Towards a Task-based Guidance in Exploratory Visual Analytics

Belgin Mutlu
Pro²Future GmbH, Austria
Graz University of Technology, Austria
belgin.mutlu@pro2future.at

Milot Gashi
Pro²Future GmbH, Austria
milot.gashi@pro2future.at

Vedran Sabol
Know-Center GmbH, Austria
Graz University of Technology, Austria
vsabol@know-center.at

Abstract

Exploring large datasets and identifying meaningful information is still an active topic in many application fields. Dealing with large datasets is currently not only a matter of simply collecting and structuring data for retrieval, but sometimes it also requires the provision of adequate means for guiding the user through the exploration process. Visualizations have shown to be an effective method in this context, the reason being that since they are grounded on visual cognition, people understand them and can naturally perform visual operations such as clustering, filtering and comparing quantities. However, systems which help us to create visualizations often require specific knowledge in data analysis, which ordinary users typically do not possess. To address this gap, we propose a system that guides the user in the data analysis process. To achieve this, the system observes current user behavior, tries to infer the task of the user and recommends the next analysis steps to help her to carry out the task.

1. Introduction

Recent advances in technology have resulted in an enormous increase of personal and industrial data. Although this data includes massive volumes of valuable information that can be used to better support both humans and machines, extracting and using this information to gain valuable insights and draw correct conclusions is a tedious and time-consuming task. As a result, relevant, useful information may be overlooked, possible links within the data not identified and wrong conclusions drawn. Visualizations have shown to be effective in dealing with huge datasets: because they are grounded on visual cognition, people understand them and can naturally perform visual operations such as clustering, filtering and comparing quantities. However, creating meaningful and valid visualizations is challenging and time consuming because it involves many steps [1, 2]. It is a complex process, which often requires domain knowledge, understanding of the data and knowledge about the user's goal and task [3]. The ordinary user, however, is not an expert in visualizations and can rarely generate a visualization, which provides meaningful message about data.

There exist several approaches which automatically generate and recommend appropriate visualization [4, 5, 6, 7, 8]. What they all have in common is that these approaches follow the visual encoding rules and perceptual guidelines to define, which visualization is more appropriate for the given dataset. These rules are basically formal specifications for what might be achieved when representing data visually. While these approaches are successful in reducing the knowledge barrier, they are restricted to a certain domain or technology, and more importantly, they ignore the fact that the choice of a visualization depends to a great extent on the user's visual goal and tasks [9].

The existing research covers a broad range of visual analytics tasks [10, 11, 12] a user may seek for. However, the most frequently performed visual tasks are: *gaining overview, analyzing outliers (anomalies), analyzing trends and comparing variable distribution*. Furthermore, a large number of techniques (=interactions) are available that visualizations provide and users can make use of to analyze their data: *comparing quantities, filtering data points, aggregating dimensions, zooming, coordinated views* etc. Basically, a goal or task requires the user to perform a sequence of these interactions and not a single one. With regard to Shneiderman [10], an overview task includes zooming out views of individual data points, as well as panning or scrolling through the whole dataset. Outlier detection on the other hand includes searching, filtering or querying a multivariate data[11, 13]. Knowing which visualization to choose and which interactions to perform, however, requires a significant level of visual and data analytical skills from the user. Unfortunately, average users do not possess these skills and may have serious difficulties in analyzing their data. Even domain experts within a specific area may be affected by the same issue.

HℱCSS

In order to address this gap, we propose a visualization tool that guides the user in the visual analysis process [14, 15]. To achieve this, our visual analytics tool attempts to infer which analytical task(s) the user is currently performing and guide her during this process by recommending the next analysis step or a sequence of steps (interaction(s), such as filtering, aggregating, zooming etc.). The main contribution of our tool in providing guidance to the user on the visual analysis process are: (i) our tool hides the complexity behind data and visual analytics, and (ii) it supports the users in readily gaining insights out of their data and draw conclusions, which might remain hidden by other means of analyzing them.

2. Related work

A task-based visual analytics (VA) tool aims to identify users' intention behind visualizing the data [16]. Basically, users' visual intention can be categorized into two groups: exploratory and confirmatory (=verification) [17]. The former describes an analysis process in which the user aims to summarize the main features of the data and uses them to define preliminary hypotheses. The later, however, describes the process in which the user tests her found or assumed hypotheses from the previous analysis. Depending on the stage of analysis, users rely on different kinds of visualizations. Hence, a task-based VA tool is usually provided not only with a single, but with a series of interactive visualizations to support the users through the exploration and verification process and thus in gaining new insights and findings from their data. In the following we discuss the relevant work conducted in the area on guided visual analytics tools.

