




Marjorie: Visualizing Type 1 Diabetes Data to Support Pattern Exploration

Anna Scimone , Klaus Eckelt , Marc Streit , and Andreas Hinterreiter 

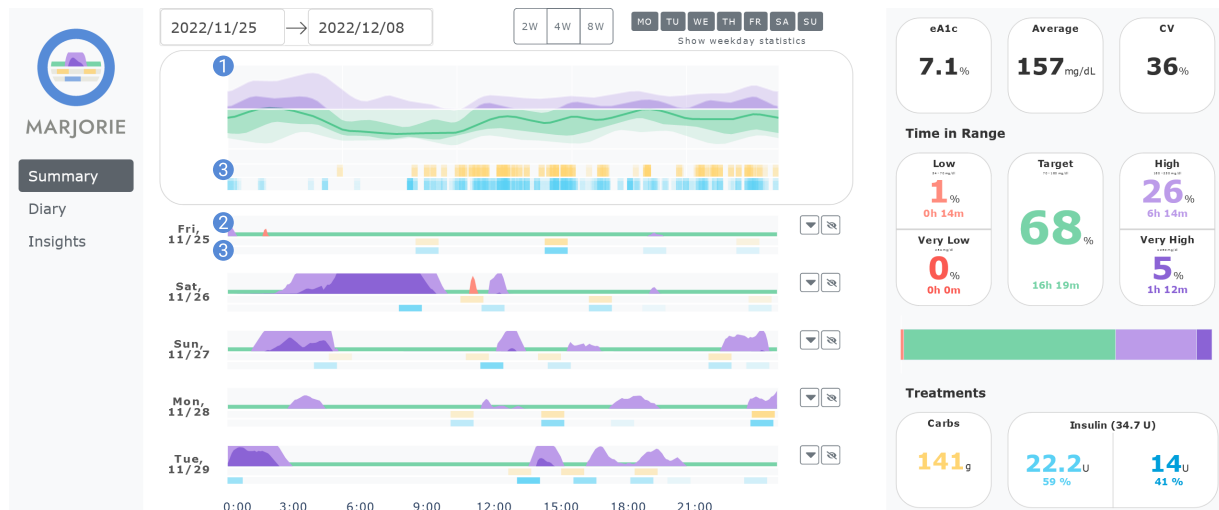


Fig. 1: User interface of Marjorie showing the overview visualization for a user-specified time range in the panel in the middle. The AGP plot at the top summarizes glucose values of the whole time range (1). Underneath, adapted horizon graphs present the data of distinct days (2). Both visualizations show data on carbohydrate and bolus insulin intake juxtaposed (3). The statistics panel on the right summarizes the most relevant metrics and additionally serves as a legend.

Abstract—In this work we propose Marjorie, a visual analytics approach to address the challenge of analyzing patients' diabetes data during brief regular appointments with their diabetologists. Designed in consultation with diabetologists, Marjorie uses a combination of visual and algorithmic methods to support the exploration of patterns in the data. Patterns of interest include seasonal variations of the glucose profiles, and non-periodic patterns such as fluctuations around mealtimes or periods of hypoglycemia (i.e., glucose levels below the normal range). We introduce a unique representation of glucose data based on modified horizon graphs and hierarchical clustering of adjacent carbohydrate or insulin entries. Semantic zooming allows the exploration of patterns on different levels of temporal detail. We evaluated our solution in a case study, which demonstrated Marjorie's potential to provide valuable insights into therapy parameters and unfavorable eating habits, among others. The study results and informal feedback collected from target users suggest that Marjorie effectively supports patients and diabetologists in the joint exploration of patterns in diabetes data, potentially enabling more informed treatment decisions. A free copy of this paper and all supplemental materials are available at <https://osf.io/34t8c/>.

Index Terms—Design study, task analysis, diabetes, time series data, visual analytics, clustering

1 INTRODUCTION

Type 1 diabetes is a chronic autoimmune disease with a prevalence of over 9 million cases globally [18]. As part of their therapy, patients need to closely monitor their blood glucose levels and administer finely tuned doses of insulin. The adjustment of therapy parameters requires a thorough understanding of past situations and patterns in the patients' blood glucose levels. To this end, patients meet regularly with their diabetologists to jointly analyze the recorded blood glucose data [22].

From a data analysis and visualization perspective, type 1 diabetes data poses interesting challenges. First, diabetologists typically have little time. Therefore, patient–physician meetings must be kept short and important patterns must be identified and understood efficiently.

- All authors are with Johannes Kepler University Linz.
- E-Mail: anna_scimone@live.de; klaus.eckelt@jku.at; marc.streit@jku.at; andreas.hinterreiter@jku.at.

Manuscript received xx xxx. 201x; accepted xx xxx. 201x. Date of Publication xx xxx. 201x; date of current version xx xxx. 201x. For information on obtaining reprints of this article, please send e-mail to: reprints@ieee.org. Digital Object Identifier: xx.xxx/TVCG.201x.xxxxxxx

Second, the patient data combines regular or continuous measurements of blood glucose levels (i.e., time series data) with potentially irregular data about meals and exercise (i.e., event data) [19]. Third, specific patterns, such as periods of hypoglycemia, where the glucose levels fall below the normal range, require increased attention as they pose serious health risks [22]. Fourth, new visualization approaches have to be compatible with standardized visualizations that have been established for certain subtasks of analyzing diabetes data [9, 22].

Existing visualization tools for type 1 diabetes data focus on displaying glucose time series but struggle to effectively visualize additional attributes (such as meal or insulin data) without losing clarity. Understanding the interactions between the different attributes requires an integrated visualization of all attributes in a simple overview, which allows users to drill down and identify patterns on varying temporal scales. In this work we introduce Marjorie, a visual analytics solution that addresses these challenges. We named our tool in honor of the experimental dog that Frederick Banting and Charles Best first successfully treated with insulin, enabling clinical trials thereafter [7, 38].

The main contribution of this work is *Marjorie*, a visual analytics solution for analyzing diabetes data. Marjorie includes (i) a compact

visualization of blood glucose levels based on horizon graphs with a range of baselines; (ii) a concise representation of aggregated carbohydrate and bolus insulin data, which combines hierarchical clustering with semantic zooming; and (iii) the automatic identification of characteristic glucose patterns occurring during periods of hypoglycemia or mealtimes, enabled by dynamic time warping and hierarchical clustering. Secondary contributions of our work are a detailed description of tasks involved in the visual analysis of type 1 diabetes data, developed in consultation with diabetologists; a case study with real patient data; and a discussion of feedback by diabetologists and patients.

2 BACKGROUND

The term *diabetes mellitus* encompasses two types of diseases that differ in clinical course and treatment. The immune systems of type 1 diabetes patients attack and destroy the β -cells in the pancreases. The destruction of the β -cells, which are responsible for insulin production, leads to a lack of insulin in the patients' bodies. This chronic autoimmune disease usually manifests in childhood or early adulthood, and requires lifelong treatment. Type 2 diabetes, in contrast, is a metabolic disorder characterized by insulin resistance and insufficient insulin production. It usually develops in adults and is often associated with obesity, inactivity, and unhealthy eating habits. Our work focuses on type 1 diabetes data. Hence, we use the term *diabetes* throughout our paper to refer to type 1 diabetes.

Patients with diabetes need to keep their blood glucose levels in a target range most of the time to avoid long-term complications such as permanent damage to organs and nerves [33]. Deviations, both towards low blood glucose (hypoglycemia, colloquially referred to as “hypo”) and high blood glucose (hyperglycemia, colloquially “hyper”), should be kept to a minimum. The patients' glucose levels can be affected by various factors. The most common and predictable influence factors are food and insulin intake: food intake—especially carbohydrate intake—increases the glucose levels, while administered insulin decreases them [6]. Other influencing factors that can influence glucose levels are physical exercise, alcohol intake, menstruation, stress, and infections [6].

A patient's therapy involves regular monitoring of glucose levels and administration of insulin. The glucose levels can be measured either by a blood glucose meter or with a continuous glucose monitoring (CGM) system. A blood glucose meter is a device that measures blood sugar levels at a single point in time. A CGM system, in contrast, measures glucose levels continuously through a sensor placed under the skin. Patients also need to administer insulin on a regular basis, either through injections (using a syringe or insulin pen), or with an insulin pump. An insulin pump is a small device worn on the body, which delivers insulin continuously through a cannula. The pump is programmed to deliver small amounts of insulin (known as “basal” insulin) throughout the day, mimicking the natural “background” insulin production of the pancreas. When the patient consumes carbohydrates, an additional dose of insulin, known as “bolus insulin”, is required. The dose is determined based on the amount of carbohydrates and the patient's current glucose levels.

Appropriate adjustment of therapy parameters requires a thorough analysis of past situations, which is usually done in regular time intervals with consultation of a diabetologist. The treatment parameters vary depending on the time of day, as the body's insulin sensitivity changes in a circadian rhythm. Parameters may also need to be based on the patient's changing needs and response to treatment. The following are the most important therapy parameters: (i) the *basal rate* is the programmed basal insulin dose; (ii) the *carbohydrate/insulin factor* is the amount of carbohydrates covered by one unit of insulin; (iii) the *prebolus time* is the time between administering insulin and eating a meal; (iv) the *insulin sensitivity factor* is the amount by which one unit of insulin decreases the level of glucose—it is influenced, among others, by physical activity, stress, and illness. Patients record their glucose, carbohydrate, and insulin data either manually, or automatically if they use a CGM sensor and insulin pump. At regular intervals, typically every three months, they meet with their diabetologist to review the collected data and assess the quality of their therapy. Based on the results, the treatment parameters are adjusted to optimize glucose levels.

