MineTime Insight: Visualizing Meeting Habits to Promote Informed Scheduling Decisions

Marco Ancona, Marilou Beyeler, Markus Gross, Tobias Günther

Abstract—Corporate meetings are a crucial part of business activities. While numerous academic papers investigated how to make the scheduling process of meetings faster or even automatic, little work has been done yet to facilitate the retrospective reasoning about how time is spent on meetings. Traditional calendar applications do not allow users to extract actionable statistics although it has been shown that reflection-oriented design can increase the users' understanding of their habits and can thereby encourage a shift towards better practices. In this paper, we present *MineTime Insight*, a tool made of multiple coordinated views for the exploration of personal calendar data, with the overarching goal of improving short and long-term scheduling decisions. Despite being focused on the working environment, our work builds upon recent results in the field of Personal Visual Analytics, as it targets users not necessarily expert in visualization and data analysis. We demonstrate the potential of MineTime Insight, when applied to the agenda of an executive manager. Finally, we discuss the results of an informal user study and a field study. Our results suggest that our visual representations are perceived as easy to understand and helpful towards a change in the scheduling habits.

Index Terms—Scheduling, calendar, personal visual analytics, casual information visualization, virtual assistant.

1 Introduction

In the corporate environment, time is a precious resource. As such, it is desirable to find good time management practices to distribute the workload among several activities, while maximizing productivity and ensuring a good work-life balance. Depending on the role within the company, a significant portion of the employee time can be invested in meetings, which is especially true for team leaders, managers and executives [1], [2]. For this reason, analytics tools promoting and facilitating self-monitoring of scheduling habits can be a valuable resource in the work environment.

In recent years, the field of Personal Informatics (PI) [3], has consolidated best practices to develop tools that aim at helping people to collect personally relevant information for the purpose of self-reflection. This is motivated by the observation that self-monitoring and self-reflection often affect behavior, and this change typically goes in the desired direction of improvement [4], [5].

The visualization community has also increased its attention towards a broader audience of users that are not visualization or data analytics experts [6]. Personal Visual Analytics (PVA) [6] emphasizes the need for facilitating user reasoning by visual representations within a personal context. While we are only interested in scheduling habits within the working environment, which is outside the realm of *personal context* as defined by Huang et al. [4], we still argue that Personal Visual Analytics is the appropriate way to address the problem, given the broader audience and the different motivations than in traditional information visualization [6].

Nowadays, calendar tools like Google Calendar or Microsoft Outlook are ubiquitous but still offer users little to no support for self-reflection about their own scheduling habits, i.e., they are arguably not suitable for analyzing aggregated scheduling trends over months, for individuals or teams.

• *M. Ancona, M. Beyeler, M. Gross and T. Günther are with ETH Zurich. E-mail: marco.ancona, mabeyele, grossm, tobias.guenther@ethz.ch*

We address the following question: *How can calendar views used in the corporate environment be made appropriate by the use of visualizations to increase the users' self-awareness about the management of their meeting time - including users with little experience in data visualization and analysis?* To tackle the problem, we formulate three Design Goals suggested by related literature and refined through preliminary interviews.

We present MineTime Insight, a visual analytics tool consisting of four coordinated views to reveal hidden patterns and meeting distributions in multivariate calendar data. MineTime Insight can be used to answer questions about how much time, in absolute and relative terms, the user spends with single persons or groups. It also allows users to evaluate the number of hours spent weekly on meetings over the past months and drill down to see statistics about individual people or groups. Finally, it provides a simplified representation of the meeting periodicity to enable people with little experience in statistical reasoning to evaluate not only how often, but also how regularly, people have been met in the past. In order to seamlessly integrate this into the daily user workflow, we deployed and tested this tool as part of *MineTime*, a proprietary calendar application.

We perform a case study, where MineTime Insight is applied to the real agenda of an executive manager, showing the power of our visualization techniques with respect to traditional calendar views. Finally, we present quantitative and qualitative feedback from interviews with eight recruited participants as well as from users "in the wild".

The major contributions of this paper are as follows. First, we provide, to the best of our knowledge, a novel characterization of the problem of analyzing meeting time in the corporate environment, including Design Goals and Visualization Tasks. Secondly, we discuss the visualization design of multiple linked views to support the analysis of personal meeting time, and validate it through an informal user study as well as a field deployment; thirdly, we reflect

©2019 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works. DOI 10.1109/TVCG.2019.2941208

on the scope of Personal Visual Analytics and its role for our design study, showing that these ideas indeed apply well also to the analysis of a working environment.

2 Related Work

Visual Analytics for Non-Expert Users. Personal Informatics (PI) [3] first formalized the idea of helping non-expert users to collect and reflect on personal information for the purpose of self-knowledge and self-reflection. More recently, Huang et al. [6] described the taxonomies of a growing field within the visualization community called Personal Visual Analytics (PVA), dealing with the challenges in designing visualizations for use in non-professional situations (personal context). The target audience of PVA is characterized by i) general lack of expertise in data analysis and statistical reasoning; ii) limited time budget and low priority in performing analytical tasks and iii) different goals, motivations and expectations than those of their professional role. Similarly, Casual Infovis [7] highlights the necessity of a smooth integration of visualization tools into the users' routines as well as the importance of building *pleasing* tools, delightful without regard whatsoever for their utility.

Even though we target a professional context, we strongly benefit from the recent results and considerations of PI, PVA and Casual Infovis. In fact, in dealing with time management, our target users are arguably better described by their personal context than their professional one. First, for most people, managing and distributing time is a side effect of carrying out daily work-related tasks, as opposed to managing time as an explicit practice. Secondly, people usually have lower priority and time budget to analyze their time and this is not part of their work-role expectations. Finally, the vast majority of people trying to improve scheduling decisions are not experts in visualization or data analysis.

PVA has been applied to a variety of data sources including personal finance [8], localization [9], physical activity logs [10], [11], energy consumption [12] and computer-related activities such as email exchanges [13], application usage [14] and web navigation [15]. To the best of our knowledge, no previous work has focused on calendar data and meeting time, through the lens of PVA.

Experiments to display non-calendar related information, such as fitness data or daily activities, on an augmented calendar view showed the potential of visualizing personal information within the daily user workflow [11], [16]. Our work differs in that the calendar is not a visualization tool but rather the source of information to be explored.

Calendars and Time-oriented Data. Early work on calendar data visualization focused on providing rapid access to an individual's daily agenda or facilitating the scheduling of group meetings. Spiral Calendar and Time Lattice [17] used a 3D environment to organize different levels of detail in the user agenda or the schedules of a group of people, respectively. With a similar goal of facilitating groupware scheduling, Augmented Daily Calendar [18] and Availability Bar [19] augmented the calendar view with additional information about other people's availability, scheduling preferences or attendance likelihood. All these works highlighted the importance of providing the user additional information while preserving the calendar usability as a productivity tool. However, contrarily to our work, their focus was primarily set on facilitating short-term decisions while we are more interested in analyzing long-term scheduling habits.

Particularly relevant for our work is Interactive People Cloud (IPC) [20]. This tool has been proposed to analyze whom a person is spending time with, using calendar data. The tool shows a "people-cloud", similar to a word-cloud, where photos of calendar contacts are sized proportionally to the number of meetings with each person. It also provides other statistics in textual form, like the average and total duration of the meetings. Compared to our work, however, IPC only provides basic statistics about meeting time, being part of a larger system to analyze how time is spent on a broader range of activities. For example, one of the limitations of IPC, as emerges from their user study, is the lack of a temporal dimension for most of the metrics displayed to the users. As scheduling habits change over time, we believe the temporal evolution is crucial to assess the improvement, regression or stagnation towards personal goals, which is at the center of our investigation.

Visualization techniques for generic multivariate time-oriented data has been also extensively investigated in the past [21], [22] and benefit a number of applications. For instance, PostHistory [23] aims at visually uncovering patterns from email activity, such as closer contacts and exchange frequency over time. Similar to our application, PostHistory adopts a user-centric approach, focusing on a single user's interaction with other people. However, their analysis is only based on emails and motivated by the exploration of social circles. By looking at calendar data with a focus on productivity, our work deals with different data, challenges and objectives.

3 Scope and Calendar Data

In this section, we first define the scope of the project, then we describe our data.

Scope. The overarching goal of this work is to develop a visual analytics tool that facilitates the exploration, self-awareness and improvement of short and long-term scheduling habits towards best-practices and/or personal goals. In particular, our analysis focuses on meetings within the corporate environment, with the ultimate goal of incrementing the daily productivity.

Although meetings might also occur in other professional contexts, meetings within the corporate environment have been extensively studied for decades, both in terms of cost and impact on productivity [1], [2], [24], [25], [26] and in terms of opportunities and challenges for the adoption of groupware systems [27], [28], [29], [30], [31], [32].

We leverage this vast literature to characterize the context of our work and therefore exclude from the analysis meetings that only involve professionals outside of the corporate setting (ie. lawyers, dentists, doctors, contractors, independent workers, etc.), as these might have different goals and requirements. We also exclude from our research scope meetings that might occur among private individuals and registered on their private calendars, as meeting analysis has arguably limited applicability or impact in this context.

Calendar data. Since there is no unique naming convention in related literature or industry, we describe our data using definitions taken from the Internet Calendaring and Scheduling Core Object Specification [33] as well as from the Microsoft Exchange documentation [34].

Calendars are collections of *events*, organized chronologically. Automatic processing of events is a complex task because they contain a mixture of structured data (*date*, *time*, *duration*, *timezone*, *creation timestamp*, *organizer* and *attendees* [33]) and unstructured information given in form of free text (*title*, *description* and *location*).

Events can be further partitioned into *appointments* and *meetings* [34]. Appointments are calendar events that users create for themselves and that have no listed attendees other than the organizer. While *appointments* might involve other people (e.g., an appointment with a doctor), this information is not stored in the event metadata. Contrarily, *meetings* are events that involve multiple attendees, explicitly listed in the calendar entry and therefore immediately processable. We take into account both recurring and not recurring meetings.