Casner [18] introduces BOZ, which models the visual tasks as a set of logic rules. These rules are used to define equivalent perceptual operators and to create appropriate visualizations. BOZ does not provide guidance in performing a task, it is a relatively logical procedure that aids the definitions of appropriate visualizations. Similar to BOZ, VizAssist [19] supports the user in finding the most appropriate visualization, but for data mining tasks such as *classification*, *clustering* and *regression*. When the user selects a particular task, the system defines a list with relevant mapping combinations and presents them to the user as visual recommendations. The user can either accept the recommendations or go to the previous step and change the task. In this sense, VizAssist lacks the provision of a dynamic support for performing a mining task and supports the user only in exploring the mapping space. Voyager [7] provides guidance in defining appropriate visualizations

using explicit user feedback, which are defined via variable selection. Voder [20] uses interactive data facts to recommend visualizations and to support users in interpreting the recommended visualizations: e.g., when the user hovers on data facts, the corresponding parts of the visualization are dynamically highlighted. Similar to ViZAssist, these tools only address one user need which is *finding the most relevant visualizations for the underlying data* and they do not guide the user throughout a visual task. Cool et al. [21] present a set of design guidelines for the development of task-model-based visual analytics tools. However, the task-models are formulated in terms of specific data sources and do not adapt easily to new ones. Silva et al. [22] use an eye-gazed based recommendation model to guide the users in identifying time-series patterns. Similar to this, Shao et al. [23] track gaze movements to guide the user to the most interesting areas of large scatter plot matrices. Yet, these two approaches are highly interface-dependent as the use of an eye-tracker is essential in order to collect the gaze information.

Basically, the users can be considered to have different analysis behavior when performing exploration or confirmation activities/tasks [24]. Nazemi et al. [25] propose a system, which tracks user behavior (interactions on visualizations) to infer her goal/task and adapts a set of applicable visualizations on user intention behind visualizing the data. Given that the target of this approach are the subjects (i.e., bibliographical notes and publications) previously extracted from digital libraries, this tool is bound to a specific domain and its specific use cases. Gotz et al. [9] present a system, that logs current user's behavior and tries to match it to a set of interaction patterns. These patterns aid in identifying the intended task and in recommending alternative visualizations. The biggest limitation of this work, however, is that the interactions patterns are defined by experts a priori and thus do not reflect the individual user characteristics.

In this paper, we propose a VA tool that further improves the existing work along three main contributions: (i) our VA tool guides the user throughout a visual task by recommending the next analysis steps in form of interactions, (ii) the behavior patterns we use to identify the implicit tasks result from experimental observations reflecting the individual user characteristics, and (iii) our tool is not bound to a certain technology (e.g., eye tracker) or restricted to a single domain (e.g., digital libraries, life-science).
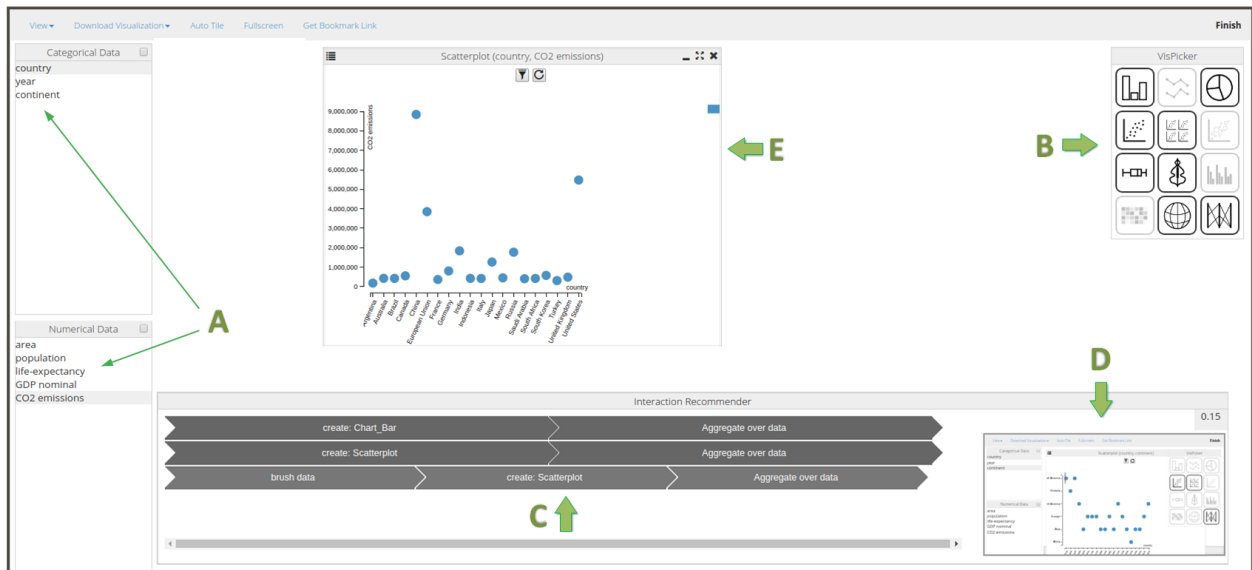
Figure 1. Task-based visual analytics tool: Stage A shows how to choose the attributes to be visualized. Stage B lists the appropriate visualizations that are automatically recommended by our rule-based system for the selected attributes. Stage C lists the recommended sequence of interactions (paths) for user's current task. Stage D provides a preview of a visualization that shows the final outcome of the selected path. Once the user accepts the recommendation, the final visualization is shown in stage E.