When physicians and patients analyze CGM data, they typically use the standardized Ambulatory Glucose Profile (AGP). The AGP is a visualization developed by an international panel of diabetes clinicians and researchers [9], which is also featured in the consensus report by the American Diabetes Association [22]. It supports the exploration of seasonal patterns by aggregating glucose readings of multiple days or weeks across a 24-hour period using the 10th, 25th, 50th, 75th, and 90th percentiles.

3 RELATED WORK

In this section, we first review relevant literature on the analysis of time series and event sequence data. We then discuss works specifically addressing diabetes data. Given our approach's utilization of pattern detection and time series clustering, we also explore related works in these areas.

3.1 Time Series and Event Sequence Visualization

The patient data analyzed by our approach combines time series data (i.e., continuous measurements of blood glucose levels) with event data (e.g., exercise or meals) [19]. Time series and event sequence visualizations naturally play an important role in the analysis of electronic health records [36, 45]. Visual representations of time series can make it easier for users to find trends, patterns, and anomalies in the data by taking advantage of people's natural ability to process and interpret information [32]. Based on our interviews with diabetologists (see Sec. 4.1), we focus on techniques that feature a dedicated time axis. We thus disregard approaches for time series visualization that rely on implicit temporality along trajectories [3, 11, 16, 21, 40].

Following the categorization by Aigner et al. [2], we reviewed two-dimensional static visualizations for abstract uni- or multivariate data in the TimeViz Browser [44]. We consider point plots, connected scatterplots and line plots, area graphs, heatmaps, and horizon graphs [35] most suitable for visualizing diabetes data. The grouping of blood glucose values into ranges in the AGP [9, 22] further relates diabetes data visualization to works that use qualitative ranges.

Bade et al. [4] visualize time series data with a qualitative scale and use a semantic zoom approach to transition from compact heatmaps to area diagrams and line diagrams based on the available space. In VisuExplore [37] data of patients with chronic diseases, such as diabetes, are shown in multiple views that share a common time axis and adapt to the available space using Bade et al.'s semantic zooming approach [4].

Qualizon Graphs [13] extend the idea of horizon graphs to use qualitative ranges, instead of equally sized areas, that are centered around a baseline. In our approach, we adapt Qualizon Graphs such that one of the qualitative ranges (i.e., the target range) serves as the baseline. Changes within the target range are not of primary interest and can thus be omitted.

3.2 Diabetes Data Visualization

Diabetes-related data from wearable devices is commonly viewed directly inside the web, phone, or desktop applications of the manufacturers [1, 10, 28, 29]. Furthermore, several integrated platforms allow consolidation of data from different sources and variables. Examples of such platforms are the commercial application *Glooko* [17] and the open source applications *Nightscout* [41] and *Tidepool* [43, 47]. The review of visualization approaches for diabetes management by Zhang et al. [50] revealed that only a few tools support CGM data. Additionally, some of the few tools that display carbohydrate and bolus insulin data in their daily detail views do not visually represent the quantity of these attributes; Tidepool [43] and IDMViz [49] only show the quantity of bolus insulin, while indicating the time of carbohydrate consumption with glyphs. Other tools, such as Nightscout [41], suffer from overplotting, making it difficult to read the quantities.

For identifying seasonal daily patterns and gaining insight into a patient's typical daily glucose profile, the AGP is a well-established visualization in the diabetes domain (see Sec. 2). However, since the AGP does not take into account carbohydrate and insulin data, it does not allow physicians to determine the factors contributing to a patient's glucose profile. As a result, physicians must rely on daily detail views

to gain a comprehensive understanding of the relationship between glucose patterns and treatment habits.

The review by Zhang et al. [50] also describes temporal folding strategies for exposing recurrent patterns in the data. This folding typically allows users to switch between daily, weekly, bi-weekly, or monthly views of the data. To identify non-seasonal patterns, a common approach is to align the time series data by important events to display patterns of precursor, co-occurring, and aftereffect events [50]. IDMVis [49] additionally allows users to align multiple time series between two events by justifying or stretching them.

In contrast to existing tools, our approach supports automatic clustering of meal and hypoglycemia patterns, and the AGP is juxtaposed with carbohydrate and insulin data.

3.3 Pattern Detection

Patterns in temporal data can be classified as *trend*, *seasonal*, or *cyclic* patterns [24]. *Trend* patterns mark a long-term increase or decrease in the data [24], e.g., in the context of diabetes data, a long-term improvement or deterioration of therapy. *Seasonal* patterns occur at an always fixed and known frequency, such as higher blood glucose levels on weekends or lower blood glucose levels during summer, whereas *cyclic* patterns are repeated sequences without a fixed frequency [24], for example, hyperglycemia after a meal or hypoglycemia during exercise.

The wearable device manufacturers *Dexcom* (“Patterns Report” [10]), *Medtronic* (“Pattern Snapshot” [29]) and *Abbott* (“Glucose Pattern Insights” [1]) include sections in their applications, where non-periodic patterns are shown. Patterns such as “nighttime highs” or “pre-breakfast highs” are listed, along with suggestions for possible causes and considerations. Additionally, *Dexcom* shows a visualization of the found patterns by highlighting the respective sequences in the daily view.

The algorithms used by these diabetes device manufacturers are not publicly available. However, various algorithms have been developed and tested for diabetes data analysis in research papers. Woldaregay et al. [46] performed a comprehensive review of machine learning methods for discovering patterns in diabetes data, with a focus on works that detect glucose anomalies, predict glucose levels, and/or classify hypoglycemia, hyperglycemia, and glycemic variability. Hall et al. [20] classify the magnitude and degree of glucose variability by segmenting CGM data into sliding windows and then using dynamic time warping and spectral clustering to classify the windows into clusters of low, moderate or severe glucose variability. Lobo et al. [27] classify single days of CGM data by matching them to pre-defined representative daily profiles. *PSEUDO* [48] allows to visually search multivariate time series data for patterns similar to a user-defined example. To identify patterns visually, the *Sequence Braiding* technique [5] gives an overview of blood glucose values of several days along important events.

In our approach, rather than clustering whole days, we focus on specific time ranges around meals or episodes of hypoglycemia. We also do not rely only on glucose variability, but make use of the detailed glucose data.

4 PROBLEM CHARACTERIZATION AND ABSTRACTION

In this section, we lay a foundation for our design study by characterizing the problem on a domain level, abstracting it into visualization tasks and analyzing the occurring data types.

4.1 Domain Analysis

Our domain analysis is based on contextual inquiry interviews [23] and think-aloud protocols [8]. We conducted seven 90-minute sessions with two diabetologists (*D1* and *D2*). For our domain analysis, we showed the physicians a sample patient dataset inside the Nightscout diabetes visualization tool [41] and asked them to explore the data, while communicating their thoughts and insights. We supplemented the interview data with domain literature on the management of type 1 diabetes in adults [22, 25].

The target users addressed within this design study are diabetologists and their type 1 diabetes patients who wear a CGM sensor and an insulin pump. The objective of a diabetologist is to assist patients in managing their disease. They do this by educating patients about

patterns in their diabetes data and adjusting treatment parameters to maintain blood glucose levels within a safe range. The appointments are usually repeated in three-month intervals and typically last between 20 and 40 minutes. During the sessions, the diabetologist loads the data from different devices into a software and observes the metrics and visualizations.

The primary challenge during the analysis, according to both interviewed diabetologists, is to quickly gain an overview and make sense of the vast amount of data within a short time frame. As *D1* stated, “*The sheer amount of data is overwhelming, and it’s impossible to review it all in the limited time frame. This can make patients lose interest as they often don’t understand the complicated diagrams [...]*.” *D2* expressed similar thoughts: “*Normally, I only pick 1 to 2 situations to go through with the patient; we can’t handle more during that time frame anyway.*”

Current Workflow of Experts During the interviews we observed that both diabetologists currently proceed similarly during the analysis of patient data, acting in accordance with the officially standardized workflow of data interpretation suggested by expert physicians and researchers [25]. *Table 1* describes this typical analysis workflow with state-of-the-art diabetes visualization tools, focusing on the *What* (analysis objective), *Where* (specific visual component), and *How* (exact procedure). Notably, the physicians predominantly focus on three components of the tools: metrics, the AGP, and the daily detail views (see *Sec. 3*). They mentioned that these were also the most standard components contained in all tools.

In summary, diabetologists search for recurring daily patterns of hypoglycemia, hyperglycemia, and high glucose variability within the AGP, and investigate the daily detail views to find potential weekly patterns. After having identified critical blood glucose patterns, they further analyze the daily views to uncover potential causes. They particularly focus on the recorded carbohydrate intake and insulin doses, as well as the timing between them. This information is then used to optimize the patient’s treatment plan (e.g., by adapting one of the treatment parameters described in *Sec. 2*).

Requirements Based on the insights from our domain analysis, we outlined five design requirements for our visual analytics approach.

R1: Make periods of low and high glucose easy to identify.

