G. P. Knudsen, et al. Cohort profile update: the norwegian mother and child cohort study (MoBa). *Int J Epidemiol*, 45(2), Apr. 2016. doi: 10/f8pdwv 1

[27] D. McCandless. The microbescope – infectious diseases in context. https://informationisbeautiful.net/visualizations/the-microbescope-infectious-diseases-in-context/, 2021. 2

[28] M. Meuschke, L. A. Garrison, N. N. Smit, B. Bach, S. Mittenentzwei, V. Weiß, S. Bruckner, K. Lawonn, and B. Preim. Narrative medical visualization to communicate disease data. *Computers & Graphics*, 107:144–157, Oct. 2022. doi: 10/gt25bm 2

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

[30] Q. V. Nguyen and M. L. Huang. A focus+ context visualization technique using semi-transparency. In *Proc. CIT*, pp. 101–108. IEEE Computer Society, USA, 2004. doi: 10/c69r4b 4

[31] S. Nusrat, T. Harbig, and N. Gehlenborg. Tasks, techniques, and tools for genomic data visualization. In *Comput Graph Forum*, vol. 38, pp. 781–805, 2019. doi: 10/ggqfk5 2

[32] J. O'Daniel and S. Haga. Public perspectives on returning genetics and genomics research results. *Public Health Genomics*, 14(6):346–355, May 2011. doi: 10/fj6724 3

[33] L. Pauwels. *Visual cultures of science: rethinking representational practices in knowledge building and science communication*. UPNE, 2006. 2

[34] J. Peng, H. Xue, W. Hui, J. Lu, B. Chen, Q. Jiang, X. Shang, and Y. Wang. An online tool for measuring and visualizing phenotype similarities using hpo. *BMC genomics*, 19(571):89–97, Aug. 2018. doi: 10/ptv7 2

[35] P. Rautek, S. Bruckner, E. Gröller, and I. Viola. Illustrative visualization: new technology or useless tautology? *ACM SIGGRAPH Computer Graphics*, 42(3):1–8, Aug. 2008. doi: 10/fmjcxt 2, 5

[36] E. Segel and J. Heer. Narrative visualization: Telling stories with data. *IEEE Trans Vis Comput Graph*, 16(6):1139–1148, Nov. 2010. doi: 10/dnqx5m 2, 5

[37] C. Stoiber, F. Grassinger, M. Pohl, H. Stitz, M. Streit, and W. Aigner. Visualization onboarding: Learning how to read and use visualizations. Aug. 2019. doi: 10/gh38zd 2, 5

[38] C. Stoiber, C. Walchshofer, F. Grassinger, H. Stitz, M. Streit, and W. Aigner. Design and comparative evaluation of visualization onboarding methods. In *Proc International Symposium on Visual Information Communication and Interaction*, pp. 1–5. Germany, 2021. doi: 10/gpngs3 2, 5

[39] H. Taherdoost. Sampling methods in research methodology; How to choose a sampling technique for research. Apr. 2016. doi: 10/ggskz3 6

[40] Y. Tanahashi, N. Leaf, and K.-L. Ma. A study on designing effective introductory materials for information visualization. In *Comput Graph Forum*, vol. 35, pp. 117–126, 2016. doi: 10/f9c3ns 2, 5

[41] J. Trumbo. Visual literacy and science communication. *Science communication*, 20(4):409–425, June 1999. doi: 10/bw82ns 2

[42] S. D. Turner. qqman: an r package for visualizing gwas results using qq and manhattan plots. *Biorxiv*, p. 005165, May 2014. doi: 10/gkv9qp 2

[43] UK Biobank. UK Biobank - UK Biobank — ukbiobank.ac.uk. https://www.ukbiobank.ac.uk/, 2006. 2

[44] I. Viola, E. Gröller, M. Hadwiger, K. Bühler, B. Preim, M. Costa Sousa, D. Ebert, and D. Stredney. Tutorial 4: Illustrative visualization. *Visualization, VIS*, 5:1–188, 2005. 2, 5

[45] S. Wang, D. Dvorkin, and Y. Da. Snpevg: a graphical tool for gwas graphing with mouse clicks. *BMC bioinformatics*, 13(319):1–6, Nov. 2012. doi: 10/gb8vn4 2

[46] K. E. Zhang, M. Vaudel, and L. A. Garrison. The manhattan wheel: A radial visualization story for genome-wide association study data, 2024. 2