3. Approach

We propose a task-based visual analytics tool, which observes current user behavior to infer the visual task she is currently working on (distribution, correlation, finding anomalies, finding input/output relation). Once the task has been recognized, the tool recommends the next interaction (aggregation, filtering, linking/brushing, zooming/panning) to assist the user in completing this task successfully. To do so, the proposed tool relies on two stages: an offline stage and an online stage. The offline stage takes over the task to model the user behavior when performing a particular visual task. To achieve this, we use a repository of user interactions (e.g., filtering, zooming, etc.) previously collected through a user study in which users were asked to perform several tasks. The details of the study can be taken from the Section 4.1. In contrast, during the online stage, the tool observes the current user behavior and attempts to recognize the task she is currently working on. When the task is recognized, the tool recommends to the user a sequence of interactions that should provide her guidance on completing the task and thus on extracting valuable information from the underlying dataset.

In the following, we first detail our approach for collecting, preprocessing and customizing user interactions. Following this we introduce the algorithms used to model the user behavior and the four types of visual

tasks based on collected user data. Finally, we present how we apply these models to generate recommendations and provide guidance to the user on performing a certain task.

3.1. Collecting user interactions

The (twelve) visualizations (bar chart, line chart, pie chart, scatterplot, scatterplot matrix, bubble chart, box plot, violin chart, grouped bar chart, heat map, map, parallel coordinates) in our visual analytics tool provide several interactions (filtering, aggregating, sorting, zooming/panning, linking/brushing) each covering a certain purpose (e.g., filtering, highlighting, unified view). Whenever the user selects a visualization and uses the integrated interactions, the "logger process" is initialized that, in real time, collects and saves the interaction events in a repository. The collected interactions are used to match users' current behavior to a set of interaction patterns derived from previous user behaviors. The following section details the definition of these patterns.

3.2. Defining user behavior patterns

To investigate which interactions the users perform when a certain visual task is given, we conducted a user study and asked the users to work on predefined tasks (see Section 4.1). The tasks were defined with regard to

the four visual task categories (gaining overview, analyzing outliers, analyzing trends and comparing variable distribution) the users usually need to carry out in the analytical process. While the users worked on the tasks, we logged their behaviors and used them to define behavior patterns. Note that we only consider the behavior of the users who have successfully carried out the given task.

We made use of Markov Chain (MC) model for the pattern definition because it has been proven as effective for defining dynamic sequential behavior patterns [26, 27, 28]. Basically, MC is a stochastic model which describes a sequence of possible actions in which each action depends on the previous actions of the user: $P(X_{n+1} = x | X_1 = x_1, X_2 = x_2, \ldots, X_n = x_n) = P(X_{n+1} = x | X_n = x_n)$. This model consists of three components: the state space, the transition matrix, and the initial vector [29]. In our case, the set of all states represents the set of interactions performed within our visual analytics tool. For instance, the interactions "selecting the data attributes", "aggregating", "filtering" etc., are interpreted as states. When combined, they define the state space $S = \{s_1, s_2, s_3, \ldots, s_n\}$, where s_i denotes a particular state with the state index i, $i = 0, 1, 2, \ldots n$. The transition Matrix P, however, describes the probability that a user, whose current state is s_i, will select the next state s'. Finally, the transition probability (p_{ij}) describes the probabilities resulting from various state changes, and is calculated using the Maximum Likelihood Estimation (MLE) [30, 31, 27]:

$$p_{ij} = \frac{n_{ij}}{\sum_j n_{ij}} \tag{1}$$

where $n_{i,j}$ denotes the total number of transitions between the states $s_i \in S$ and $s_j \in S$. For instance, the following transition matrix shows that if a user brushes (i.e., considering different aspect of the same data on different charts) over a visualization, she is more likely to zoom as next.

$$\mathbf{P} = \begin{array}{c} \\ \begin{bmatrix} 0.2 & 0.1 & 0.7 \\ 0.3 & 0.5 & 0.2 \\ 0.6 & 0.4 & 0.0 \end{bmatrix} \end{array} \begin{array}{l} \text{Brush} \\ \text{Filter} \\ \text{Zoom} \end{array}$$

with columns Brush, Filter, Zoom.