As can be derived from *Tab. 1*, the primary focus in the analysis session are periods of blood glucose outside the target range, namely hypoglycemia, hyperglycemia, and high blood glucose fluctuation (alternating hypo- and hyperglycemia within short intervals). Swift identification of these segments in the dataset is crucial for the analyst.

R2: Support finding recurring daily and weekly patterns.

After segments of hypo- and hyperglycemia have been identified, the diabetologist usually browses the AGP and daily detail views to make out seasonal patterns such as daily or weekly recurrences in the data (see *Tab. 1*). These patterns are easy to identify and efficient to eliminate, as they are usually caused by unreflected habits of the patients. *D1*: “*For example, I had a patient who always had high blood sugar levels after lunch and gave a big amount of correction insulin every day to get back into the target range. So I suggested him to adjust the insulin factor for lunch and give higher doses right away.*”

R3: Visualize all variables in one place.

Another requirement is to collect and visualize multivariate and heterogeneous datasets in one place. When asked about the limitations of current diabetes data analysis tools, the diabetologists stated that they require a comprehensive view of both glycemic data and additional data such as meal and insulin information, to better understand the connections between blood glucose fluctuations and potential causes. *D1*: “*In the AGP there is only glucose data, and I have to click back and forth 10 times to understand why the patient, for example, always has high values in the morning. I always go into the daily views to see what he ate around that time.*” Moreover, the diabetologists expressed the wish to load more attributes into the tool. *D2* stated: “*Additional health data such as exercise or sleep would really help to understand the blood sugar’s behavior and to advise patients how to handle specific situations better.*” *D1* wished to see more detailed nutrition data.

Table 1: Procedure when analyzing diabetes data, summarized from expert interviews and complementing domain literature.

What	Where	How
Assess Overall Therapy Quality	Metrics	Is the eA1c < 7%? Does the time in range (TIR) lie above 70 %? (D2)
	AGP	Observe last 2 weeks' AGPs, paying attention to nighttime values and low glucose. Identify stable/unstable glucose time frames and compare with TIR to detect inconsistencies. (D1 & D2)
Analyze Patterns of Hypoglycemia	Metrics	Is the time below 70 mg/dL > 4%? Is the time below 54 mg/dL > 1 %? If yes, immediate action is required. (D2; [22])
	AGP	Observe the lower area of the AGP, where the 5 % or 25 % lower line touches the 70 mg/dL target line indicating a (severe) low glucose. Immediate action is required. (D1 & D2; [25])
	Daily Views	Double-check low glucose patterns. [25] Look for potential causes, e.g., meals, insulin dosing and timing, exercise etc. (D1 & D2) Check for weekly seasonal patterns of low glucose, e.g., clusters on weekends or special days. [25]
Analyze Patterns of Hyperglycemia	AGP	Where do the AGP areas rise above range? (D1 & D2)
	Daily Views	Observe the relation between high glucose time periods and meals. What was the timing of insulin administration (before/after meal)? Has insulin been forgotten? (D2) Check for weekly seasonal patterns of low glucose, e.g., clusters on weekends or special days. [25]
Analyze Patterns of Glucose Variability	Metrics	Is the standard deviation > 40 mg/dL? (D2)
	AGP	Where are the AGP areas particularly wide? [25]
	Daily Views	Observe timing of food intake and insulin, periods of exercise. [25]

R4: Bundle and compare similar data.

The ability to collect similar situations of the patient and make them comparable helps the physician to identify recurring behavior that leads to blood sugar imbalances. D2 stated: “A single area to analyze sport days would be great for patients; many wonder how they can better regulate their blood sugar during these activities. For example, if they tend to end up in a hypoglycemia two hours after sports, I would suggest them to eat a banana in between.” In another statement, D2 also said: “It would be really great if it was possible to collect and compare all instances of hypoglycemia in order to understand the different reasons contributing to them.”

R5: Display a detailed view of specific days on request.

The diabetologists mentioned that patients sometimes remember a specific day that they want to analyze further. D1: “Recently a patient told me about his half marathon in April. He had low blood sugar for the entire two hours and had to constantly eat while running. We looked up the day and found out that he had lunch only 1.5 hours before the run. With foresight on the exercise, he had reduced the insulin amount for the meal, but obviously not enough for such a strenuous exercise.”

4.2 Data Abstraction

The attributes we incorporated into our approach encompass blood glucose levels, carbohydrate consumption, and insulin (both bolus and basal insulin). Sleep, nutrition, and exercise data requested by the diabetologists (see R3) are not included due to the difficulty in obtaining such information. Figure 2 displays the categorization of the attributes into the design aspects of time by Aigner [2].

Blood Glucose Levels The blood glucose levels stem from continuous glucose monitor (CGM) systems, wearable devices that provide point-based real-time glucose readings every 5 minutes, yielding approximately 288 readings per day. Each CGM data point consists of a quantitative value and an associated point-based timestamp. The optimal range of glucose values is usually between 70 mg/dL and 180 mg/dL [26]. The blood glucose data can be aggregated and mapped to different granularities, such as days, weeks or months.

Carbohydrate Consumption Patients manually enter the estimated carbohydrate intake in grams (g), making it more susceptible to inaccuracies compared to other attributes. The data consists of discrete log entries, including a quantitative value indicating the amount of carbohydrates and an associated timestamp.

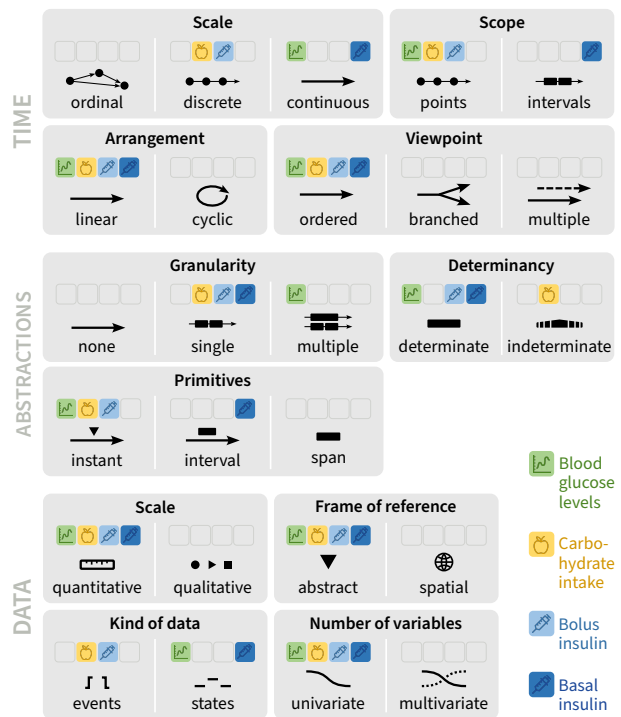


Fig. 2: Categorization of data attributes visualized in Marjorie into the design aspects of time by Aigner et al. [2].

Bolus Insulin Insulin data is automatically logged by insulin pumps, and bolus insulin is required to prevent blood glucose from rising after meals. Similar to carbohydrate data, bolus insulin is recorded in event logs, including quantitative values measured in international insulin units (U) with corresponding timestamps. From bolus insulin data, the continuous insulin activity curve can be calculated using a formula to estimate insulin metabolism in the body [42].

Basal Insulin Basal insulin is required to counteract the liver's constant glucose release and maintain stable blood sugar levels. The basal insulin dose fluctuates in a circadian pattern, with a roughly equal

progression each day. Insulin pumps are programmed with a continuous basal rate expressed in U/h, varying at hourly intervals. Users can configure temporary basal rates to adjust for activities like exercise, which enhances insulin sensitivity. The modified basal rate is expressed in event logs containing a start time, duration, and a percentage value of the original basal rate in the pump data.

4.3 Task Abstraction

After the domain analysis, we translated the domain-specific problems of the experts into abstract visualization tasks based on Munzner's taxonomy [31]. This resulted in five tasks that correspond to the requirements R1–R5. We refined and confirmed the tasks through a series of iterations by consulting the domain experts.





- T1 Identify extreme values** of blood glucose data (i.e., periods of hypo- and hyperglycemia).
- T2 Identify seasonal patterns.**
 - a Aggregate** the blood glucose raw data **by temporal granularity.**
 - b Identify extreme values in aggregated blood glucose data** to identify specific time granules of interest.
- T3 Compare multiple attributes** of the dataset (i.e., glucose, carbohydrates, and insulin) to locate potential correlations.
- T4 Summarize and compare multiple similar time segments.**
 - a Cluster similar time segments** of the raw data.
 - b Summarize similar time segments** within a cluster.
- T5 Lookup time segments of interest** by specifying the relevant date.

5 MARJORIE APPROACH

Marjorie is a visual analytics approach for type 1 diabetes data that aims to support experts in their search for seasonal and non-periodic patterns. It uses a combination of visual and algorithmic approaches discussed in Sec. 5.1 and Sec. 5.2, respectively.

5.1 Visual Encoding & Interaction Design

Our approach has been iteratively refined with input from the two diabetologists at various stages of development. During the sessions, we presented our observations and insights in the form of written notes, sketches, mockups, and working software prototypes, seeking feedback from the specialists. Our resulting web app prototype was developed using the *Plotly Dash* library [34]. User data can be loaded into the app as CSV files. The prototype is available at <https://marjorie.jku-vds-lab.at/>.