In this work, we focus on *meetings* because our processing pipeline requires an explicit list of attendees. Although, in practice, some users record their meetings as *appointments* (possibly including the attendees' names only in the title of the event) the extraction of the attendees from a textual string is not straight-forward, sometimes ambiguous and prone to errors and therefore left out of the scope of this work.

Attendees are identified by their email address, which constitutes a globally unique identifier. We organize attendees in *groups*. Groups are bootstrapped automatically based on the *email domain*, which is often tied to the company the attendee belongs to. Generic domains, such as *@gmail.com*, mostly used for private accounts, are automatically grouped into "Others". Users can later create new groups, edit them by adding or removing attendees, or delete them.

4 Design Goals

In the following, we present three Design Goals. When discussing meeting habits, the literature considers three typical corporate roles: Managers/Executives (EX), Personal Assistants (PA) and Staff members (SM) [1], [31], [35]. In this context, "Staff members", sometimes called "individual contributors" [25], [32], generically refers to most employees to whom no one reports. It has been noticed that these professional roles can significantly differ for i) the number of meetings attended weekly [1], [2], ii) the extent of the calendar management delegation, with Managers often supported by a Personal Assistant in the organization of their agenda [25] and iii) the characterization of work-related activities and their impact on other employees [25].

As a productivity tool targeting the corporate environment, we designed MineTime Insight based on the characterization of these roles as our prototypical users. The Design Goals are grounded on the aforementioned works, which investigate both the meeting patterns and the use of calendar software in the corporate environment, as elaborated below. Additionally, we conducted semi-structured interviews with two Senior Managers (EX1-2), two Personal Assistants (PA1-2) and four Staff members (SM1-4). Each interview lasted between 30 and 60 minutes and served us as Task Analysis according to Hackos' principles [36]. We encouraged the participants to walk us through their daily time management duties, which differ significantly depending on their role. We took notes of the observations and analyzed these later for comparison with previous studies.

From the literature and as further refined by the informal user interviews, we derived the following three Design Goals, organized according to the *Supporting Awareness for Action* pattern for Personal Visual Analytics [6]. Applications of this kind aim to "provide in-the-moment or on-going awareness with respect to personal behavior" and are usually characterized by 1) functionalities for *looking up data* with a quick glance, 2) in-the-moment feedback or suggestions to *support immediate action* and 3) *support for continuous awareness* over time. Our Design Goals follow these three principles.

G1 Visually display actionable statistics about the time invested in meetings

Mitntzberg's study suggests managers spend up to 70% of time in meetings (only 10% unscheduled) and on average they attend 8 meetings per day [37]. The economic impact of meetings is enormous and this makes tracking of *how much* time is spent on meetings an important cost analysis for companies [1]. One of the Personal Assistants we interviewed explained how their manager's calendar was manually analyzed at the end of each year to estimate how much time had been dedicated to different activities and to plan corrections for the following year. With this in mind, we identified the first visualization task as follows:

T1.1 Quantify the amount of overall working time dedicated to meetings over a long time range.

One of the managers we interviewed told us that it is sometimes difficult to keep track of all people in a large team. The manager does not have scheduled recurrent 1-on-1 meetings with each of the Staff members but rather relies on them to request a meeting on a reasonably regular basis. On the other hand, a study conducted on 2.5 million manager-led teams in 195 countries [26] found that employees whose managers hold regular meetings with them are almost three times as likely to be engaged as employees whose managers do not hold regular meetings with them. Given the importance of regular meetings, we identified additional three visualization tasks to analyze *with whom*, *how often* and *how regularly* meetings have been scheduled in the past:

T1.2 Quantify the amount of working time dedicated to specific people or groups during previous months.

T1.3 Identify the frequency and regularity of meetings involving specific people and groups.

T1.4 Compare meeting patterns (e.g., identify outliers) among and within groups using the metrics of T1.2 and T1.3.

It has been noticed that the manager's activities are characterized by "brevity, variety and discontinuity" [24]. More than 50% of meetings are scheduled a week in advance or more [27] and, in some cases, secretaries handle all meeting scheduling and answer all meeting invitations [29]. All these factors make it difficult for a person to recall the details of past and upcoming meetings [27]. We aim at providing contextual tools to avoid poor preparation and the lack of a meeting agenda which have been reported among the most common reasons for meeting ineffectiveness [1]:

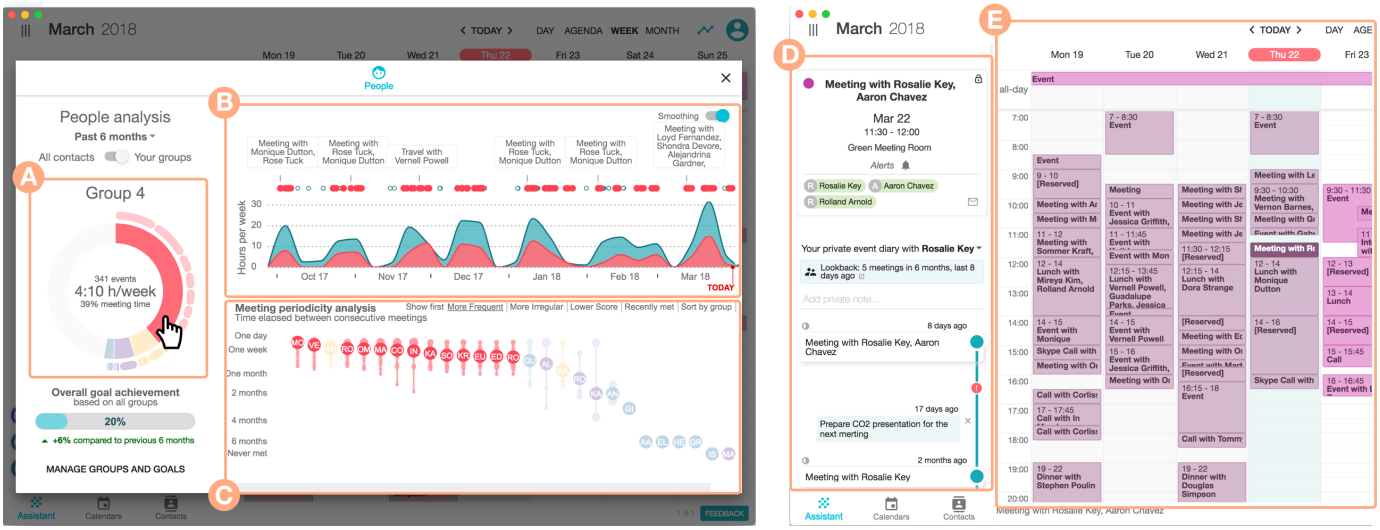

Fig. 1. The Analytics Board (left) provides three views to visualize aggregated scheduling statistics: the Meeting Breakdown (A), the Meeting Pressure (B) and the Periodicity Analysis (C). The Analytics Board is integrated into a calendar application (right). The Contact Diary (D) is available next to the calendar view (E) or from the Analytics Board upon selection of a contact.

T1.5 Get an overview of a specific person's meeting history, i.e., quickly identify when meetings with a person occurred in the recent past, what was discussed and when upcoming meetings will take place.

G2 Provide tools to set and track goals towards improvement of one's meeting schedule Being able to set goals and monitor personal achievements towards these, has been shown to be extremely helpful for motivation, focus, persistence and ultimately self-improvement [38], [39]. Having tools that enable the setting and tracking of goals has been found important by users in previous studies to turn scheduling statistics into actionable data [20]. Notice that personal goals can be either *self-set* or *assigned* [39], for example by a supervisor. During the interviews, one of the managers told us that he/she targets at least one 1-on-1 meeting with each member of the team (Staff) every month (self-set goal). On the other hand, the Personal Assistant explained that it is hardly possible to schedule recurrent events because of the many changes the Manager's schedule undergoes within four weeks. Moreover, it is too time consuming to track the last meeting date for all Staff members and schedule new meetings for them. Instead, Staff members are encouraged to request a meeting time on a regular basis, about once a month (assigned goal). In terms of design choices, this translates into supporting awareness while preserving the schedule flexibility, without increasing the user workload. As for the visualization tasks, we therefore focus on displaying goals and action recommendations:

T2.1 Identify improvement, retrogression or stagnation towards personal meeting goals.

T2.2 Identify what actions should be undertaken in order to improve towards the target goals, including identifying critical scheduling cases that should be addressed.

G3 Make the analysis accessible to a broader audience over time In order to achieve the long-term goal of shifting the scheduling habits towards the user's goals, early-stage interest needs to transition into a solid regular usage. This process, of course, depends on the audience. Our users are not professional meeting analysts but rather people for whom meeting-related data is not the central focus of their work. This means that they may have low incentive to analyze it regularly and a low understanding of the data structures underlying meeting analysis. Moreover, they may have low visualization literacy in general.

Previous works also remarked the need to combine low attentional demand and just-sufficient salience, not to interrupt life routines [6]. In order to achieve this, as well as to reduce technical friction that might further damage the engagement of less motivated users, the visualization tools must seamlessly integrate into the user daily workflow. Concretely, this raises a number of technical challenges, including the integration into existing calendar tools, uncontrolled data sources, missing or noisy user-generated content, privacy concerns, scalability to support calendars containing thousands of events and computational constraints to ensure the system can run in real-time on consumer-level hardware.

Rather than identifying new visualization tasks, G3 affects our design choices, stressing the importance of visualizations that are novel and aesthetically appealing, the avoidance of information clutter and the need for intuitive tools targeting people with limited visualization expertise [40].

5 MineTime Insight

MineTime Insight is an extension of an existing calendar application that provides tools to explore and analyze one's past meeting behavior. The interface is designed to allow users to get a sense of their overall scheduling habits in the past months, while enabling easy access to contact-specific information on demand. The mixture between aggregated and event-specific information allows users to reminisce past events in their context and promote introspective behaviors. In the remainder of this section, we discuss the layout, interaction, navigation and deployment of this tool.