MC of various orders are available. The differentiating factor between them is the numer of previous states the current state depends on. For instance, for the first order MC the next step depends on the current state only and not on past states [32]. In contrast, when using the second order MC the next step follows the two past

states etc. Yet, the studies presented in [33, 34, 35] revealed that the first order MC is plausible due to the high number of parameters needed for higher order models. In order to avoid the shortcomings in available data, we follow this example and apply a first order MC to model user behavior within a task. Note that using the first order MC, we trained four behavior models each for one of the four task categories.

3.3. Task recognition

To identify the different tasks considering the m first interactions in users' data, we made use of the random forest algorithm (RF). The main idea behind this algorithm is to apply the "divide and conquer" principle: the data is sampled randomly in small parts (subsets) and then used to grow a decision tree predictor within each small data subset. We use RF, because of its ability to deal even with the small datasets. The decision tree model is learned using the data from the user study and is applied to recognize four different visual tasks in real time applications. To recognize the intended task, our system first observes the current user's interactions on the visual analysis tool. After the user has performed three sequential interactions, the system applies the learned model and attempts to classify the current task. The advantage of using this data-driven approach is that the patterns associated to the tasks are learned automatically (Section 3.2). Consequently, with each new dataset and new user, new patterns are learned which over time would improve the performance of the task recognition algorithm. Once the task can be classified, the system finally applies the four MC models and recommends the user the next analytical step.

The process of task guidance is challenging as it consist of the two aforementioned parts (task recognition and recommending the next analytical step). To tackle this challenge and improve the results we considered a hybrid approach consisting of RF and MC algorithms.

3.4. Visual guidance interface

Our visual analytics tool is primarily a rule-based system which is built upon visual encoding rules and perceptual guidelines [36, 37, 38, 39], and automatically recommends a list of appropriate visualizations for a given dataset. For the recommendations, the user first uploads her dataset and selects the attributes (i.e., specification that can be measured, observed or logged such as age, price, temperature) which should be visualized. After this, the system defines a list with visual mappings, valid for the selected data and presents these mappings to the user as recommendations. A mapping is valid only if there is a data type compatibility between the at-

tributes and the visual components (e.g., axes) of a visualization. The approach of our rule-based recommender system is described in detail in our previous paper and is beyond the scope of this paper. For further details please refer to [8].

Once presented, the user can select any of the recommended visualizations and use it for exploring and analyzing her data. This process is illustrated in Figure 1A and B. When the system now recognizes the task the user is currently working on, it recommends a sequence of interactions that are presented on the UI as an arrow-formed chain (in further text, path) with each arrow representing a particular interaction (see Figure 1C). The direction of the arrows describes the order of the interactions that must next be performed to succeed with the task. However, if there is more than one possible path, a ranking algorithm ranks them based on how likely they will be performed by the user as next and presents them to the user in form of a list. When the user now clicks on a path, she is provided with a preview of a visualization that shows the final outcome of the path, recommended interactions respectively (see Figure 1D). The user can accept the preview by clicking on the recommended visualization (see Figure 1E). However, it is also possible to select a single interaction within a path and consider the visualization that illustrates the outcome of it. To make a short recall, guidance is a dynamic process (i.e., user is included in the loop) that can be decomposed into a series of actions or decisions [15], made not only by the system but also by the user. Actions taken by the user in the guidance process are of enormous importance since these promote user's understanding and also the generation of new knowledge about the unknown dataset. Thus, the interactive exploration of the recommended interactions enables the user to gather new insights about the data and interpret them in the context of the current task. It also paves the way for an in-depth analysis of the various options for performing a particular task and thus fosters the user's perception of the relevance of that option.

4. Evaluation

In this section we investigate the performance of our task-based visual analytics tool in generating recommendations for a specific task. The quality of suggesting tasks while observing user behavior may depend on a number of different factors, including the complexity of a task the user is currently working on, the available set of visualizations and the supported visual operations. In order to prepare measurements for observing this quality, we conducted a user study in which we collected the initial data for training the models for our task-based

visual analytics tool. This section describes in detail (i) the data source and the methods used to perform the user study, and (ii) the evaluation of the task-recognition algorithm.