Color Encoding We use color to differentiate attributes and indicate glucose level ranges consistently. As there is no consensus on color coding for blood glucose value ranges [50], we based our scheme on that of Tidepool [43] with modified glucose hues and additional intermediate shades. Despite the continuous nature of glucose values, we use a categorical color scheme to clearly distinguish different blood glucose levels: . Red represents hypoglycemia, green represents the normal range, and purple represents hyperglycemia. Different color intensities represent the severity of each state. Furthermore, yellow  indicates carbohydrates, while insulin is depicted in blue, with bolus insulin using a lighter shade  and basal insulin a darker shade of blue . The statistics panel to the right of the summary view serves as a legend for this color scheme. The glucose color scheme was tested for compatibility with color blindness, and the ranges for hypoglycemia, normal range, and hyperglycemia are visually distinguished [15]. All other colors are used only in isolation, and titles and tooltips provide additional aids to differentiate and identify the visualized data types.

Layout Figure 1 shows an overview of Marjorie's layout, which is divided into three separate views. The first view, *Summary*, provides a comprehensive summary of the therapy quality over the preceding weeks and helps to detect any daily seasonal glucose patterns (see Fig. 3). It contains the patient's **AGP** (Sec. 5.1.1), overall **Statistics** (Sec. 5.1.3), and an **Overview** (Sec. 5.1.2) of multiple days at a glance. The second view, *Diary*, provides the **Daily View** (Sec. 5.1.4),

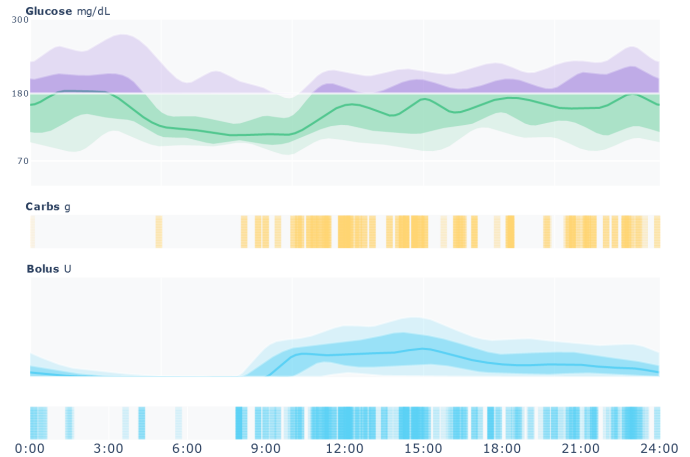

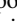
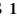



Fig. 3: The summary visualization shows the data for a user-defined time range with an AGP plot for glucose values (hypo  hyper ) and barcode plots for carbohydrate ) and bolus insulin ) intake. Additionally, insulin activity is shown as a curve.

a detailed visualization of specific days that can be selected and queried using a date picker, along with statistics related to that particular day. The third view provides **Insights** (Sec. 5.1.5) regarding non-periodically occurring glucose profiles in the data. It consists of two tabs for *Meals* and *Hypos*. In each tab, the user can explore and compare the glucose profiles of the corresponding situation.

5.1.1 AGP

The AGP view (Fig. 3) extends the original AGP visualization [9, 22]. It supports T2 (identify seasonal patterns) by letting the user explore hourly patterns. Like the original AGP, it aggregates the ranges between the 10th and 90th, and the 25th and 75th percentiles of two weeks of CGM data, supporting T2a (aggregate by temporal granularity). The aggregated data is displayed as superimposed area charts across a 24-hour period. On top, the 50th percentile is visualized as a line graph. To support T2b (identify extreme values in aggregated data), the AGP highlights areas that fall outside the target glucose range in different colors, according to the color scheme described above.

We enhance the original AGP by adding carbohydrate and bolus insulin data, supporting T3 (compare multiple attributes), by adding three vertically juxtaposed charts sharing the same time axis. We use two barcode plots for carbohydrate and insulin administration events, where the position on the horizontal axis encodes the timing of the events, and the opacity of the bars corresponds to the amount. This enables the user to identify time periods during which the frequency and/or amount of carbohydrate or insulin consumption was high. We chose barcode plots because they use little screen space while conveying most of the important information. As a third chart, we added a qualitative line chart visualization of the insulin activity in the bolus section to simplify the search for correlations with the glucose data.

Previous Versions. Our first AGP draft presented to the diabetologists did not include the insulin activity curve. D1 mentioned that the barcode plots alone did not provide enough information, as the effects of insulin on the body can last up to eight hours and accumulate with each bolus dose. To gain a better understanding of how insulin and glucose interact, they requested a visualization of the insulin activity.

In addition, our initial version of the AGP had buttons that allowed users to aggregate the data by different granularities aside from hours, such as by weekday, month, or quarter. The software displayed the results in the form of bar charts (see Appendix). The diabetologists' feedback was that they were most interested in seeing the data on an hourly basis and that they did not need to see aggregations by different time granularities, except for the weekday view. Therefore, we made the weekday option accessible separately through a *Show weekday statistics* button and removed the other aggregation options.

5.1.2 Overview

By clicking on the *Explore days in detail* button located below the AGP, the size of the AGP is reduced, and a detailed visualization of the days used to calculate it is presented below (see Fig. 1). The data is partitioned into distinct days which are juxtaposed underneath each other, sharing the same time axis. The glucose data is represented using adapted horizon charts, where the baseline is not a single value, but instead the target range of 70 mg/dL to 180 mg/dL. The values outside the target range are displayed like in typical horizon charts, while the values within the range are shown as a green band. This makes it simple to pinpoint days and times that require closer examination, helping to achieve T1 (identify extreme values). Clicking the right arrow () expands the graph to show a detailed line plot of the glucose data.

The carbohydrate and bolus insulin data are presented in a heatmap below the glucose data, addressing T3 (compare multiple attributes). It often happens that event entries occur within short time frames, for instance when a patient administers multiple insulin corrections within half an hour to address high blood glucose levels. To address this, we cluster and summarize these values and display the median time point of consumption or administration in the heatmap. The opacity of the field encodes the total amount of carbs or insulin administered. More detailed information on the distinct events is available in a tooltip window. We decided to represent the quantity using opacity instead of length (such as in bar charts) to save vertical screen space in favor of displaying a larger number of days.

Previous Versions. Earlier drafts included glucose line plots instead of horizon graphs, superimposed with differently colored semi-transparent bar charts that encoded the amount of carbohydrates and bolus insulin. However, with this choice of superimposing the data to save screen space, diabetologists found it difficult to get an overview of the data, particularly when the patient recorded multiple events within a short timeframe. In such cases, the bars would overlap, causing confusion about the amount of insulin and food consumed. For example, D2 asked, “How much insulin did the patient inject here? The bars are all on top of each other. [...] Is this food or insulin?”, when presented with overlapping blue (insulin) and yellow (carbs) bars.

5.1.3 Statistics

To provide a quick overview of the therapy quality, we included a statistics sidebar (see Fig. 1). It is visible in the AGP, overview, and daily views. In the AGP and overview views, it shows statistics for a selected time range, while in the daily view, it displays statistics for the selected day. From our interviews with diabetologists, we discovered that the metrics they find most significant are the estimated hemoglobin A1c level (eA1c), average glucose, coefficient of variation, time in range (TIR), and the daily sums of consumed carbohydrates, bolus and basal insulin. The statistics sidebar also serves as a legend for the colors used in Marjorie’s charts.

Previous Versions. A discarded design for aggregated daily data was the timeline graph (see Appendix), which showed the time in range in the form of stacked bar charts, together with heatmaps for the aggregated sum of carbohydrates and bolus insulin on a weekly basis for the entire dataset. Range sliders allowed the user to select the time frame for calculating the AGP. However, the diabetologists expressed that they were unlikely to use it. D1 stated: “I am not interested in older data. It doesn’t make sense to analyze it because the patients naturally only remember the last few weeks.” D2 mentioned: “What I am missing in this visualization are the concrete numbers. I always check the time in range values, especially the time in hypoglycemia.”

5.1.4 Daily View

The daily view (Fig. 5) is designed to support both T5 (lookup time segments of interest) and T3 (compare multiple attributes). Users can select a specific day with a date picker, and the view displays juxtaposed glucose, carbohydrate, bolus, and basal data arranged on a single time axis. Glucose is presented as a scatter plot, while carbohydrates and bolus insulin are presented by bar charts, and basal insulin by an area plot. To avoid overplotting with overlapping bars, we cluster the

carbohydrate and bolus insulin data in the same way as the data in the overview. When hovering over a bar, a tooltip window displays the distinct events that are summarized within it. We incorporated a semantic zoom feature that adjusts the level of summarization of carbohydrate and bolus insulin data based on the zoom level. Upon zooming into the graph, the summarized bars split up into their components, providing more details. In Sec. 5.2, we elaborate on the details of the algorithm.

5.1.5 Insights

The *Insights* view addresses T4 (summarize and compare multiple similar time segments). Here, situations from a selected category (*Meals* or *Hypoglycemia*) are clustered by glucose similarity (T4a; cluster similar time segments) and presented to the user in the form of typically occurring glucose profiles (T4b; summarize similar time segments). This allows users to discover non-periodic glucose patterns without the need to manually browse and compare multiple individual daily views. The decision to include the two situation categories was made after consulting the diabetologists about which aspects of the data they found most relevant to observe. The identification of the time segments in the data that match these categories, and the subsequent clustering of these time segments are explained in more detail in Sec. 5.2.