5.1 Layout Design

MineTime Insight has four linked views, three of which dedicated to medium and long-term analysis and one focused on

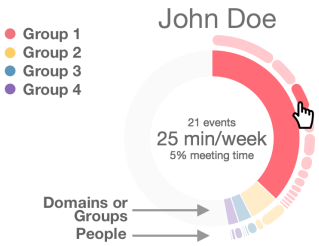

Fig. 2. The Meeting Breakdown shows how time is partitioned among different companies (based on company domain) or custom groups. It also shows to which people within these groups more time is dedicated.

daily productivity. Fig. 1 shows an overview of the system. Within the Analytics Board, the user can monitor aggregated statistics over several months and access contact-specific information on demand. Here are three interconnected views: the Meeting Breakdown (A), the Meeting Pressure (B) and the Periodicity Analysis (C). The Contact Diary (D) provides an overview of the scheduling history for a single contact. In designing the layout, we focused on making it accessible to people with no visualization expertise (G3). This includes reducing the information clutter, for which we use some of the techniques illustrated by Ellis et al. [41]. In the following, we explain the four views in detail.

5.1.1 *Meeting Breakdown*

Displayed on the left sidebar of the Analytics Board (Fig. 1 - A), the Meeting Breakdown provides the main entry point for the analysis and the highest level of aggregation. As shown in Fig. 2 in more detail, it consists of a two-layer Sunburst plot [42], where the inner layer represents contact categories (email domains or customized groups) and the outer layer represents individual contacts within these categories. A unique color, chosen from a predefined palette, is assigned to each category and inherited by its contacts.

Sunburst plots have been shown to be an effective technique for visualizing and comparing hierarchically-structured data [42]. In this case, the total meeting time is partitioned among categories and, in turn, time dedicated to each category is partitioned among its members. A slight variation in style on the outer layer emphasizes the difference between groups and people. The plot is scaled such that the full circle represents the total sum of hours the user spent on meetings in the analyzed time window. Each arc in the Sunburst chart represents the fraction of time dedicated to a category or a specific person in the same period. When a meeting involves multiple attendees, we equally partition the duration of the meeting among them. This choice ensures that the time spent with each person sums up to the total meeting time or, in other words, that the juxtaposition of the plot arches always covers the full circle.

This plot serves both as a legend to map colors and categories and as breakdown analysis of the total meeting time within the considered period (T1.1). It also visually facilitates the comparison between time invested in different companies (domain mode) or teams/roles (group mode), which builds towards one of our desired Visualization Tasks (T1.4). By hovering the mouse on one of the arches, the user highlights a group or a contact, which triggers the display of details: number of events involving the group or person

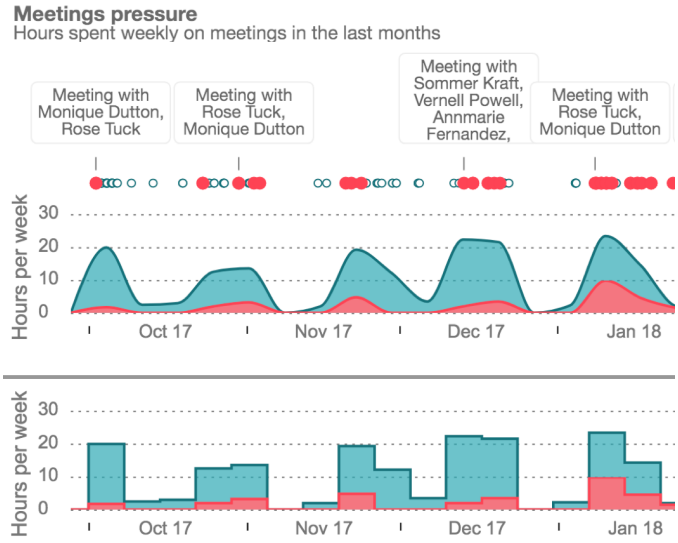

Fig. 3. The Meeting Pressure shows the number of hours invested in meetings over time, in total (blue curve) and for a subset of contacts (red curve). Smoothing can be disabled if necessary (bottom). Meeting titles can be accessed on demand by interacting with the mouse.

in the period, average weekly meeting time and percentage with respect to total meeting time (T1.2).

5.1.2 *Meeting Pressure*

The aggregated metrics of the Meeting Breakdown across several months, although quick to read, might hide non-uniform trends over time. The Meeting Pressure, displayed on the top-left of the Analytics Board and shown in more detail in Fig. 3, aims at providing a time-oriented context for the analysis of how many working hours per week have been dedicated to meetings over the past months (T1.1-1.2).

As meetings have been found to have a positive correlation with daily fatigue and subjective perception of workload [43], we call the aggregated amount of time invested in meetings within a specific time frame, "meeting pressure". For the same reason, we claim this is a useful metric to monitor in order to build self-awareness and aim at improving scheduling habits (G1). In this case, the scale on the vertical axis shows meeting pressure in hours per week. Although the choice of a weekly baseline is arbitrary, working time is conventionally measured week-to-week in many countries [44]. For this reason, by displaying weekly hours, the user can compare against a standardized baseline, e.g., the 40-hours work week. This piece of information cannot be easily read from a traditional calendar view because i) meetings and other types of events are hardly distinguishable unless the user manually assigns different colors to them, ii) the standard grid representation for month-week-day calendars is only appropriate for representing the cyclical nature of time that assumes a fixed granularity [22] and iii) traditional calendar systems usually do not allow to display events from year-long periods of time together.

The meeting pressure is encoded using an Area Chart (blue) that spans horizontally over a time window which is adjustable by the user. Area Charts are usually perceptually preferred over Line Plots for wide aspect ratios [21], as in our case. A secondary pressure curve (red), superimposed on the first one, enables a direct comparison between the total

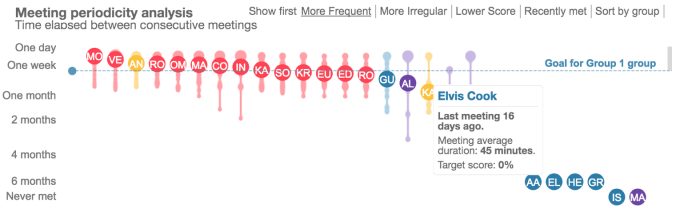

Fig. 4. The Periodicity Analysis enables the comparison among and within groups. The average meeting period is encoded on the y-axis while the meeting regularity is represented by the spread of the period distribution. Several criteria can be chosen to sort contacts horizontally.

meeting time and the time dedicated to a subset of contacts. While Area Charts generally suffer from occlusions, in our case the secondary curve never occludes the principal one, since it is always lower.

By default, the meeting pressure is displayed as a smooth curve. Feedbacks from our test users encouraged the use of smoothing, which they considered a helpful abstraction for the overall trend compared to weekly-specific values, as well as more aesthetically appealing. We use a cubic spline that never overshoots in the y-axis and passes through any of the underlying data points [45] since smoothing suggests a trend between adjacent data values. We also allow the user to disable smoothing and switch to a stepwise constant curve any time they might seek weekly-specific figures.

Individual meetings are represented on the chart as small dots, distributed chronologically. No other information about these meetings is displayed, unless the user highlights a contact on one of the other views, triggering the display of floating annotations with the titles of the related events. To avoid clutter, we heuristically choose a subset of annotations to be displayed, but we allow the user to explore the calendar by moving the cursor along the timeline.

5.1.3 Periodicity Analysis

The Meeting Periodicity Analysis, shown in Fig. 4 and positioned on the right-lower part of the panel, provides the highest level of detail for the medium and long-term analysis of meeting time. The chart aims at providing the user a tool to analyze how *frequently* and how *regularly* contacts have been met in previous months (T1.3), as well as to enable the comparison between contacts (T1.4). From the user perspective, the overarching goal of this analysis is to evaluate how much attention, broadly intended as interest and support, is dedicated to contact groups and people within these groups. While surely not the only characterizing factor, in this context we assume that more frequent and more regular meetings with a person are an indication of more attention dedicated to this person. Nevertheless, metrics like meeting frequency and distribution are not easily accessible in traditional calendar systems, unless manually computed.

Frequency vs. Periodicity. The visualization of frequency and regularity rises some design and readability challenges. On the one hand, frequency, formally defined as the number of occurrences of a repeating event per unit of time, requires to choose such unit in a meaningful way. Days, weeks, months or quarters are some of the possible candidate units. In order to enable an effective comparison between contacts, the same unit needs to be applied to all. This might cause some frequency values to be lower than one (e.g., *0.25 meetings/week*), which is arguably difficult to interpret in terms of absolute time intervals. For this reason, on the vertical axis we plot the *period* instead, expressed as time interval between two consecutive events (e.g., *one month*). As reciprocal of frequency, the period provides the same information while being better suited to display a wider range of values (one day to one year, including *Never* to indicate an infinite period).

Visual Mapping. In the Periodicity Analysis, we represent each contact with a circle colored according to the group or domain the user belongs to. The initials of the name are shown in order to facilitate the identification of the person of interest. On the vertical axis, each contact is positioned according to its average scheduling period, i.e., the average interval between two consecutive meetings. The average period is computed as the total number of days in the active time interval divided by the total number of meetings that occurred in that same interval. This definition holds also for those contacts who have been scheduled only once in the considered interval. Periods on the axis are labeled in a human-readable fashion (e.g., *every day*, *every week*, *every month*, etc.). The lowest line is reserved for contacts that have never been scheduled in the selected interval (*Never met*), corresponding to an infinite period. Name, time since the last meeting, meeting average duration and target score are displayed for each contact on demand (Fig. 4).