4.1. User study

For the user study, we used the rule-based recommender, without task-recognition and guidance, that integrates twelve interactive visualizations. Using the visualizations, the participants had to perform four individual tasks each with three sub-tasks. The tasks resulted from a real-live dataset about the *CO$_2$ emissions in G20-Countries within the last 10 years*. We defined one task (with three sub-tasks) for each visual task category (see Section 3.2). For instance, the first task was *"Which countries are represented in this dataset? Can you recognize a pattern?"* This task can be linked to the high-level task definition "gaining overview", because in order to succeed in it, one must first explore the visualizations (e.g., zoom out views), visually navigate through the data, and finally identify the patterns (i.e., only G20-Countries are shown). With regard to Shneiderman [10], these are the steps to follow for *gaining overview* about the data. For the remaining tasks, we refer to the document.[1]

The participants had three minutes for the tasks 1& 2 and five minutes for the tasks 3&4. Once a task was completed or the allotted time was consumed, the participants had to submit their answers using a dialogue provided on the user interface (UI). While a participant worked on a task, every interaction on the UI was logged and saved. The collected data have been used to train the models for the task-based guided analytics. In the training, we only used the data from participants who provided the correct answer before the allotted time expired. Subjective feedback has been collected through a post-task questionnaire consisting of a 10-point likert NASA TLX[2] scale covering six dimensions of workload (mental demand, physical demand, temporal demand, effort, frustration, perceived performance).

4.2. Task recognition quality

In the second part of our evaluation, we performed an offline experiment to estimate the performance of the task-based visual analytics tool in generating recommendations for a specific task. The performance of the recommender highly depends on the accuracy of the trained model for task-recognition. To this end, we used the data from the user study and randomly split it into

---

[1]https://bit.ly/2RIbPwv
[2]https://humansystems.arc.nasa.gov/groups/tlx/

training and validation set performing five-fold cross validation. We trained the classifier with the training set and compared the generated recommendations (prediction set) with the validation set. To report the accuracy of the model, we used the confusion matrix and presented the results of the classification outcomes [40].

4.3. Results and discussion

This section presents the results of the (1) user study conducted to collect user data, and (2) offline evaluation conducted to estimate the performance of the task-based visual analytics tool.

4.3.1. User study: data elicitation The user study was carried out with twenty participants (M=15, F=5), ages 18-57. The users had experiences with visual- and data analytics tools. Task 1 (incl. three sub-tasks) was successfully completed by 5 participants, in 2 minutes (m) and 40 seconds (s) and with 12.65 interactions. Task 2 was successfully completed by 15 participants, in 1m 30s and with 8.64 interactions. In a similar manner task 3 was successfully completed by 20 participants in 1m and 50s and with 9.67 interactions. Finally, task 4 was successfully completed by 18 participants, in 2m and 18s and with 11.19 interactions. Note that all reported samples represent average values. The workload has been estimated using the results of the NASA TLX questionnaire. The results have revealed low overall workload on all the tasks (mean below 50): (T1) $\mu = 26$, $\sigma = 12.37$, (T2) $\mu = 18.6$, $\sigma = 9.72$, (T3) $\mu = 20$, $\sigma = 9.71$, and (T4) $\mu = 26$, $\sigma = 14.41$.

In addition to basic behavior interactions collected in the course of this study, the collected initial results have also brought an important insight: the first task was much more difficult to complete on time compared to the other three tasks. We assume, this was caused by the nature of the task (rather than its complexity), i.e. the user had first to gain an overview of the data. After gaining familiarity with the data, the users became faster (less time and fewer interactions) and more successful (although tasks 3&4 were more cognitively demanding). This is an important outcome for the behavior and task recognition models. It suggests that the recognition of the different tasks may require different handling in terms of e.g., the number or type of interactions per task. The relatively small quantity of the collected data, however, prevented us from investigating this assumption further and establishing appropriate statistical inferences. This study thus needs a follow-up involving more participants.

4.3.2. Task recognition quality To represent the classification outcomes we used the confusion matrix which is commonly used to represent the accuracy results of a classification model [40]. The confusion matrix is illustrated in Figure 2. It can be observed on examining the results that the trained model could correctly classify task 2 and task 3 but not task 1 and task 4 (default probability threshold is set to 0.50). To be more accurate, task 1 has been classified as task 3 whereas task 4 could not be classified at all. The misclassification of task 1 can be understood since both tasks (1&3) were formulated in a very similar manner. To give an example, task 1c was *"How many countries have a GDP greater than 4 000 000 USD?"*, whereas task 3c was *"Is there a trend of increasing/decreasing GDP value in a certain country from 2005-2015? If yes, name these countries?"*. In contrast, the formulation of task 4 was too general compared to task 2 and 3 leading to a high variety of interpretation and interactions by the participants (see Section 4.1). In a nutshell, the sub-tasks of task 4 were *"What is the distribution of life expectancy/GDP/population per country between 2005 and 2015?"*. In summary it can be said that while our task-based approach seems to be promising, we face some challenges in applying it. The major challenge concerns the definition of the tasks we use to train our classifier. An overly general description (i.e., overly high level of description) might mislead the user and cause a variety of unrelated activities (i.e., analytical steps) which would hinder the performance of the classifier. By contrast, if the tasks are too similar to each other the same analytical steps may appear in multiple tasks and bias the classifier. Another challenge concerns the choice of the algorithm used. We did not evaluate the performance of the random forest classifier against a baseline algorithm, such as SVM and logistic regression, to assess which algorithm is more appropriate for our purpose. Hence, this study needs a follow-up involving different datasets, tasks and different algorithms for task recognition to make more accurate statements about the task recognition quality.