We fixed the time frame of *meal* situations to span from one hour before to three hours after the meal. For *hypoglycemia* situations, we fixed a time frame from two hours before to two hours after the event. We determined these values in collaboration with the diabetologists. For meals, the postprandial glucose response is relevant to determine if the insulin dose was correct, while for hypoglycemia it is both important to understand what caused it and how the patient treated it.

To select a desired situation category, the user can click on the respective tab at the top of the screen. Each tab has a similar structure, featuring a main view that displays a list of the found glucose clusters (“patterns”), where each cluster is represented by a graph along with characteristic statistics. In the sidebar on the right, the user can view the glucose curves of all situations that match the selected category and filter them by different criteria, such as time of day. Figure 4 shows example patterns from the *Meals* and *Hypoglycemia* tabs.

The graph of each cluster consists of a superimposition of all glucose line graphs within it. The area between the lowest and highest line is filled in. The filling hue corresponds to the glucose ranges, which facilitates the identification of low or high glucose patterns. The start of the meal or hypoglycemia is highlighted with a vertical dashed line across all graphs. Below the graph, a heatmap aggregates the hourly glucose values by mean, unless hypoglycemia occurs. For hypoglycemia, we emphasize low glucose values with a higher weight on the mean. This feature was requested by D2 to get a quick overview of the pattern: “I want to see how the blood sugar was before the meal, during the meal and one, two, three hours after the meal. If the sugar goes up after the meal but returns to the initial value three hours later, the insulin dose was correct. Then I would advise the patient to alter the timing of the insulin dose.” The time and amount of carbohydrates and insulin are encoded below by position and opacity of circles.

To the right of the graph, a selection of metrics further characterizes the pattern. These metrics indicate, among other things, at which time of the day the pattern typically occurs, the average glucose value in the beginning and end of the time window, and the average amounts of carbohydrates and insulin consumed. The metrics are specific to the pattern’s category. The *Meals* category includes the number of minutes between the insulin injection and the meal (*prebolus*) and the average *carbohydrate/insulin ratio* factor. The *Hypoglycemia* category distinguishes between the mean amounts of carbohydrates or bolus insulin consumed before and after the hypoglycemia event, to gain insights into its cause and treatment.

Previous Versions. Detecting and presenting non-periodic patterns in diabetes data is a novel and challenging task that required multiple iterations to arrive at the final version. We adopted a top-down approach in which users first select a category and then observe the glucose patterns in that category. Before settling on this approach, we explored various iterations using a bottom-up approach. We tested solutions which allowed for an objective search for patterns without being constrained

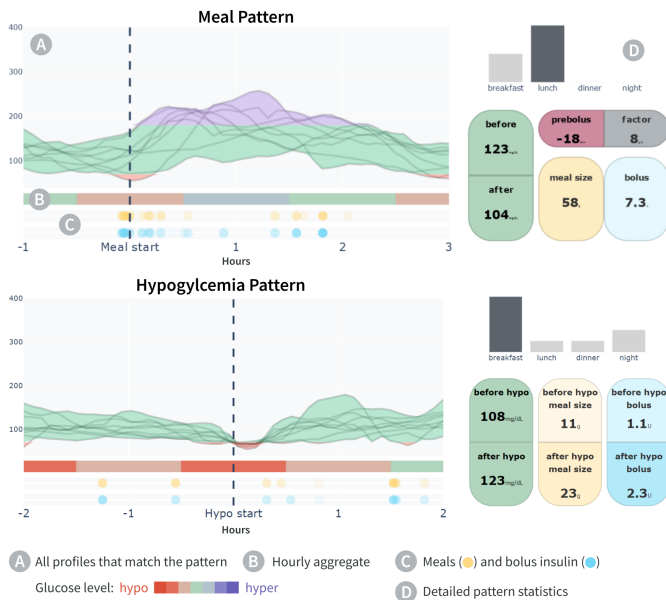


Fig. 4: Top: High glucose variability after lunch. Bottom: Pattern of high glucose before and after the hypoglycemia.

by pre-determined categories. However, the resulting patterns could be more challenging to interpret and may not be clinically relevant.

In an early bottom-up iteration, we segmented the entire glucose dataset into equal-length windows and applied PCA to project each window onto two dimensions. This resulted in a scatter plot where the axes corresponded to the levels and slopes of the glucose data in the windows. We paired this scatterplot with another one that presented the dataset using the horizontal axis to encode time and the vertical axis to encode glucose values. Linking between these plots allowed the user to select regions in the PCA plot and see the corresponding time intervals highlighted in the glucose profile beneath (see Appendix). Our idea was to integrate the PCA plot into a filter view, allowing users to search for specific glucose profiles and view the query results next to the filter (see sketch in Appendix). However, when we shared our sketches with D1, they were confused by the lack of an explicit time axis in the PCA chart. Additionally, they felt that this design would be too complicated to explain to patients during the short meetings.

In a different bottom-up iteration, we applied PCA to cluster time series windows with similar glucose profiles, which were then presented in aggregated area charts using the same percentiles as in the AGP. The average sums of carbohydrates and insulin were also displayed with bars (see Appendix). However, the diabetologists stated that context data was missing. D1: “Many different factors can lead to a similar glucose curve. For example, you can have a rising glucose because you just ate a meal, or because you are in a stressful situation. What exactly are these patterns telling me?”. This feedback led us to develop the more detailed insights representation now used in Marjorie.

5.2 Algorithm Design

To generate the visualizations in the Marjorie frontend, various data processing steps are executed in the backend. The glucose data is aggregated by minute for the AGP, the bolus insulin event data is transformed into a continuous insulin activity series, and the carbohydrate and bolus insulin data is clustered to avoid overplotting. Additionally, for the Insights view, periods of meals and hypoglycemia are detected, which are later clustered into patterns. In the following, we provide details on the clustering of event entries and on pattern detection.

5.2.1 Clustering Event Entries

The overview and daily view summarize event data by grouping data elements that are located near each other. This is applied to both carbohydrate and bolus insulin data. To achieve the multiple levels of

summarization necessary for semantic zooming, we use agglomerative hierarchical clustering. We apply the clustering to each day separately, using a one-dimensional dataset that includes event timestamps converted into numerical decimal values representing the time of the day. An example day of bolus insulin data is used to illustrate the data preparation in a figure in the Appendix.

We used the *average linkage* criterion for clustering, as it has been shown to perform better on datasets with varying cluster sizes and non-periodic data [14]. This is supported by our own observations, as the cluster sizes of the datasets we analyzed were uneven. High blood glucose levels often resulted in frequent, small bolus insulin corrections in short intervals, while normal glucose levels were typically managed with a single dose for meals. The mealtime clustering results from our sample dataset are shown in the Appendix. The clusters are related hierarchically, as indicated by the dendrogram shown in the Appendix. By following a path down the dendrogram, clusters are recursively split into smaller units, until the granularity level matches that of a data sample. The resulting clusters of our example are displayed in the Appendix. Each of the child branches below a determined distance threshold represents an individual cluster. We set this distance threshold, which corresponds to the maximum allowed inter-cluster distance, to a default value of 1.5 hours (i.e., the minimum time between two meals) and decrease it proportionally with increasing zoom level.

5.2.2 Pattern Detection

The first step in generating patterns for the insight view is to identify time intervals that correspond to each situation category. The glucose data of these intervals is then extracted for subsequent clustering.

For finding periods of hypoglycemia, we scan the glucose data for at least three consecutive readings below the threshold of 70 mg/dL. To account for noise in CGM sensors, we allow a tolerance of up to two values that may lie above the threshold. We then define the timestamp of the first reading below the threshold as the start of the hypoglycemic event. From there, we define a hypoglycemic interval that spans from two hours before the start of the event to two hours after.

To detect the meal intervals in the dataset, we use the hierarchical carbohydrate clustering described in Sec. 5.2.1 and aggregate entries in close temporal proximity. Aggregation is necessary for the detection of meal situations since users may occasionally split up a carbohydrate recording into multiple smaller ones, especially when the total carbohydrate intake is high. Entries with a total carbohydrate intake of less than 40 g are filtered, as diabetologists classify them as snacks rather than complete meals. Meal start times are determined from the first carbohydrate entry within a cluster. The meal interval is set to range from one hour before the start of the meal until three hours after.

We then cluster the glucose intervals using *dynamic time warping* (DTW) as the similarity measure. We chose DTW based on its ability to group time series by their shapes, even if they are not synchronized in time. Additionally, DTW enables the clustering of time series with unequal lengths, which is necessary in the case of CGM data gaps. We use hierarchical, agglomerative clustering with average linkage. Clusters are extracted using a distance threshold of 200. We opted for this method because we did not want to pre-specify the number of resulting patterns, which may vary significantly between patients.

6 EVALUATION

We evaluated Marjorie in a case study with a diabetes patient and real data. We also conducted feedback sessions with the diabetologists and a second patient.

6.1 Case Study

To assess Marjorie’s performance on real-world patient data, we enlisted a patient to gain a deeper understanding of how people use our tool. The patient evaluated the latest prototype in two one-hour sessions during which we collected field logs. The user analyzed their own dataset which included glucose, carbohydrates, and bolus and basal insulin data. We summarize the case study in the supplementary video.