Meeting Regularity. On the other hand, meetings are not necessarily a periodic process, (with the exception of those belonging to a recurrent series), which leads to the notion of "regularity". While regularity has no unique scientific definition, here we define meetings as regularly scheduled when the number of days between two consecutive ones is consistent over time. If meetings occur with little regularity, the average period can be misleading. Low regularity, and in particular too long time intervals between consecutive meetings, are an important warning flag to be visualized in this context (G2). In the Periodicity Analysis, we convey the notion of regularity by plotting the discrete period distribution, for which we evaluated several visualization options. Min Max Average Plots, which use a line to show the minimum and maximum of a distribution and a dot to mark the average, are easier to read but only provide limited information about the distribution. On the other hand, Box Plots [46], although widely adopted in the scientific community to visualize minimum, 25th percentile, median, 75th percentile and maximum in a compact space, are not intuitively readable without training [47]. Furthermore, it has been shown that error bars, like those frequently associated with Bar Charts, are easily misunderstood by the general audience, and should therefore be avoided in favor of an explicit representation of the distribution [48].

Our final visual encoding has been inspired by Violin Plots [49]. First, we discretize periods into bins of 72 hours. Then, we plot each bin as semi-transparent circle whose size is proportional to the number of periods that fall into the bin, normalized to sum up to one. The radius is chosen according to Stevens's power law for visual areas [50], with a coefficient of 0.71 ($r \propto x^{0.71}$, where r is the radius and x is the normalized period count). Circles that happen to be

close to each other fuse together with a Gooey effect. The final result provides the information of a Violin Plot with the difference that in our case spikes in the distribution produce more sharp responses in the visualization. The horizontal layout allows us to fit a large number of calendar contacts on the screen, to facilitate their comparison. Horizontal scrolling enables scaling beyond the screen space.

Horizontal ordering. The user can sort this plot at any time by average period (default), variance, date of the last meeting, group or by a score which indicates how well the user is performing against a personal goal, as further defined in Section 5.3. These dimensions have been selected based on our tasks (T1.3, T1.4): the default ordering allows users to easily rank and compare how frequently people have been met; sorting by variance enables the identification of people whose meetings have been scheduled with higher irregularity; sorting by date of the last meeting allows users to quickly identify people with whom no meeting has been scheduled for a long time; sorting by lower score allows users to identify people that require attention based on the user's goals; finally, sorting by groups allows users to analyze the previous metric within the same group of contacts by co-locating its members. This is possible thanks to the use of a stable sorting algorithm, which allows users to sort by groups while maintaining the previous ordering.

5.1.4 Contact Diary

Contrarily to the previous views, which provide aggregated statistics, the Contact Diary is designed to provide details for specific contacts. Fig. 1 (D) shows the contact diary next to the calendar view. Besides being the main location for all operational information regarding a contact's entry including name, address, company, phone number and email addresses, the Contact Diary also provides a timeline that summarizes past and future meetings with a specific person. The duration of each meeting is represented through circular gray glyphs, each circle representing one hour or part of an hour, enabling the pre-attentive processing of the durations during fast scrolling. Events are spaced proportionally to the time interval between them, with a cap at six months to contain outliers. The goal is to enable users to visually capture the regularity of meetings, rather than to construct a precise chronological timeline. Intervals longer than the user's pre-defined goal, if set, are marked with a warning sign to highlight violations.

The user can also add private notes and comments, displayed chronologically together with the calendar events. This view is designed for everyday use, as a tool to quickly recall when previous meetings occurred, what was discussed and whether future follow-up meetings are scheduled already. It serves the overall goal to ease the preparation of an upcoming meeting and provides context for it (T1.5).

This view is deeply integrated into the calendar workflow and is accessible from both the calendar view and the Analytics Board, serving as a bridge between the two.

5.2 Interaction

Meeting Breakdown, Meeting Pressure and Periodicity Analysis are coordinated through *linking and brushing* [51] to provide context on demand. Hovering over one of the segments of the Meeting Breakdown reveals details about a specific group or contact (Fig. 2) and triggers the plot of a second Meeting Pressure curve, while some of the event titles are displayed (Fig. 3). The selected contact or group is highlighted in the Periodicity Analysis as well, while filtered-out elements are displayed with reduced opacity to provide contextual information. If a contact is clicked in any of the views, the Contact Diary is displayed on the left side.

5.3 Goals Definition and Review

While visualizing aggregated statistics is a fundamental step towards building self-awareness, we also want to allow users to define their own, personalized meeting goals (G2). Contacts can be organized into groups, where each can be assigned a specific goal. By setting a goal, the user defines the desired maximum number of days between two consecutive meetings with each person belonging to the group.

Since our goals represent average periods, a natural location to display them is on the Periodicity Analysis. Here we display goals as horizontal dashed lines. The relative position of a contact with respect to the line indicates how the meeting history relative to the contact has fulfilled the predefined goal: all contacts under the goal line should be scheduled more frequently (see blue line in Fig. 4).

To engage users towards improving their scheduling habits, we compute an aggregated goal achievement score, in a scale that ranges from 0% (no goal achieved) to 100% (all goals achieved). The motivation for an aggregated score is to provide users a unified measure of their own performance. We compute this score by comparing the meeting period distribution, the same visualized on Periodicity Analysis, with an ideal distribution: since we assume that regularity is desirable, the benchmark distribution is a delta around the desired period but we allow a 10% delay to take into account possible schedule interferences that normally occur in the everyday life. Ultimately, the final score is computed by averaging all meetings and all contacts for which a goal has been assigned. We display the score computed in the time period selected by the user, together with variation with respect to the previous period with same length (T2.1).

In order to identify actions to improve towards the target goals (T2.2) the user can consult the Periodicity Analysis as described above or, alternatively, rely on the visual reminders that we place directly on the main calendar view.

5.4 Navigation

Fig. 5 illustrates the navigation. Launching the application, the user first sees his/her weekly agenda. On the left sidebar, we display a prominent but not obtrusive visual reminder for people that are late in the schedule with respect to the predefined goals (1). This has been added on a second design iteration based on user feedback.

Whenever the user clicks on a reminder, the details of the contact are displayed together with the Contact Diary, while keeping the agenda visible (2). Similarly, the Contact Diary is displayed when a scheduled meeting is selected on the weekly agenda. In this case, the diary refers to the attendees of the event. The Contact Diary provides an entry point for the meeting analysis, specifically for the selected contact. For more details about the past scheduling history and to get an

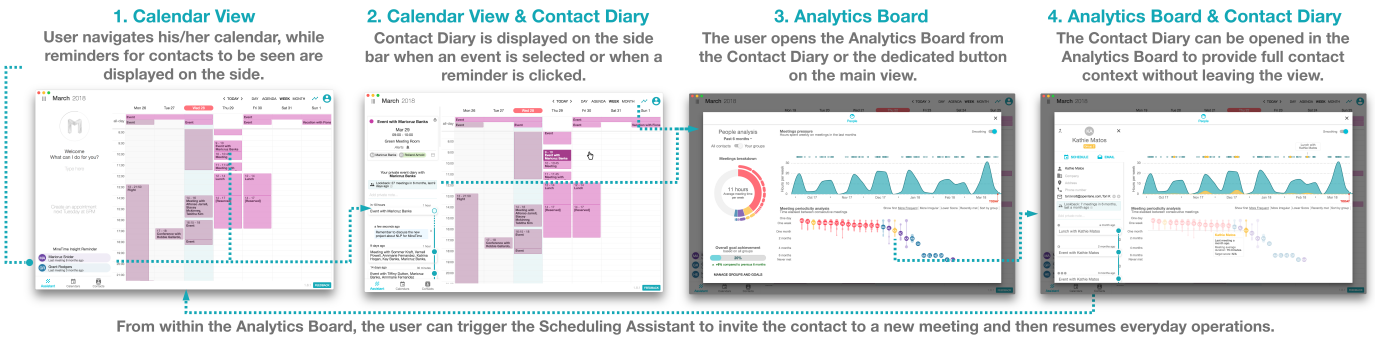

Fig. 5. Example of navigation between operational and analytics views. The user is alerted about people that should be scheduled (1). Details about these contacts are consulted in the Contact Diary (2). The user is encouraged to consult the Analytics Board for a more general overview (3). Necessary scheduling actions are triggered from the Analytics Board (4) and the user is forwarded back to the calendar and the scheduling assistant.

overview of groups and contacts, the Analytics Board (3) can be accessed directly from the Contact Diary (4).

Within the Analytics Board, the user explores statistics in aggregated form. By clicking on a contact, in the Meeting Breakdown or in the Periodicity Analysis, the Contact Diary appears on the left side. From here, the user can either go back to the analysis, email the selected person or directly schedule a new meeting. This last option, will bring the user back to the calendar view (1) and trigger MineTime's integrated scheduling assistant to find an appropriate time.

5.5 Implementation

MineTime Insight is integrated into MineTime, a fully-featured calendar application that provides natural language interaction and a scheduling assistant that is currently subject of active research. MineTime and MineTime Insight can be used on Windows, macOS and Linux and connect to all major calendar providers. All data is stored and processed locally to fulfill data protection regulations. These choices generally aim at satisfying our deployment challenges (G3).

6 Evaluation

6.1 Informal User Study

In order to better evaluate the achievement of our Design Goals and the overall readability of the tool, we conducted an informal user study.

Procedure. We scheduled a second in-person interview with the 8 participants that had taken part in the preliminary discussion. For this test, we showed an anonymized analysis of the same calendar data to all participants. Each study session lasted 30 minutes and was audio and screen recorded. In the first 8-10 minutes, participants were given an overview of the research project, including a brief description of the tools provided by our system. During the overview, the system was operated by us as we walked the participants through the Analytics Board and the Contact Diary. The explanation provided for each view was minimal: we did mention what information each plot was designed to provide, but we did not explain how to read it. After the introduction we handed over the tool to the participants, leaving them a few minutes to explore it freely. In order to collect readability feedback, we encouraged the participants to share their own interpretation of the different views.