5. Conclusion and future work

Our proposed visual analytics tool guides the user throughout a visual task by recommending the next analysis steps in the form of interactions. To do so, our tool logs the current behavior of the user and tries to match it to a set of interaction patterns derived from previous user behaviors. These patterns are used to identify the implicit task of the user and to recommend the next visual interactions, which should help the user to carry out the task. This work, however, is lacking in a thoroughgoing
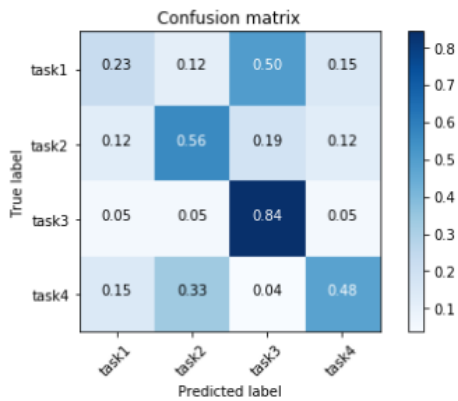
Class classification based on the first three interactions

Figure 2. Confusion matrix showing the accuracy results of the task-recognition model (trained with random forest classifier).

investigation of the usability and usefulness of a guided visual analytics tool. While we have made the hypothesis that a task-based user guidance would positively influence the performance and experience of users, we did not conduct an online user study to find out how people would respond when working with such a tool. We thus plan to conduct a user study in the future to evaluate the value (significance) of our guided analytics tool compared to our regular visual analytics tool which does not assist the user in exploring the data.

It stands to reason that when guiding the user towards a task, considering the different user preferences (e.g., visual taste) or needs (e.g., knowledge gap, background, expertise) is paramount, because this helps us to assess how much guidance is required by the current user. Recent research has provided strong experimental evidence supporting the practical applicability of personalization in visual analytics processes [8, 41]. Although the current version of our guided visual analytics tool does not address personalization- it only uses a global user model, resulting from an empirical user study, and ignores individual differences among users 3.2)- it may substantially enhance the quality and accuracy of the recommendations and thus the impact of the guidance on user performance. The privacy and security risks associated with collecting and processing personal data, however, cannot be overestimated. In order to make use of the power of personalization, we will thus focus in the future on user preferences/needs oriented user guidance, integrating privacy protection techniques to fully protect the user's data while generating high quality personalized recommendations that are adapted to different user needs (e.g., knowledge-gap)

and preferences. This approach will build on the philosophy of considering local information of user interaction with the system and using them to retrain the behavior models. Extending beyond that, we plan to perform a crowd-sourced study with more tasks for each category (*gaining overview, analyzing outliers (anomalies), analyzing trends and comparing variable distribution*), different datasets and longer sessions to collect a substantial number of real user data to train and evaluate the algorithms. Another purpose of this study will be to validate and justify the representativeness of the given tasks to the abstract task categories. To do this, participants will be asked to assign each question to one of the four categories after completing the task, without first stating to which category the task was originally assigned.

In the context of guidance, it is also crucial to define correct timing of guidance so that the analysis flow will not be interrupted [42]. As highlighted by Ceneda et al. [15] considering e.g., the long stall times or absence might not be enough to assess if guidance is needed. Revisiting previously seen states [43] or performing a non-systematic analysis [44], however, may indicate situations in which guidance might be appropriate/needed. Hence, in our future work we will use heuristics to assess the degree and the optimal timing of user guidance [14] and thus avoid distraction and decrease of user engagement in the analysis process.

References

[1] T. Munzner, *Visualization Analysis and Design*. A.K. Peters visualization series, Peters/CRC Press, 2014.

[2] M. S. T. Carpendale, "Considering visual variables as a basis for information visualisation," vol. 2001-693-16, 2003.

[3] L. M. Blaha, D. L. Arendt, and F. Mohd-Zaid, "More bang for your research buck: Toward recommender systems for visual analytics," in *Proceedings of the Fifth Workshop on Beyond Time and Errors: Novel Evaluation Methods for Visualization*, BELIV '14, (New York, NY, USA), pp. 126–133, ACM, 2014.

[4] J. Mackinlay, P. Hanrahan, and C. Stolte, "Show me: Automatic presentation for visual analysis," *IEEE Transactions on Visualization and Computer Graphics*, vol. 13, pp. 1137–1144, Nov 2007.