At the beginning of the first session, we briefly described Marjorie’s interface and functionalities. We then asked the user to explore their




Fig. 5: Daily view of November 26. The basal insulin delivery had been interrupted between 1 am and 8 am, leading to hyperglycemia.

data in a think-aloud style. The user first opened the *Summary* view, which summarized their data over the previous two weeks, and observed the **Statistics** panel (see Fig. 1, right). The user noticed that their “*time in range is more or less okay, but could be much better.*” They found “*a lot of hypers, even really high ones, but at least no hypos.*” The user then observed that their eA1c value was slightly high and worse than usual, correlating with their time in target range. Still in the *Summary* view, they realized that their glycemic variability was elevated: “*35 % means that my sugar is relatively unstable.*”

They proceeded with analyzing the **AGP** visualization (Fig. 3): “*I see that my problem seems to be mainly during the nighttime. Even the median line is above range there.*” They also included the bolus activity curve in their analysis: “*It’s interesting that there’s barely any insulin activity during the night despite the high sugar. Maybe my [carbohydrate/insulin] factor is wrong at dinner time.*” They continued with analyzing the glucose graph section-by-section. They observed that the glucose decreases and becomes more stable during the morning hours, and related this finding to the bolus insulin corrections in the early mornings that they spotted in the barcode graph. From 11 am on, the glucose remained rather high until the evening: “*The sugar is constantly high and oscillating.*”

The user now clicked on the “Explore day in detail” button (“*Let’s see what this gives me*”) and opened the **Overview** (Fig. 1): “*Wow this is great to analyze my hypers!*” They found that nights exhibited longer stable hyperglycemic periods while hyperglycemia during the day was shorter but more frequent. To analyze this in detail, they compared glucose, carbohydrate, and bolus data using both the compact horizon graphs and the detailed line graphs. They made the following observations: (i) On some evenings (November 28 and 30), they consumed complete meals despite having elevated glucose levels, leading to periods of hyperglycemia during the following nights: “*It was a bad idea to eat here. Even though I injected insulin, the sugar just continued to rise during sleep.*” (ii) They also noticed consecutive instances of hyperglycemia in the afternoons (e.g., on November 27) due to consuming food immediately when glucose levels began to decline: “*It’s like a ping pong effect. I’m in a hyper and make a high insulin correction because I want to get my sugar in range as soon as possible. Then I’m scared to land in a hypo and eat again, but somehow I overeat and land in a hyper again.*” (iii) During the night of November 26, there was a long and intense episode of hyperglycemia with no apparent cause.

To understand this glucose outlier, the patient opened the **Daily View** and navigated to November 26 (Fig. 5) using the date picker. There, they discovered that there was an interruption of basal insulin delivery in the basal area graph: “*Right, my pump tube was recently clogged during the night and I woke up with very high sugar. I immediately changed the tube and gave an insulin correction.*”

Next, the patient clicked on the **Insights** tab and inspected the meal patterns. They found two of the patterns particularly interesting: (i) Hyperglycemia after dinner: “*My sugar is quite okay here before the meal, but three hours later it is always high. And I perform a lot of insulin corrections to get it back in range again.*” (ii) High glucose variability after lunch (Fig. 4): “*Here my sugar rises high directly after eating, but gets back eventually to a value even lower than before the meal.*” The user then noticed that the average prebolus time of this pattern was negative: “*So I only bolused 18 minutes after eating. That’s why the sugar rises so fast. And then I panicked and corrected with a lot of insulin and arrived low in the end.*”

After analyzing the meal patterns, the user clicked on the **Hypoglycemia** tab to observe the hypoglycemia patterns. Since their data contained relatively few instances of low glucose levels, only two patterns were displayed: glucose levels mostly within target range before and after the hypoglycemia, and high glucose before and after the hypoglycemia (Fig. 4). The user quickly understood that “*these are really two different situations. In the first I handle the hypo well, in the second I apparently overeat and land in a hyper.*”

After the user had completed their exploration, we asked them for feedback about their experience with Marjorie.

6.2 Diabetologist & Patient Feedback

We conducted final feedback sessions with the diabetologists (D1, D2) to evaluate the design of Marjorie. We gathered feedback by sharing our screen and going through Marjorie’s views, loaded with case study data. In addition to the case study patient (P1), we also asked another type 1 diabetes patient (P2) for feedback.

All persons found the inclusion of carbohydrate and bolus insulin data in the AGP highly beneficial. D1 noted that the weekday filter is important to distinguish weekends from workdays. The two patients (P1, P2) particularly liked the barcode plots and the bolus insulin activity curve, which provided a clear understanding of how insulin impacts glucose levels.

Likewise, horizon graphs in the overview were well received. D1 and P2 appreciated that problematic values can be seen at a glance. D2 noted that the horizon graphs lead to a much more structured overview compared to other tools. While D1 was initially confused by values of hypo- or hyperglycemia both going upwards in the horizon graph, they quickly understood the encoding due to the applied color scheme. D1 also mentioned that they would like to use the horizon graphs and statistics for explanations in meetings with patients.

The diabetologists (D1, D2) liked the large display of statistics and their placement next to the visualizations. D1 said they use the statistics in all meetings with patients, so they need to be sufficiently large and highlight the most important values. P2 also mentioned the high relevance of statistics to get an overview of the data. D2 commented that other applications display the statistics separately, so they always have to scroll back and forth between the AGP and statistics.

In the daily view, the diabetologists and P2 particularly welcomed the temporal aggregation of carbohydrate and insulin data, as their effect also aggregates. D1 believed that this can be used to train patients but criticized that the bolus insulin data is displayed too far from the glucose data—they would have preferred a superimposition of the data.

The *Insights* view got the most feedback, and we received several suggestions for additional situations and data types that could be considered. D1 liked that they could use the patterns to review recurring scenarios with the patients, adjust therapy parameters, and discover systematic problems. D1, D2, and P2 found that the patterns were well differentiated and meaningful, with D1 and P2 stating that they would typically search and identify similar patterns manually in their analysis. However, P1 felt slightly overwhelmed by the lengthy list of extracted patterns. D1 and P2 also compared the pattern detection to other tools they had used. None of the tools allowed clustering or grouping of meals but instead displayed all data on top of each other, leading to overplotting. D2 suggested *nights* as another pattern category. The participants additionally stated that information on sports activity (D1, D2, P2), other macronutrients (D2), or glycemic index of carbs (P2)

would be helpful in their analysis. However, all were aware of the hurdles in recording these types of data.

Overall, the feedback of all participants was generally positive. They found it easy to understand the components of Marjorie and appreciated the visual and algorithmic innovations introduced in comparison to the state-of-the-art. Based on this feedback, we conclude that Marjorie is suitable for use by both diabetologists and patients, making it applicable and valuable in patient–physician meetings.

7 DISCUSSION

In this section we discuss Marjorie’s strengths and weaknesses and reflect on general insights gained from the design process.

7.1 Tasks & Requirements

Based on our domain analysis in Sec. 4.1, we identified five design requirements for Marjorie. In the following, we describe the extent to which Marjorie fulfills these requirements, qualitatively comparing Marjorie to other state-of-the-art methods. We also discuss limitations and resultant directions for future work.

R1: Make periods of low and high glucose easy to identify. State-of-the-art visualization diabetes tools address this requirement by daily glucose profile views [50]. To provide an overview of multiple days, they plot glucose values on a small vertical scale and distinguish extreme values by color based on target range, which can make it difficult to determine the extent of hypo- or hyperglycemia. In contrast, we propose a daily glucose profile view that uses horizon graphs (see Fig. 1.2). Inspired by Qualizon Graphs [13], our approach visualizes only the curves outside the target range, while displaying the section within the range as a green block. This design effectively highlights periods of hypo- and hyperglycemia and retains information on their extent, while still ensuring a compact layout.

R2: Support finding recurring daily and weekly patterns. As discussed in Sec. 2, the AGP is an effective tool for identifying daily glucose patterns, and has been routinely used in related work. We extend the AGP by incorporating carbohydrate and insulin data. This allows analysts to directly associate the glucose patterns with possible causes, such as behavioral patterns that are linked to carbohydrate and insulin intake. This seemingly simple extension of the AGP was very well received by diabetologists and patients alike.

R3: Visualize all variables in one place Our approach fulfills this requirement through all its components by visualizing glucose together with carbohydrate and bolus insulin data (e.g., alongside the AGP). Especially Marjorie’s overview and daily view visualizations enable an effective comparison of attribute quantities through a combination of hierarchical clustering and semantic zooming. In contrast, other state-of-the-art tools that incorporate variables apart from glucose either do not visualize the quantity of all variables or suffer from overplotting.

Diabetologists and patients both mentioned that it would be helpful to incorporate even more data into the analysis. For instance, sleep and activity information can be extracted from fitness trackers and included as additional situation category tabs in the *Insights* view. The visualizations currently used for the event-data attributes could be readily adopted for activity- or sleep-related data. This enhancement would enable a more comprehensive analysis of potential causes of blood glucose patterns. Sleep stages in particular were a critical aspect of the patient’s investigation in the case study. However, the limited availability of this data for most patients is a challenge.

R4: Bundle and compare similar data Other state-of-the-art tools attempt to address this requirement by (i) algorithmically identifying micro-level hyper- or hypoglycemia events within specific time frames and displaying them in a dedicated view [10], or (ii) enabling users to select a specific event, such as lunch, with which their glucose data is aligned vertically [49]. We experimented with various methods to fulfill this requirement before settling on a hybrid approach. We enable the user to define an alignment event category, such as *Meals* or *Hypoglycemia*. Then, we cluster time intervals that match the selection, based on glucose similarity. This provides an overview of

typical glucose profiles, allowing users to establish connections with treatment decisions and derive insights about meal intake and insulin administration habits. It should be possible to extend this method of pattern identification to additional data attributes mentioned above.