Finally, we asked the participants to solve specific visualization tasks. These tasks included finding how many hours per week have been dedicated to a given group of people or to a specific person; identifying the busiest week of the last year, locating two people in the Meeting Periodicity Analysis, comparing frequency and regularity of these and identifying outliers within specific groups of contacts. We expected these tasks to serve as active learning of MineTime Insight and we encouraged participants to provide their feedback while exploring the tool. Overall, each participant operated the tool for 15-20 minutes. After the session, we asked the participants to fill a short questionnaire to express their agreement with 8 sentences on a 7-point Likert scale (1: strongly disagree, 7: strongly agree).

Findings. The results of this questionnaire are reported in Table 6.1. Overall, all participants confirmed that our data representation was easy to understand after a short introduction to the tool (Q1: mean=6.1; σ=0.6). Although EX1 argued that more training might be necessary for some people (*"As this is a professional tool, one could organize training sessions for administrative assistants"*), after a few minutes of exploration, recognized that a guided on-boarding when the tool is launched for the first time might suffice to grasp the goal of the different views. Taking into account this feedback, we decided to add a descriptive subtitle to the Meeting Pressure and the Periodicity Analysis in our final version.

MineTime Insight was also unanimously judged aesthetically pleasant (Q5: mean=6.8; σ=0.4). Although we did not ask questions specifically on engagement, PA2 commented the animations: *"These transitions are fun."*.

We also asked whether the tool could potentially help in analyzing scheduling decisions retrospectively (Q2). The participants were generally positive, but while Executives and Personal Assistants expressed particular interest in this respect, some Staff members pointed out how their meeting schedule might not be busy enough to take advantage of this kind of analysis: *"I only have many meetings during short periods of the year"*, and *"I would rather like to be able to keep track how productive meetings are for retrospective analysis, rather than whom I meet with, because I do not meet many different people"* (SM4). Question Q3 asked participants to imagine whether the regularity of future meetings could be improved by the use of this tool. In this case, PA1 showed skepticism: *"We tried in the past to enforce regular meetings using a rolling-based approach. However, there was not always something to be*

Question	EX1	EX2	PA1	PA2	SM1	SM2	SM3	SM4
Q1. It was easy to understand the presented information	5	6	6	7	6	6	7	6
Q2. The tool could help me in analyzing past scheduling decisions retrospectively	6	7	6	6	7	7	7	5
Q3. The tool could help me make sure meetings are scheduled more regularly	7	7	3	7	5	6	7	7
Q4. Overall, MineTime Insight is a useful tool	7	7	6	7	7	7	7	7
Q5. Overall, MineTime Insight is aesthetically pleasant	6	7	7	7	7	7	7	7
Q6. It was fast to find the information requested in the practical tasks	6	6	6	6	5	6	6	7
Q7. I would likely use this tool, if it was integrated with my current calendar application	7	5	6	7	7	7	7	7
Q8. I would likely use this tool, if it required to install a new dedicated application	7	7	4	6	4	7	7	6

Legend: *strongly disagree* 1 2 3 4 5 6 7 *strongly agree*

TABLE 1
Results of the User Study on a 7-point Likert scale. EX: Executive, PA: Personal Assistant, SM: Staff member (various roles)

discussed and eventually we dropped this idea. We haven't found a good solution yet to solve this issue [of too many weeks between consecutive meetings]". PA2 mentioned a similar experiment: *"We had some blocks of time reserved for these meetings every week but it failed. I guess it was partially because some of these meetings were not productive and partially because it was annoying to track the next person in queue with the tool we were using"*. On the other hand, he/she confirmed that low regularity and delayed meetings had been a concrete problem in the team and recommended visual alerts of critical situations, which we included in the final design iteration.

To conclude, we asked the participants how likely they would adopt this new system in two scenarios: as part of their daily calendar application (Q7) or as a separate dedicated application (Q8). The integration on existing calendar tools is perceived as more appealing (Q7: mean=6.6, σ=0.7) than as a separate service (Q8: mean=5.8, σ=1.2).

6.2 Case Study

To better illustrate the utility of MineTime Insight and its potential in improving scheduling management on a real scenario, we conducted a case study. The discussion that follows is based on the agenda of EX1, who works for two companies and collaborates with several others. During the preliminary interviews, EX1 confirmed the difficulty in tracking time dedicated to different companies and to people in different roles within these companies. The major difficulty seems to arise with people that are expected to be scheduled less often or on an on-call basis. In this case, the lack of a pre-scheduled recurring series of meetings often leads to excessively long periods of time between two consecutive updates with a specific person.

Procedure. We connected MineTime Insight to the Manager's real calendar and engaged in a conversation with EX1, as well as with PA1 and PA2, who manage this calendar on a daily basis. These informal conversations, done individually with each participant and lasting 20 to 30 minutes, happened after the informal user study but as part of the same sessions.

For the sake of generality and due to evident privacy concerns, company, person, domain and group names have been anonymized in all the figures, as well as in our discussion. The anonymization procedure includes replacing the event titles with basic keywords (e.g., event, meeting, lunch, call, etc.) and all contact names with fictitious ones, making sure the mapping between the real and fictitious names stays consistent over time. This is fundamental to preserve the original data distribution for each person in the anonymized data. In the remainder of the paper, we refer to companies generically as Company A, Company B, etc., and to contact groups as Group 1, Group 2, etc. Groups potentially interesting for the analysis emerged during the preliminary interviews with EX1 and we setup the analytics board with these groups before starting the discussion.

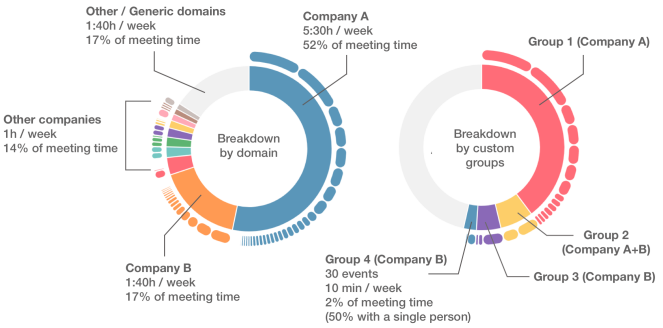

Fig. 6. Time Breakdown for the considered case study. Domain-based partitioning of all contacts (left) reveals the imbalance of time invested in Company B with respect to Company A. Group-based partitioning (right) highlights how Group 4 are allocated only 2% of the meeting time.

Analysis of Calendar Data. The calendar we consider contains a total of 1166 events (excluding full-day events) over 12 months, equivalent to an average of 22 events, or 39 hours, per week. Out of all events, 648 (55%) are meetings and they represent 32% of the total event time. Weekly, this corresponds to 12 meetings and an average of 10 hours. Overall, 256 unique attendees took part in those meetings, and 107 (42%) of them participated in two or more distinct ones. The number of attendees for each meeting varies significantly, from a minimum of 2 (*one-to-one meetings*) to a maximum of 19. Most meeting attendees are from Company A and Company B, and they cover a variety of professional roles and different positions in the company organigrams. There are also contacts from other, external partners. Overall, the calendar satisfies our quality assumptions: EX1 estimates that 99% of the meetings are actually captured in the calendar and the great majority of them include attendee metadata.

Findings. The Time Breakdown (Fig. 6) reveals a significant imbalance between time assigned to Company A (52%) and to Company B (17%). About 14% of the time is dedicated to other companies while the remaining 17% represent an aggregation of other domains, including generic ones like *gmail.com* or *yahoo.com* that cannot be associated with any

company. Although the imbalance was qualitatively known, EX1 was not aware of its extent and expressed concern about the small percentage of time dedicated to Group 4, as highlighted by the group-based analysis.

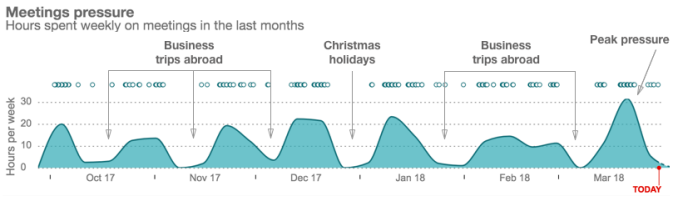

Fig. 7. Meeting Pressure over 6 months, reaching up to 30 hours/week.

The Meeting Pressure in Fig. 7 is affected by the partial unavailability of calendar data during business trips, managed on a different calendar we had no access to. While data originating from multiple calendars can be analyzed concurrently with MineTime Insight, PA2 noticed that being able to apply the analysis on this specific calendar might be useful to have an overview of traveling patterns. Because of the partially missing information, it is clear that a global average of weekly hours spent on meetings, as provided by the Meeting Breakdown, is not particularly meaningful. This is also true when the working time is inherently not well distributed over time. The Meeting Pressure, on the other hand, unrolling on a timeline, shows that meetings can easily take 20 hours a week, with a peak of 30 hours in March.

Finally, we asked our study participants to examine the Meeting Periodicity Analysis. Fig. 8 shows this visualization tool applied to three groups of particular interest, here called Group 1, Group 3 and Group 4. From the analysis, EX1 recognized that Group 1 has been scheduled for meetings more often and more regularly with respect to the other two categories, likely due to the fact that most of the meetings in this category are part of a recurring series. Furthermore, the analysis of single persons within the same group reveals irregularities that could potentially raise a warning flag. For example, within Group 3, one person had never been scheduled in the analyzed time window of six months. Similarly, among members of Group 4, only three of them have been met more than once. Less of a severe issue, the plot also reveals a positive outlier in this category, with one person that has been met significantly more frequently than all others. From the discussion, we learned that EX1 was aware of the outlier and could explain the motivations for it.

Feedback. Finally, we asked the three participants whether they noticed patterns and/or meeting distributions they were not aware of. Both assistants explained that they were already partially aware of the general time distribution and of the critical cases. On the other hand, they also explained how this requires a time effort from their side, which could be reduced by MineTime Insight: *"At the end of each year, we manually go through the agenda of the past twelve months and we roughly estimate the time dedicated to different activities. This is the reason why for me, personally, there was not much new qualitatively."* (PA2). On the other hand, EX1 told us that it was rather surprising to see some of the results and expressed great interest in MineTime Insight: *"Such tool, just on meetings and people, is extremely helpful in management."*. EX1 also suggested to extend it to track time spent on different projects.