[5] A. Key, B. Howe, D. Perry, and C. Aragon, "VizDeck: Self-organizing dashboards for visual analytics," *Pro-*

ceedings of the 2012 international conference on Management of Data - SIGMOD '12, p. 681, 2012.

[6] M. Vartak, S. Madden, A. Parameswaran, and N. Polyzotis, "Seedb: Automatically generating query visualizations," *Proc. VLDB Endow.*, vol. 7, pp. 1581–1584, Aug. 2014.

[7] K. Wongsuphasawat, D. Moritz, A. Anand, J. Mackinlay, B. Howe, and J. Heer, "Voyager: Exploratory Analysis via Faceted Browsing of Visualization Recommendations," *IEEE Transactions on Visualization and Computer Graphics*, vol. 22, no. 1, pp. 649–658, 2016.

[8] B. Mutlu, E. Veas, and C. Trattner, "Vizrec: Recommending personalized visualizations," *ACM Transactions on Interactive Intelligent Systems*, vol. 6, pp. 31:1–31:39, Nov. 2016.

[9] D. Gotz and Z. Wen, "Behavior-Driven Visualization Recommendation," *IUI '09 Proceedings of the 14th international conference on Intelligent user interfaces*, pp. 315–324, 2009.

[10] B. Shneiderman, "The eyes have it: a task by data type taxonomy for information visualizations," in *Proceedings 1996 IEEE Symposium on Visual Languages*, pp. 336–343, Sep. 1996.

[11] H.-J. Schulz, T. Nocke, M. Heitzler, and H. Schumann, "A design space of visualization tasks," *IEEE Transactions on Visualization and Computer Graphics*, vol. 19, pp. 2366–2375, Dec. 2013.

[12] M. Brehmer and T. Munzner, "A multi-level typology of abstract visualization tasks," *IEEE Transactions on Visualization and Computer Graphics*, vol. 19, pp. 2376–2385, Dec 2013.

[13] D. Gotz and M. X. Zhou, "Characterizing users' visual analytic activity for insight provenance," in *2008 IEEE Symposium on Visual Analytics Science and Technology*, pp. 123–130, 2008.

[14] C. Collins, N. Andrienko, T. Schreck, J. Yang, J. Choo, U. Engelke, A. Jena, and T. Dwyer, "Guidance in the human-machine analytics process," *Visual Informatics*, vol. 2, no. 3, pp. 166–180, 2018.

[15] D. Ceneda, T. Gschwandtner, T. May, S. Miksch, H. Schulz, M. Streit, and C. Tominski, "Characterizing guidance in visual analytics," *IEEE Transactions on Visualization and Computer Graphics*, vol. 23, pp. 111–120, Jan 2017.

[16] M. Vartak, S. Huang, T. Siddiqui, S. Madden, and A. Parameswaran, "Towards Visualization Recommendation Systems," *1st Workshop on Data Systems for Interactive Analysis (DSIA)*, no. Section 2, p. 1, 2015.

[17] D. Sacha, A. Stoffel, F. Stoffel, B. C. Kwon, G. Ellis, and D. A. Keim, "Knowledge generation model for visual analytics," *IEEE Transactions on Visualization and Computer Graphics*, vol. 20, pp. 1604–1613, Dec 2014.

[18] S. M. Casner, "Task-analytic approach to the automated design of graphic presentations," *ACM Trans. Graph.*, vol. 10, pp. 111–151, Apr. 1991.

[19] F. Bouali, A. Guettala, and G. Venturini, "Vizassist: An interactive user assistant for visual data mining," *Vis. Comput.*, vol. 32, pp. 1447–1463, Nov. 2016.

[20] A. Srinivasan, S. M. Drucker, A. Endert, and J. T. Stasko, "Augmenting visualizations with interactive data facts to facilitate interpretation and communication," *IEEE Trans. Vis. Comput. Graph.*, vol. 25, no. 1, pp. 672–681, 2019.

[21] K. Cook, N. Cramer, D. Israel, M. Wolverton, J. Bruce, R. Burtner, and A. Endert, "Mixed-initiative visual analytics using task-driven recommendations," in *2015 IEEE Conference on Visual Analytics Science and Technology (VAST)*, pp. 9–16, Oct 2015.

[22] N. Silva, T. Schreck, E. Veas, V. Sabol, E. Eggeling, and D. W. Fellner, "Leveraging eye-gaze and time-series features to predict user interests and build a recommendation model for visual analysis," in *Proceedings of the 2018 ACM Symposium on Eye Tracking Research & Applications*, ETRA '18, (New York, NY, USA), pp. 13:1–13:9, ACM, 2018.

[23] L. Shao, N. Silva, E. Eggeling, and T. Schreck, "Visual exploration of large scatter plot matrices by pattern recommendation based on eye tracking," in *Proceedings of the 2017 ACM Workshop on Exploratory Search and Interactive Data Analytics*, ESIDA '17, (New York, NY, USA), pp. 9–16, ACM, 2017.