One aspect that requires increased attention in future work is that the lengthy list of detected patterns can overwhelm users. A more easily manageable experience for the user could be achieved by making use of the hierarchical order of the clusters. Top-level clusters could be presented to the user first, with the option to drill down into more detailed clusters on demand. Another potential direction for future work is to allow users to also compare *different* patterns with each other. For example, users could compare situations of hypo- or hyperglycemia to cases with similar carbohydrate intake but without hypo- or hyperglycemia. This could help them identify better strategies.

R5: Display a detailed view of specific days on request Similar to existing tools, we address this requirement with a dedicated *Diary* view, which provides separate daily detail views that users can access with a date picker. Marjorie’s views differ from other approaches in the summarization of event data and the use of semantic zooming.

7.2 Lessons Learned

In this section we reflect on lessons learned from the design process of Marjorie. We expect some of these insights to be transferable to other problem domains with similar constraints.

In one of our earlier designs for the *Insights* view we used a projection-based visualization of time series snippets to allow users to find patterns (see *Previous Versions* in Sec. 5.1.5). Thanks to the iterative design process we soon realized that users were confused by a chart without an explicit time axis. By trying to make use of projections, which we had applied successfully in previous work [11, 21], we could have run into what Sedlmair et al. call an “alluring pitfall for researchers accustomed to technique-driven work” [39]. Early feedback helped us avoid moving further in the direction of “Application Bingo” [30].

In the end, we implemented the *Insights* view by clustering event entries in the background and visualizing them with more established charts and statistics. The use of automated data processing in the background may raise concerns about trust (for a recent discussion see, e.g., the work by van den Elzen et al. [12]). However, we realized that users had to validate any patterns visually regardless of whether they were found in an exploratory or automatic way. Diabetologists appreciated the efficiency that was gained by automating the step of finding potentially interesting patterns and did not express much concerns about distrusting the clustering algorithm.

The appreciation of the increased efficiency relates to the tight time constraints of the patient–physician meetings. Once we had realized that the limited time for the joint analysis and discussion of the results in these meetings was a major bottleneck, we tried to strip away unnecessary details. While this generally led to a cleaner and more effective design, we sometimes overshot by adhering too closely to the idea of “details on demand”. Feedback from the diabetologists made it clear that certain details do need to be shown right from the start.

8 CONCLUSION

This paper described a design study about visually analyzing type 1 diabetes data. The resulting application, Marjorie, enables physicians and patients to search for seasonal and non-periodic patterns, using a combination of visual and algorithmic approaches. Through carefully designed visualizations and interactions, Marjorie addresses important analysis tasks distilled from interviews with two diabetologists. Our solution offers a unique and concise representation of glucose and treatment data, which enables a comprehensive overview and allows carbohydrate and bolus insulin data to be quickly associated with glucose levels. In addition, Marjorie incorporates the algorithmic detection of characteristic glucose profiles that occur during specific pre-selected situations, which simplifies the identification of non-seasonal patterns. In a case study, Marjorie successfully revealed glucose- and therapy-related insights for real-world patient data. Based on feedback by diabetologists and patients, we conclude that Marjorie effectively supports the exploration of patterns in diabetes data.

SUPPLEMENTAL MATERIALS

All supplemental materials are available on OSF at osf.io/34t8c. In particular, they include (1) a full version of this paper with all appendices, (2) a video to demonstrate the workflow and interactions in Marjorie, and (3) the source code of our prototype together with the dataset used in the case study and all figures of our presented approach.

ACKNOWLEDGMENTS

This work was supported by the Austrian Science Fund (FWF DFH 23-N), and the Austrian Research Promotion Agency (FFG 881844). Pro²Future is funded within the Austrian COMET Program under the auspices of the Austrian Federal Ministry for Climate Action, Environment, Energy, Mobility, Innovation and Technology, the Austrian Federal Ministry for Digital and Economic Affairs, and of the States of Upper Austria and Styria. COMET is managed by the Austrian Research Promotion Agency FFG.

REFERENCES

- [1] Abbott Diabetes Care Inc. LibreView Portal. <https://www.freestyleprovider.abbott/us-en/libreview.html>, 2023. Accessed: 2023-08-07. 2, 3
- [2] W. Aigner, S. Miksch, H. Schumann, and C. Tominski. *Visualization of Time-Oriented Data*. Springer, 2011. doi: 10.1007/978-0-85729-079-3_2, 4
- [3] B. Bach, C. Shi, N. Heulot, T. Madhyastha, T. Grabowski, and P. Dragicevic. Time Curves: Folding Time to Visualize Patterns of Temporal Evolution in Data. *IEEE Transactions on Visualization and Computer Graphics*, 22(1):559–568, 2016. doi: 10.1109/TVCG.2015.2467851_2
- [4] R. Bade, S. Schlechtweg, and S. Miksch. Connecting Time-Oriented Data and Information to a Coherent Interactive Visualization. In *Proc. CHI*, pp. 105–112. ACM, Vienna, 2004. doi: 10.1145/985692.985706_2
- [5] S. D. Bartolomeo, Y. Zhang, F. Sheng, and C. Dunne. Sequence Braiding: Visual Overviews of Temporal Event Sequences and Attributes. *IEEE Transactions on Visualization and Computer Graphics*, 27(2):1353–1363, 2021. doi: 10.1109/TVCG.2020.3030442_3
- [6] N. A. Bazaev and K. V. Pozhar. Classification of Factors Affecting Blood Glucose Concentration Dynamics. *Biomedical Engineering*, 47(2):100, 2013. doi: 10.1007/s10527-013-9344-7_2
- [7] C. H. Best and D. A. Scott. The Preparation of Insulin. *Journal of Biological Chemistry*, 57(3):709–723, 1923. doi: 10.1016/S0021-9258(18)85482-5_1
- [8] S. Carpendale. Evaluating Information Visualizations. In J. T. Stasko, J.-D. Fekete, C. North, and C. North, eds., *Information Visualization: Human-Centered Issues and Perspectives*. Lecture Notes in Computer Science, pp. 19–45. Springer, 2008. doi: 10.1007/978-3-540-70956-5_2_3
- [9] T. Danne, R. Nimri, T. Battelino, R. M. Bergenstal, K. L. Close, J. H. DeVries, S. Garg, L. Heinemann, I. Hirsch, S. A. Amiel, R. Beck, E. Bosi, B. Buckingham, C. Cobelli, E. Dassau, F. J. Doyle, III, S. Heller, R. Hovorka, W. Jia, T. Jones, O. Kordonouri, B. Kovatchev, A. Kowalski, L. Lafel, D. Maahs, H. R. Murphy, K. Nørgaard, C. G. Parkin, E. Renard, B. Saboo, M. Scharf, W. V. Tamborlane, S. A. Weinzimer, and M. Phillip. International Consensus on Use of Continuous Glucose Monitoring. *Diabetes Care*, 40(12):1631–1640, 2017. doi: 10.2337/dc17-1600_1, 2, 5
- [10] Dexcom Inc. Dexcom Clarity Reports. <https://provider.dexcom.com/education-research/cgm-education-use/product-information/dexcom-clarity-reports-overview>, 2023. Accessed: 2023-08-07. 2, 3, 9
- [11] K. Eckelt, A. Hinterreiter, P. Adelberger, C. Walchshofer, V. Dhanoa, C. Humer, M. Heckmann, C. Steinparz, and M. Streit. Visual Exploration of Relationships and Structure in Low-Dimensional Embeddings. *IEEE Transactions on Visualization and Computer Graphics*, 29(7):3312–3326, 2023. doi: 10.1109/TVCG.2022.3156760_2, 9
- [12] S. v. d. Elzen, G. Andrienko, N. Andrienko, B. D. Fisher, R. M. Martins, J. Peltonen, A. C. Telea, and M. Verleysen. The Flow of Trust: A Visualization Framework to Externalize, Explore, and Explain Trust in ML Applications. *IEEE Computer Graphics and Applications*, 43(2):78–88, Mar. 2023. doi: 10.1109/MCG.2023.3237286_9
- [13] P. Federico, S. Hoffmann, A. Rind, W. Aigner, and S. Miksch. Qualizon Graphs: Space-Efficient Time-Series Visualization with Qualitative Abstractions. In *Proc. AVI*, pp. 273–280. ACM, Como, 2014. doi: 10.1145/2598153.2598172_2, 9
- [14] L. Ferreira and D. B. Hitchcock. A Comparison of Hierarchical Methods for Clustering Functional Data. *Communications in Statistics-Simulation and Computation*, 38(9):1925–1949, 2009. doi: 10.1080/03610910903168603_7
- [15] D. Flück. Coblis: Color Blindness Simulator. <https://www.color-blindness.com/coblis-color-blindness-simulator/>, 2021. Accessed: 2023-08-07. 5
- [16] T. Fujiwara, Shilpika, N. Sakamoto, J. Nonaka, K. Yamamoto, and K.-L. Ma. A Visual Analytics Framework for Reviewing Multivariate Time-Series Data with Dimensionality Reduction. *IEEE Transactions on Visualization and Computer Graphics*, 27(2):1601–1611, 2021. doi: 10.1109/TVCG.2020.3028889_2
- [17] Glooko Inc. Glooko. <https://www.glooko.com>, 2023. Accessed: 2023-08-07. 2
- [18] A. Green, S. M. Hede, C. C. Patterson, S. H. Wild, G. Imperatore, G. Roglic, and D. Beran. Type 1 Diabetes in 2017: Global Estimates of Incident and Prevalent Cases in Children and Adults. *Diabetologia*, 64(12):2741–2750, 2021. doi: 10.1007/s00125-021-05571-8_1
- [19] Y. Guo, S. Guo, Z. Jin, S. Kaul, D. Gotz, and N. Cao. Survey on Visual Analysis of Event Sequence Data. *IEEE Transactions on Visualization and Computer Graphics*, 28(12):5091–5112, 2022. doi: 10.1109/TVCG.2021.3100413_1, 2
- [20] H. Hall, D. Perelman, A. Breschi, P. Limcaoco, R. Kellogg, T. McLaughlin, and M. Snyder. Glucotypes Reveal New Patterns of Glucose Dysregulation. *PLOS Biology*, 16(7):1–23, 2018. doi: 10.1371/journal.pbio.2005143_3
- [21] A. Hinterreiter, C. Steinparz, M. Schöfl, H. Stitz, and M. Streit. Projection Path Explorer: Exploring Visual Patterns in Projected Decision-Making Paths. *ACM Transactions on Interactive Intelligent Systems*, 11(3–4):Article 22, 2021. doi: 10.1145/3387165_9
- [22] R. I. Holt, J. H. DeVries, A. Hess-Fischl, I. B. Hirsch, M. S. Kirkman, T. Klupa, B. Ludwig, K. Nørgaard, J. Pettus, E. Renard, J. S. Skyler, F. J. Snoek, R. S. Weinstock, and A. L. Peters. The Management of Type 1 Diabetes in Adults. A Consensus Report by the American Diabetes Association (ADA) and the European Association for the Study of Diabetes (EASD). *Diabetes Care*, 44(11):2589–2625, 2021. doi: 10.2337/dci21-0043_1, 2, 3, 4, 5
- [23] K. Holtzblatt and S. Jones. Contextual Inquiry: A Participatory Technique for System Design. In *Participatory Design*. CRC Press, 1993. doi: 10.1201/9780203744338-9_3
- [24] R. J. Hyndman and G. Athanasopoulos. *Forecasting: Principles and Practice*. OTexts, 2018. <https://otexts.com/fpp2/>. 3
- [25] M. L. Johnson, T. W. Martens, A. B. Criego, A. L. Carlson, G. D. Simonson, and R. M. Bergenstal. Utilizing the Ambulatory Glucose Profile to Standardize and Implement Continuous Glucose Monitoring in Clinical Practice. *Diabetes Technology & Therapeutics*, 21(S2):S17–S25, 2019. doi: 10.1089/dia.2019.0034_3, 4
- [26] B. P. Kovatchev. Metrics for Glycaemic Control—from Hba1c to Continuous Glucose Monitoring. *Nature Reviews Endocrinology*, 13(7):425–436, 2017. doi: 10.1038/nrendo.2017.3_4
- [27] B. Lobo, L. Farhy, M. Shafiei, and B. Kovatchev. A Data-Driven Approach to Classifying Daily Continuous Glucose Monitoring (CGM) Time Series. *IEEE Transactions on Biomedical Engineering*, 2021. doi: 10.1109/TBME.2021.3103127_3
- [28] Medtronic Inc. Medtronic IPRO 2 System for CGM. <https://www.medtronic.com/de-de/fachreise/diabetes/produkte/cgm/ipro2.html>, 2018. Accessed: 2023-08-07. 2
- [29] Medtronic Inc. Medtronic CareLink Software. <https://www.medtronicdiabetes.com/products/carelink-personal-diabetes-software>, 2023. Accessed: 2023-08-07. 2, 3
- [30] T. Munzner. Process and Pitfalls in Writing Information Visualization Research Papers. In *Information Visualization: Human-Centered Issues and Perspectives*. Lecture Notes in Computer Science, pp. 134–153. Springer, 2008. doi: 10.1007/978-3-540-70956-5_6_9
- [31] T. Munzner. *Visualization Analysis and Design*. CRC Press, Taylor & Francis Group, 2014. doi: 10.1201/b17511_5
- [32] W. Müller and H. Schumann. Visualization Methods for Time-Dependent Data - An Overview. In *Proc. WSC*, vol. 1, pp. 737–745. IEEE, New Orleans, 2003. doi: 10.1109/WSC.2003.1261490_2
- [33] D. M. Nathan. Long-Term Complications of Diabetes Mellitus. *New England Journal of Medicine*, 328(23):1676–1685, 1993. doi: 10.1056/