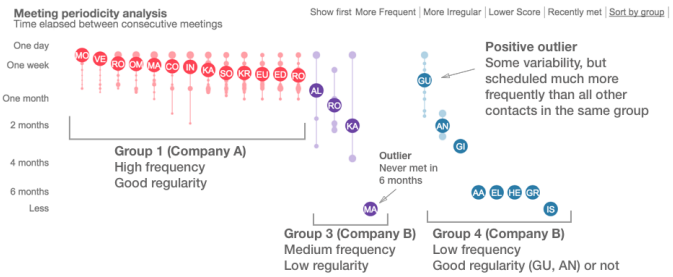

Fig. 8. Periodicity Analysis applied to group-based analysis. The color of each group corresponds to those in Fig. 6.

6.3 Field Deployment

While face-to-face interviews allowed us to engage in informative conversations, a controlled environment does not always allow us to evaluate whether a system meets a casual user's goals and needs over time [40]. Moreover, it is often biased by *experimental demand effects* [52] - well-known tendency of lab studies' participants to offer positive and favorable feedback to the researchers' questions. We therefore conducted a survey-based field study to assess user engagement and system performance "in the wild".

Procedure. We shipped our analytics tools with the latest version of MineTime, making it available for free to users worldwide. Over a period of two months, a survey has been automatically proposed to all users who engaged with MineTime Insight at least once. The participation to the survey was voluntary, anonymous and not remunerated. We collected 65 answers, from as many users that completed the survey. As part of contextual information, we asked the participants to estimate how many meetings, on average, they had per week, whether they were using MineTime for their personal or work calendars and their profession.

As for the internal user study, we asked to rate five statements (Q1-Q5) on a 7-point Likert scale. We used the same statements of the previous study, but we excluded Q6, as we could not design equivalent tasks with no information about the user data, and Q7-Q8 as the users were already using a dedicated application. We finally encouraged users to provide an optional comment to support their rating.

Findings. The participants only using MineTime for their personal calendars (27%) expressed little or no interest in our tool (*"I am using MineTime for personal use. Analytics are not relevant."*, *"This is being used for a personal account. It would be more useful if I were using it for work."*). Quantitatively, the perceived utility of the tool (Q4) was low (mean=2.5; σ=2.1) compared to the remaining users (mean=4.6; σ=2.0). On the other hand, this comes at no surprise as the tool was designed for the corporate environment. This also confirms our initial assumption that meeting analysis is not particularly relevant in the private context. The analysis that follows focuses on the participants using MineTime, at least in part, for professional reasons. Fig. 9 summarizes the results of the remaining 47 answers, including two significant sub-groups: i) people with intense meeting activity (defined as having 11+ meetings/week on average) and ii) people declaring to be a Senior Manager or Executive. Notice that a specific user could be considered in both sub-groups. We could not

collect enough answers from Personal Assistants to report statistically significant results about this category.

Compared to the lab study, we observed a higher deviation in the ratings for all five questions, which could be explained by a more diversified background, pool of professions, age or expertise of the participants. Overall, ratings were mostly positive and in line with our previous findings, but we also observed more critical feedback. Compared to the average, managers and executives found our tool significantly more likely to be helpful in scheduling meetings more regularly (Q3) and in general more useful for their time management (Q4). They gave slightly higher scores also regarding the readability of the visualizations (Q1), the utility in analyzing past scheduling decisions (Q2) and the visual appearance (Q5). Interestingly, people declaring 11+ meetings/week did not value the utility of MineTime Insight (Q2, Q3, Q4) significantly different than the average professional user. For this sub-group, we also reveal below-average ratings for both readability (Q1) and aesthetic appeal (Q5). We hypothesize that, with more data displayed in the Analytics Panel, some users might feel overwhelmed.

We were interested in understanding the reasons for the most critical feedback. We found that negative ratings come mostly from sole practitioners, self-employed people, researchers and students. Overall, we identified a number of motivations across different roles: 1) **number of meetings** (*"I don't have a lot of meetings, so there isn't much to glean"*); 2) **lack of interest** (*"I'm not interested in keeping track of how much time I spend in meetings"*, *"Fine, but that's not what I use MineTime for."*); 3) **lack of motivation** (*"It is interesting, but I am not sure what to do with it"*, *"I work in a small team, who all are in the same room. The rest of my meetings are incidental, or company-wide."*, *"I normally do not use this type of screen"*); 4) **security** (*"Interesting, concerned with privacy"*) and 5) **technical** (*"It's nice however its very very slow and makes the app almost unusable when you have a ton of events"*, *"It's an amazing tool! However it isn't perfect. I'm at college, so I don't attend all club meetings, but MineTime assumes I have."*).

On the positive side, users mentioned that MineTime Insight can help them to *"have more effective meetings"* (manager), *"plan better, have smarter meetings"* (executive) and *"find space in the week to focus on work"* (team leader). Finally, some people suggested possible extensions: *"It would be more helpful to have the tool track how much time I've spent on specific projects"* (student), *"Please, make it not only for meetings but for any type of activities"* (executive).

7 Discussion

In this section, we discuss what we learned along the process of designing MineTime Insight, as well as from the user studies, including limitations and ideas for future work.

7.1 Lessons Learned and Design Implications

Perceived utility. Even though MineTime Insight was designed to be used by all employees in the corporate environment, the perceived utility is clearly shifted towards users with a managerial role. On the one hand, this comes at no surprise since managers and executives, besides having more filled agendas, might feel responsible for the success of the employees they supervise, giving them an additional motivation to keep track of their meeting activity according to best practices [26]. The value for Personal Assistants seems to derive more from the time saving brought by the automatic computation of some metrics rather than from the insights themselves, often qualitatively already known to them. Some Staff members, on the other hand, showed lack of motivation for self-reflection (*"Fine, but that's not what I use MineTime for"*) or no interest (*"I only need to know where to show up"*). Some employee might not have an intense meeting activity as part of their role which, not surprisingly, can explain a low perceived utility. Another explanation might be users not being aware of the potential benefits of more regular meetings or the importance of their preparation [1], [26]. This can be addressed with designs more mindful of the user's awareness of the motivations supporting the analysis. Finally, we notice that some employee see value in tracking their meeting activity but have no need to scale to hundreds of calendar contacts. In this case, a design tailored to one or few contacts (e.g. the supervisor or the team), might be more appropriate. Shneiderman et al. [53] theorized an user interface that adapts to the specific user role, where each role has a vision statement that reminds the users of their goals, which is an interesting direction to investigate.

Data quality. In this work, we assumed high quality calendar data (i.e. most meetings tracked in the calendar, attendees explicitly listed in the metadata). When the calendar is managed by a Personal Assistant, as for the calendar showcased in Section 6.2, this is often the case but the assumption does not hold in general. If too many meetings happen extemporaneously and are not recorded, the statistics we provide lose significance. This problem could be mitigated by providing easier and more engaging ways to track time in order to encourage a more consistent tracking. On the other hand, many users do not register the meeting participants either. We believe this happens because of 1) lack of motivation, as the operation might not provide additional value and 2) the design of most calendar applications, which bind this operation to the delivery of meeting invitations, not always desired. We also found several cases when meetings are scheduled in the calendar but not actually attended. While most calendar applications allow to flag events as "free time" to indicate that the user will not take part in the event, we found that only a minority of people use this option consistently. We hope that the availability of tools like MineTime Insight could motivate users to use their calendar more rigorously. This would facilitate data analysis by a large extent and, in turn, be beneficial to the users.

Attentional demand. One of our primary goals was to make the tool accessible to a broader audience, including people with no visualization expertise (G3). During our interviews, the participants had no difficulty reading the Meeting Breakdown and the Meeting Pressure, but the Periodicity Analysis required significantly more time and additional explanation from our side to let users interpret it effectively. As part of the iteration process for this tool, we added a vertical line spanning the periodicity range, which helped users to recognize it as a distribution. We also included a line of textual explanation above the plot and removed all statistical terms replacing, for example, *"Sort by: Variance"*