[24] D. Sacha, I. Boesecke, J. Fuchs, and D. A. Keim, "Analytic behavior and trust building in visual analytics," in *EuroVis 2016 - Short Papers* (E. Bertini, N. Elmqvist, and T. Wischgoll, eds.), pp. 143–147, The Eurographics Association, 2016.

[25] K. Nazemi, R. Retz, J. Bernard, J. Kohlhammer, and D. Fellner, "Adaptive semantic visualization for bibliographic entries," *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, vol. 8034 LNCS, no. PART 2, pp. 13–24, 2013.

[26] A. Pentland and A. Liu, "Modeling and prediction of human behavior," *Neural Comput.*, vol. 11, pp. 229–242, Jan. 1999.

[27] S. Walk, P. Singer, M. Strohmaier, D. Helic, N. F. Noy, and M. A. Musen, "How to apply markov chains for modeling sequential edit patterns in collaborative ontology-engineering projects," *Int. J. Hum.-Comput. Stud.*, vol. 84, pp. 51–66, Dec. 2015.

[28] P. Singer, D. Helic, B. Taraghi, and M. Strohmaier, "Detecting memory and structure in human navigation patterns using markov chain models of varying order," *PLOS ONE*, vol. 9, pp. 1–21, 07 2014.

[29] N. Privault, *Understanding Markov Chains*. Springer Undergraduate Mathematics Series, 2013.

[30] S. Rendle, C. Freudenthaler, and L. Schmidt-Thieme, "Factorizing personalized markov chains for next-basket recommendation," in *Proceedings of the 19th International Conference on World Wide Web*, WWW '10, (New York, NY, USA), pp. 811–820, ACM, 2010.

[31] G. Biau and E. Scornet, "A random forest guided tour," *TEST*, vol. 25, pp. 197–227, Jun 2016.

[32] B. Berent, B. Mobasher, M. Levene, B. M. Masand, and M. Spliliopoulou, "Web navigation path pattern prediction using first order markov model and depth first evaluation," 2016.

[33] J. Borges and M. Levene, "Data mining of user navigation patterns," in *Revised Papers from the International Workshop on Web Usage Analysis and User Profiling*, WEBKDD '99, (Berlin, Heidelberg), pp. 92–111, Springer-Verlag, 2000.

[34] P. L. Pirolli and J. E. Pitkow, "Distributions of surfers' paths through the world wide web: Empirical characterizations," *World Wide Web*, vol. 2, pp. 29–45, Jun 1999.

[35] R. R. Sarukkai, "Link prediction and path analysis using markov chains," *Comput. Netw.*, vol. 33, pp. 377–386, June 2000.

[36] J. Bertin, *Semiology of graphics*. University of Wisconsin Press, 1983.

[37] C. Ware, *Information Visualization: Perception for Design*. San Francisco, CA, USA: Morgan Kaufmann Publishers Inc., 3 ed., 2012.

[38] P. Rheingans and C. Landreth, "Perceptual principles for effective visualizations," in *Perceptual Issues in Visualization* (G. Grinstein and H. Levkowitz, eds.), (Berlin, Heidelberg), pp. 59–73, Springer Berlin Heidelberg, 1995.

[39] J. von Engelhardt, *The Language of Graphics: A Framework for the Analysis of Syntax and Meaning in Maps, Charts and Diagrams*. Yuri Engelhardt, 2002.

[40] R. Susmaga, "Confusion matrix visualization," pp. 107–116, 01 2004.

[41] B. Mutlu, E. E. Veas, and C. Trattner, "Tags, titles or q&as?: Choosing content descriptors for visual recommender systems," in *Proceedings of the 28th ACM Conference on Hypertext and Social Media, HT 2017, Prague, Czech Republic, July 4-7, 2017* (P. Dolog, P. Vojtás, F. Bonchi, and D. Helic, eds.), pp. 265–274, ACM, 2017.

[42] M. Chen and D. Ebert, "An Ontological Framework for Supporting the Design and Evaluation of Visual Analytics Systems," *Computer Graphics Forum*, 2019.

[43] M. Behrisch, F. Korkmaz, L. Shao, and T. Schreck, "Feedback-driven interactive exploration of large multi-dimensional data supported by visual classifier," in *2014 IEEE Conference on Visual Analytics Science and Technology (VAST)*, pp. 43–52, Oct 2014.

[44] E. T. Brown, A. Ottley, H. Zhao, Q. Lin, R. Souvenir, A. Endert, and R. Chang, "Finding waldo: Learning about users from their interactions," *IEEE Transactions on Visualization and Computer Graphics*, vol. 20, no. 12, pp. 1663–1672, 2014.