NEJM199306103282306 2

- [34] Plotly Technologies Inc. Plotly Dash User Guide & Documentation. <https://dash.plotly.com/>, 2023. Accessed: 2023-08-07. 5
- [35] H. Reijner. The Development of the Horizon Graph. In *Proc. IEEE VIS Workshop: "From Theory to Practice: Design, Vision and Visualization"*. IEEE, Columbus, 2008. <https://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.363.5396>. 2
- [36] A. Rind. Interactive Information Visualization to Explore and Query Electronic Health Records. *Foundations and Trends® in Human-Computer Interaction*, 5(3):207–298, 2013. doi: 10.1561/1100000039 2
- [37] A. Rind, W. Aigner, S. Miksch, S. Wiltner, M. Pohl, T. Turic, and F. Drexler. Visual Exploration of Time-Oriented Patient Data for Chronic Diseases: Design Study and Evaluation. In *Proc. USAB*, pp. 301–320. Springer, Graz, 2011. doi: 10.1007/978-3-642-25364-5_22 2
- [38] W. Rostène and P. De Meyts. Insulin: A 100-Year-Old Discovery With a Fascinating History. *Endocrine Reviews*, 42(5):503–527, 2021. doi: 10.1210/endrev/bnab020 1
- [39] M. Sedlmair, M. Meyer, and T. Munzner. Design Study Methodology: Reflections from the Trenches and the Stacks. *IEEE Transactions on Visualization and Computer Graphics*, 18(12):2431–2440, 2012. doi: 10.1109/TVCG.2012.213 9
- [40] H. Stitz, S. Gratzl, W. Aigner, and M. Streit. ThermalPlot: Visualizing Multi-Attribute Time-Series Data Using a Thermal Metaphor. *IEEE Transactions on Visualization and Computer Graphics*, 22(12):2594–2607, 2016. doi: 10.1109/TVCG.2015.2513389 2
- [41] The Nightscout Foundation. The Nightscout Project. <http://www.nightscout.info/>, 2023. Accessed: 2023-08-07. 2, 3
- [42] The OpenAPS Community. Understanding Insulin on Board (IOB) Calculations. <https://openaps.readthedocs.io/en/latest/docs/WhileYouWaitForGear/understanding-insulin-on-board-calculations.html>, 2017. Accessed: 2023-08-07. 4
- [43] Tidepool Project. Tidepool. <https://tidepool.org/>. Accessed: 2023-08-07. 2, 5
- [44] C. Tominski and W. Aigner. The TimeViz Browser – A Visual Survey of Visualization Techniques for Time-Oriented Data. <https://browser.timeviz.net>, 2023. Accessed: 2023-08-07. 2
- [45] Q. Wang and R. Laramée. EHR STAR: The State-Of-the-Art in Interactive EHR Visualization. *Computer Graphics Forum*, 41(1):69–105, 2022. doi: 10.1111/cgf.14424 2
- [46] A. Z. Woldaregay, E. Årsand, T. Botsis, D. Albers, L. Mamykina, and G. Hartvigsen. Data-Driven Blood Glucose Pattern Classification and Anomalies Detection: Machine-Learning Applications in Type 1 Diabetes. *Journal of Medical Internet Research*, 21(5):e11030, 2019. doi: 10.2196/11030 3
- [47] J. C. Wong, A. B. Neinstein, H. Look, B. Arbiter, N. Chokr, C. Ross, and S. Adi. Pilot Study of a Novel Application for Data Visualization in Type 1 Diabetes. *Journal of Diabetes Science and Technology*, 11(4):800–807, 2017. doi: 10.1177/1932296817691305 2
- [48] Y. Yu, D. Kruffy, J. Jiao, T. Becker, and M. Behrisch. PSEUDO: Interactive Pattern Search in Multivariate Time Series with Locality-Sensitive Hashing and Relevance Feedback. *IEEE Transactions on Visualization and Computer Graphics*, 29(1):33–42, 2023. doi: 10.1109/TVCG.2022.3209431 3
- [49] Y. Zhang, K. Chanana, and C. Dunne. IDMVis: Temporal Event Sequence Visualization for Type 1 Diabetes Treatment Decision Support. *IEEE Transactions on Visualization and Computer Graphics*, 25(1):512–522, 2018. doi: 10.1109/TVCG.2018.2865076 2, 3, 9
- [50] Y. Zhang, A. G. Parker, and C. Dunne. Information Visualization for Diabetes Management: A Literature Review. In *Proc. PervasiveHealth*, pp. 427–430. ACM, Atlanta, 2021. doi: 10.1145/3421937.3421957 2, 3, 5, 9