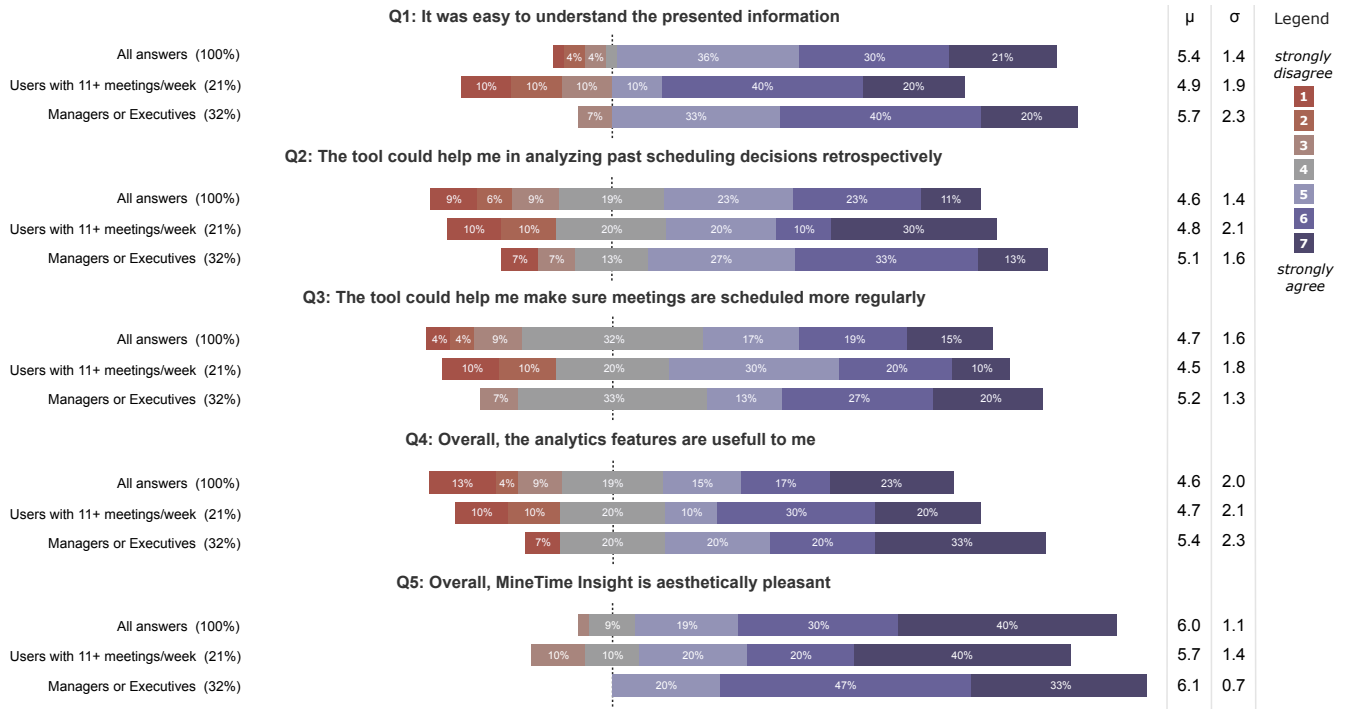

Fig. 9. Results of the field study. Only the users who declared to use MineTime for their professional calendar has been considered (47 answers = 100%). We report the mean (μ) and standard deviation (σ) over all the answers, as well as for two relevant subsets for comparison.

with *"Show first: More Regular"*, a simple but quite important usability expedient that had been overseen during the first design iteration. At the same time, we realized that simpler views, those adopting well-known plotting techniques and able to provide quick insights like the Meeting Breakdown, are particularly appreciated by most users, as they can learn basic information with virtually no investment in time and effort. Especially when dealing with the constraints of PVA, we argue that the design of professional tools might require extra care to balance attentional demand, aesthetic appeal and accuracy, as discussed earlier, for instance, for the design of the Meeting Pressure visualization.

Beyond actionable data. The majority of PVA designs proposed in the literature aim at satisfying user curiosity, reminiscing about past experiences or sharing them with others [6]. An important takeaway from our study is that our users wanted their visualizations to be actionable, i.e. reveal possible time management problems and solutions at a glance. While curiosity and social sharing might be a initial motivation for use (in fact, we learned that two users publicly shared a screenshot of their analytics dashboard on social networks), we argue this is likely not sufficient to promote long-term usage and behavioral change. Some users suggested the tool could produce textual or visual suggestions proactively and to place them in a prominent position: *"It would be nice to have some kind of visual alerts for people that need to be scheduled"*. Providing actionable metrics was part of our goals (G2). However, this feedback suggests that providing actionable insight might not be sufficient if action recommendations are not pushed *beyond* the analytics dashboard and integrated seamlessly into the tools that people use everyday as part of their routine.

7.2 Reflection on the Scope of PVA

We argue that MineTime Insight qualifies as one the "infovis edge cases" as defined by Pousman et al. [7]. Despite targeting the professional context within a working environment, we note differences compared to traditional infovis systems. In particular, using the taxonomy of Casual Infovis [7]: 1) the user population includes a wide spectrum of users, from experts to novices when it comes to data analysis and data visualization; 2) the *usage pattern*, although ultimately productivity-oriented, expands beyond the main work-role expectations as users are not professional "meeting analysts"; 3) the *data type*, coming from the user's own calendar, is typically considered personal for all but the Personal Assistants; 4) the *insight* is meant to build awareness about personal scheduling habits and interaction with other employees, rather than crystallized analytical conclusions.

The taxonomy for Personal Visual Analytics by Huang et al. [6] further characterizes these ideas, building around the notion of *personal context*, "non-professional situations, in which people may have quite different motivations, priorities, role expectations, environments, or time and resource budgets as compared to professional situations" [6]. Strictly speaking, this definition does not apply to our case, as we target professional situations. Yet, we have previously discussed how all other characterizations are well suited to describe our target users, their expertise and motivations, as further supported by the user studies. For example, as meeting analysis is normally not part of the work-role expectations, average users might not find good motivations to analyze past scheduling decisions or reflect on their own habits, which is captured by some of the feedback (*"why do I care?"*, *"not sure what to do with it"*). On the other hand, we also recognize possible differences. For example, some people may be requested to analyze their meeting

behavior without having an intrinsic curiosity in it, and it is also more likely that the data analysis may need to be shared with coworkers or managers.

Overall, our design inherits from the lessons learned by others PVA applications, where the most effective factors for adoption have been personal interest, curiosity, personal challenges, control, fantasy, boredom, aesthetic appeal, and novelty [6], [40]. Surely, MineTime Insight is not the only application where the constraints and adoption mechanisms mentioned above apply to the working environment. We believe the fields of PVA and Casual Infovis would benefit from a more systematic analysis of the professional environment, possibly generalizing the definition of *personal context* to include professional situations where the objectives are awareness and self-reflection. Further research is necessary to assess when and how the mechanism of adoption in the professional and non-professional contexts can be fully assimilated. The former is likely to harden some constraints (e.g., time budget) or add new ones (e.g., data protection policies). How this affects the design, deployment, validation and adoption of new tools are open research questions.

7.3 Limitations and Future Work

Despite the encouraging results, our design study might not cover the vast variety of priorities, motivations and expectations of all corporate employees. A formal, methodical analysis of the preliminary interviews, possibly involving a larger user base, could have led to a more diversified set of Design Goals. Regarding the validation, the feedback was collected after a few minutes of interaction with MineTime Insight, which limits its information to first impressions. In order to quantitatively measure the overarching goal of improving short and long-term scheduling decisions, an analysis over year-long periods will be necessary, as well as a larger number of participants. For this, we will continue monitoring the engagement of the users worldwide.

This work focused on the specific goal of analyzing the meeting time. There are a number of questions that might be useful to improve time management practices that we did not consider in this work. Among these, whether meetings result to be productive or not, how the scheduled duration is chosen appropriately, how often meetings are canceled or rescheduled, how often meetings are organized by the user compared to somebody else, how many participants are usually involved and statistics about the location. In this work, we focused on the tasks that seemed more relevant, based on the literature and the interviews. Some metrics, like the average meeting duration, were only provided in textual form and might benefit from a visual encoding. We do not exclude the possible utility of other metrics in a different setting, which should be investigated in the future.

Besides meetings, the studies showed that some employees, including managers, are actually more interested in analyzing the time invested in specific activities and/or projects. This problem has been partially addressed in the past [20] but a number of research questions remain open. In particular, recording activities usually requires extra effort from the user side. In order to minimize this additional work, analytics solutions should likely be integrated as part of existing software (e.g., calendars) in order to streamline the process, which seems a natural extension to our work.

8 Conclusion

We presented MineTime Insight, an information visualization tool designed to fit in the daily workflow and targeting a broader audience of users with little or no visualization expertise. First, we discussed a novel characterization of the problem of monitoring meeting time in the corporate environment, showing how ideas from the realm of PVA well extend to this context. Then, we discussed the design choices with respect to our goals and showed how MineTime Insight can be used to learn interesting information about one's meeting habits in a case study. Finally, we discussed the results of the user studies and possible future extensions.

Acknowledgments

This work has been partially funded by the Swiss CTI grant 19005.1 PFES-ES.

References

[1] N. C. Romano and J. F. Nunamaker, "Meeting analysis: findings from research and practice," in *Proc. 34th Annual Hawaii Int'l Conf. on System Sciences*, 2001.

[2] Ovum. (2014) Collaboration 2.0: Death of the Web Conference (as we know it). [Online]. Available: https://bit.ly/2yzQBGt

[3] I. Li, A. Dey, and J. Forlizzi, "A stage-based model of personal informatics systems," in *Proc. SIGCHI Conf. on Human Factors in Computing Systems (CHI 2010)*, 2010, pp. 557–566.

[4] A. Rapp and F. Cena, "Personal informatics for everyday life: How users without prior self-tracking experience engage with personal data," *nt'l Journal of Human-Computer Studies*, vol. 94, pp. 1–17, 2016.

[5] R. G. Miltenberger, *Behavior modification: Principles and procedures*. Cengage Learning, 2011.

[6] D. Huang, M. Tory, B. A. Aseniero, L. Bartram, S. Bateman, S. Carpendale, A. Tang, and R. Woodbury, "Personal visualization and personal visual analytics," *IEEE Trans. on Vis. and Comp. Graphics*, vol. 21, no. 3, pp. 420–433, 2015.

[7] Z. Pousman, J. Stasko, and M. Mateas, "Casual information visualization: Depictions of data in everyday life," *IEEE Trans. on Vis. and Comp. Graphics*, vol. 13, no. 6, pp. 1145–1152, 2007.

[8] J. Schwarz, J. Mankoff, and H. S. Matthews, "Reflections of everyday activities in spending data," in *Proc. SIGCHI Conf. on Human Factors in Computing Systems (CHI '09)*. New York, NY, USA: ACM, 2009, pp. 1737–1740.

[9] V. Kalnikaite, A. Sellen, S. Whittaker, and D. Kirk, "Now let me see where i was: Understanding how lifelogs mediate memory," in *Proc. SIGCHI Conf. on Human Factors in Computing Systems (CHI '10)*. New York, NY, USA: ACM, 2010, pp. 2045–2054.

[10] S. Consolvo, D. W. McDonald, T. Toscos, M. Y. Chen, J. Froehlich, B. Harrison, P. Klasnja, A. LaMarca, L. LeGrand, R. Libby, I. Smith, and J. A. Landay, "Activity sensing in the wild: A field trial of ubifit garden," in *Proc. SIGCHI Conf. on Human Factors in Computing Systems (CHI '08)*. New York, NY, USA: ACM, 2008, pp. 1797–1806.

[11] D. Huang, M. Tory, and L. Bartram, "A field study of on-calendar visualizations," in *Proc. 42nd Graphics Interface Conference*, 2016, p. 3.

[12] E. Costanza, S. D. Ramchurn, and N. R. Jennings, "Understanding domestic energy consumption through interactive visualisation: A field study," in *Proc. 2012 ACM Conf. on Ubiquitous Computing (UbiComp '12)*. New York, NY, USA: ACM, 2012, pp. 216–225.

[13] F. Viegas, S. Golder, and J. Donath, "Visualizing email content: Portraying relationships from conversational histories," in *Proc. SIGCHI Conference on Human Factors in Computing Systems (CHI '06)*. New York, NY, USA: ACM, 2006, pp. 979–988.

[14] G. Barata, H. Nicolau, and D. Gonçalves, "AppInsight: What have i been doing?" in *Proc. Int'l Working Conf. on Advanced Visual Interfaces (AVI '12)*, New York, NY, USA, 2012, pp. 465–472.

[15] M. d'Aquin, A. , S. Elahi, and E. Motta, "Personal monitoring of web information exchange: Towards web lifelogging," in *Proc. Web Science Conf.*, 2010.

[16] B. Pagno and L. Nedel, "Everyday Visualization," in *Elec. Proc. IEEE VIS 2015 Workshop Personal Visualization: Exploring Data in Everyday Life*, 2015, pp. 1–4.

[17] J. Mackinlay, G. Robertson, and R. DeLine, "Developing calendar visualizers for the information visualizer," in *Proc. Symp. on User Interface Software & Technology*. ACM, 1994, pp. 109–118.

[18] J. Tullio, J. Goecks, E. D. Mynatt, and D. H. Nguyen, "Augmenting shared personal calendars," in *Proc. 15th annual ACM symposium on User interface software and technology*. ACM, 2002, pp. 11–20.

[19] A. Faulring and B. A. Myers, "Availability bars for calendar scheduling," in *CHI '06 Extended Abstracts on Human Factors in Computing Systems*, New York, NY, USA, 2006, pp. 760–765.

[20] C. Dugan, W. Geyer, M. Muller, A. N. Valente, K. James, S. Levy, L.-T. Cheng, E. Daly, and B. Brownholtz, ""I'D Never Get out of This !?$Redesigning Time Management for the Enterprise," in *Proc. SIGCHI Conf. on Human Factors in Computing Systems (CHI '12)*, New York, NY, USA, 2012, pp. 1755–1764.

[21] W. Aigner, S. Miksch, H. Schumann, and C. Tominski, *Visualization of time-oriented data*. Springer Science & Business Media, 2011.

[22] M. Brehmer, B. Lee, B. Bach, N. H. Riche, and T. Munzner, "Timelines revisited: A design space and considerations for expressive storytelling," *IEEE Trans. on Vis. and Comp. Graphics*, vol. 23, no. 9, pp. 2151–2164, 2017.

[23] F. Viegas, D. Boyd, D. H. Nguyen, J. Potter, and J. Donath, "Digital artifacts for remembering and storytelling: posthistory and social network fragments," in *Proc. 37th Annual Hawaii Int'l Conf. on System Sciences*, Jan 2004.

[24] H. Mintzberg, *Mintzberg on management: Inside our strange world of organizations*. Simon and Schuster, 1989.

[25] J. Grudin, "Managerial use and emerging norms: effects of activity patterns on software design and deployment," *Proc. 37th Annual Hawaii Int'l Conf. on System Sciences, 2004*, 2004.

[26] Gallup, Inc. (2013) State of the American Manager: Analytics and Advice for Leaders. [Online]. Available: https://www.gallup.com/services/182138/state-american-manager.aspx

[27] S. F. Ehrlich, "Social and psychological factors influencing the design of office communications systems," in *ACM SIGCHI Bulletin*, vol. 18, no. 4. ACM, 1987, pp. 323–329.

[28] ——, "Strategies for encouraging successful adoption of office communication systems," *ACM Transactions on Information Systems (TOIS)*, vol. 5, no. 4, pp. 340–357, 1987.

[29] J. N. Mosier and S. G. Tammaro, "When are group scheduling tools useful?" *Computer Supported Cooperative Work (CSCW)*, vol. 6, no. 1, pp. 53–70, 1997.

[30] L. Palen, "Social, individual and technological issues for groupware calendar systems," in *Proc. SIGCHI conference on Human Factors in Computing Systems*. ACM, 1999, pp. 17–24.

[31] A. E. Blandford and T. R. G. Green, "Group and individual time management tools: What you get is not what you need," *Personal and Ubiquitous Computing*, vol. 5, no. 4, pp. 213–230, Jan. 2001.

[32] L. Palen and J. Grudin, "Discretionary adoption of group support software: Lessons from calendar applications," in *Implementing collaboration technologies in industry*. Springer, 2003, pp. 159–180.

[33] B. Desruisseaux. (2009) Internet Calendaring and Scheduling Core Object Specification (iCalendar). (Visited 29/08/2018). [Online]. Available: https://www.rfc-editor.org/rfc/rfc5545.txt

[34] Microsoft Corp. Create appointments and meetings by using EWS in Exchange 2013. Visited 29/08/2018. [Online]. Available: https://bit.ly/2onkcON

[35] J. Grudin, "Groupware and social dynamics: Eight challenges for developers," *Communications of the ACM*, vol. 37, no. 1, pp. 92–105, 1994.

[36] J. T. Hackos and J. Redish, "User and task analysis for interface design," 1998.

[37] H. Mintzberg, "The nature of managerial work," 1973.

[38] E. A. Locke and G. P. Latham, "Building a practically useful theory of goal setting and task motivation: A 35-year odyssey." *American psychologist*, vol. 57, no. 9, p. 705, 2002.

[39] D. Barua, J. Kay, B. Kummerfield, and C. Paris, "A framework for modelling goals in personal lifelong informatics," in *Prof. CHI 2012 Workshop Personal Informatics in Practice: Improving Quality of Life Through Data*, 2012.

[40] D. W. Sprague and M. Tory, "Motivation and procrastination: Methods for evaluating pragmatic casual information visualizations," *IEEE Comp. Graph. & Applications*, vol. 29, no. 4, pp. 86–91, 2009.

[41] G. Ellis and A. Dix, "A taxonomy of clutter reduction for information visualisation," *IEEE Trans. on Vis. and Comp. Graphics*, vol. 13, no. 6, pp. 1216–1223, 2007.

[42] J. Stasko and E. Zhang, "Focus+ context display and navigation techniques for enhancing radial, space-filling hierarchy visualizations," in *IEEE Symp. on Information Vis.*, 2000, pp. 57–65.

[43] A. Luong and S. G. Rogelberg, "Meetings and more meetings: The relationship between meeting load and the daily well-being of employees." *Group Dynamics: Theory, Research, and Practice*, vol. 9, no. 1, pp. 58–67, 2005.

[44] J. P. Robinson, A. Chenu, and A. S. Alvarez. Measuring the complexity of hours at work: the weekly work grid. Visited 29/08/2018. [Online]. Available: https://bit.ly/2EB292W

[45] M. Steffen, "A simple method for monotonic interpolation in one dimension," *Astronomy and Astrophysics*, vol. 239, p. 443, 1990.

[46] D. F. Williamson, R. A. Parker, and J. S. Kendrick, "The box plot: a simple visual method to interpret data," *Annals of internal medicine*, vol. 110, no. 11, pp. 916–921, 1989.

[47] A. Bakker, R. Biehler, and C. Konold, "Should young students learn about box plots," *Curricular development in statistics education*, pp. 163–173, 2004.

[48] M. Correll and M. Gleicher, "Error bars considered harmful: Exploring alternate encodings for mean and error," *IEEE Trans. on Vis. and Comp. Graphics*, vol. 20, no. 12, pp. 2142–2151, 2014.

[49] J. Hintze and R. Nelson, "Violin plots: a box plot-density trace synergism," *American Statistician*, vol. 52, no. 2, pp. 181–184, 1998.

[50] H. Schumann and W. Müller, *Visualisierung: Grundlagen und allgemeine Methoden*. Springer-Verlag, 2013.

[51] D. A. Keim, "Information visualization and visual data mining," *IEEE Trans. on Vis. and Comp. Graphics*, vol. 8, no. 1, pp. 1–8, 2002.

[52] B. Brown, S. Reeves, and S. Sherwood, "Into the wild: Challenges and opportunities for field trial methods," in *Proc. Conf. on Human Factors in Computing Systems*, 05 2011, pp. 1657–1666.

[53] B. Shneiderman and C. Plaisant, "The future of graphic user interfaces: Personal role managers," in *Proceedings of the Conference on People and Computers IX*, 1994, pp. 3–8.

Marco Ancona received the B.Sc. degree in Electrical Engineer and Information Technology at University of Padova, Italy, in 2013 and the M.Sc. degree in Computer Science from ETH Zurich, Switzerland, in 2016. He is currently a PhD candidate at the Computer Graphics Laboratory (CGL) at ETH Zurich. His research interests include information visualization, machine learning and interpretable artificial intelligence.

Marilou Beyeler completed her B.Sc. degree in Computer Science at the ETH Zurich, Switzerland, and at the Imperial College London, in 2018. She is now working towards her M.Sc. degree in Computer Science at ETH Zurich.

Markus Gross is a Professor of Computer Science at ETH Zurich, head of the Computer Graphics Laboratory and director of Disney Research Zrich. His research interests include physically based modeling, computer animation, immersive displays, and video technology. He received a master of science in electrical and computer engineering and a PhD in computer graphics and image analysis, both from Saarland University in Germany in 1986 and 1989.

Tobias Günther joined the Computer Graphics Laboratory (CGL) at the ETH Zurich as a postdoctoral researcher in 2016. He received his M.Sc. in Computer Science in 2013 and his Dr.-Ing. (Ph.D.) in 2016 both from the Otto-von-Guericke University of Magdeburg. His research interests include scientific visualization, progressive light transport and real-time rendering.