

Immersive Navigation in Visualization Spaces through Swipe Gestures and Optimal Attribute Selection

Jan-Frederik Kassel*
Volkswagen Data:Lab
University of Hannover

Michael Rohs†
University of Hannover

ABSTRACT

Exploratory data analysis is an essential step in discovering patterns and relationships in data. However, the exploration may start without a clear conception about what attributes to pick or what visualizations to choose in order to develop an understanding of the data. In this work we aim to support the exploration process by automatically choosing attributes according to an information-theoretic measure and by providing a simple means of navigation through the space of visualizations. The system suggests data attributes to be visualized and the visualization's type and appearance. The user intuitively modifies these suggestions by performing swiping gestures on a tablet device. Attribute suggestions are based on the mutual information between multiple random variables (MMI). The results of a preliminary user study ($N = 12$ participants) show the applicability of MMI for guided exploratory data analysis and confirm the system's general usability (SUS score: 74).

Index Terms: H.5.2 [Information Interfaces and Presentation]: User Interfaces; H.1.1 [Information Systems]: Systems and Information Theory

1 INTRODUCTION

In business [6] as well as in private life [7] data analysis is getting more and more important. Driven by wearable sensor technologies and mobile devices, smart objects, networked cars, and online infrastructure (e.g., cloud services) the collection of huge amounts of data in everyday life has gotten easier than ever. Yet, concepts for simple and intuitive visualization and data analysis for nonprofessional and inexperienced users are scarce [15]. Ubiquitous personal data collection requires a paradigm shift in the way we analyze and make sense of data. The scenario in which a person sits at a workstation is likely to give way for scenarios in which people are able to analyze data of interest whenever and wherever they want [20].

Exploratory data analysis (EDA) [24] strongly benefits from suitable visualizations. Visualizations perfectly serve as the first entry point for analyzing an unknown data set [4]. With suitable visualizations the user may easily recognize the distributions of data attributes and relationships within the data. However, effort, care, and knowledge are required to design an effective and useful visualization [9]. Additionally, visualization preferences are subjective [25]. Whether a visualization fits a specific scenario partly depends on the user's preferences. Concepts for simple, intuitive, and immersive visualizations have to consider several factors: Handheld devices as the main entry point, new user groups with little knowledge about data visualization, and an associated risk of frustration during the exploration. Immersive [5] data exploration concepts for mobile devices [15, 20] could stimulate the user's fun during data exploration.

In this work, we present a gesture-based concept for tablet devices that enables immersive exploratory data analysis without the

need for prior knowledge about the data set or the visualization pipeline [3]. The interface offers four simple swipe gestures to refine the visualization placed at the center of the device's screen in terms of data attributes, visualization type, and further visualization parameters (e.g., interpolation). As a mixed-initiative approach, the user controls the visualization's type and minor appearance, whereas the system controls the combination of the data attributes and the visual encoding. We use the information-theoretic measure of *multivariate mutual information (MMI)* [23] as an indicator of the potential *information gain* of an attribute for a particular visualization. This allows the user to playfully navigate through the visualization space – i.e., the set of all valid visualizations for the given data set – and immediately discover relevant data characteristics.

Our contribution is a gesture-based interaction concept for refining visualizations and the use of an information-theoretic measure to let the system optimally choose visualization parameters. This combination of interaction concept and automatic parameter optimization allows the delegation of details to the system.

In the following, we give a brief overview of related work, explain the mobile interaction concept, and present the results of a user study.

2 RELATED WORK

In the research fields of exploratory data analysis [24] as well as in exploratory search [17], researchers investigate how to facilitate the user's knowledge acquisition when exploring previously unknown data sets as well as how to decrease the uncertainty about their objectives.

Wegman et al. [26] analyzed the potential of virtual reality via head-mounted devices for EDA. May et al. [18] used a semi-automated approach for improving the feature subset selection. The user can directly see the dependencies and interact with the selection algorithm, e.g., by changing the ranking. Khurana et al. [13] decomposed the EDA process into independent components to reduce the required time. Tableau¹ automatically generates visualizations but does not help the user in the selection of suitable data attributes. Wongsuphasawat et al. [27] designed Voyager, a system for recommending visualizations for both the attributes the user selects as well as for an extended attribute set by automatically adding a non-selected attribute. Xia et al. [28] introduced a tree structure of visualizations (limited to 1D and 2D) for data structure analysis. As a measure, they use the mutual information between two samples. Compared to our concept this approach is mainly designed for ordinary desktop systems, in which space is not an issue and interactions are based on mouse and keyboard.

Our interaction approach is inspired by Nandi et al. [19], who investigated the use of gestures sequences for querying databases. Beltran et al. [1] used a gestural interface for recommending documents. In their two-stage approach, users swipe right or left to decide if they like the shown document. During the gesture execution so-called reason bins appear and the user has to swipe through one of the reason bins to indicate why they like or dislike the document. In our approach, we do not use a second decision layer, but instead we added another swipe direction for refinement.

*e-mail:jan-frederik.kassel@volkswagen.de

†e-mail:michael.rohs@hci.uni-hannover.de

¹<https://www.tableau.com>




Figure 1: The four offered gestures for navigating through the visualization space. In this example, swiping right (a) switches the axis. Swiping left (b) transforms the scatter plot into a point plot. Swiping up (c) computes the MMI, selects an attribute and maps it onto the visual variable “size”. Swiping down (d) computes the MMI and removes a data attribute.

3 SCENARIOS

We envision two potential usage scenarios of our concept. In the first scenario, a user would like to analyze data related to a news story while sitting on the couch. Conventionally, the user would have to analyze the attributes’ data types, identify potential combinations, and choose a suitable visualization, which involves a risk of weak quality and user frustration [9]. In our concept, the user just has to upload the data table. Thanks to the playful interaction flow, she can now focus on the data analysis. This scenario is the basis for the evaluation reported below.

The second scenario is a meeting situation. In this synchronous and collocated activity the communication medium is an interactive surface [11] on which the proposed user interface is running. The meeting participants can directly focus on and discuss about the data without being distracted by details of the visualizing process.

4 DESIGN

The primary objective of our concept is to facilitate immersive and personalized EDA of new data sets without the knowledge of how to create proper visualizations. Since immersion strongly depends on engagement, lowering the engagement barriers is a key factor [2]. Due to the subjective nature of visualization, we decided to give the user partial control over the visualizations. To fulfill this objective, we found design inspirations in the dating App Tinder² and in the research concept BINGO [1]. Tinder allows the user to decide whether or not he or she likes to date a recommended person by simply swiping left (dislike) or right (like). Visual exploration of a data set here is analogous to finding a date. In both cases users have to navigate through a large but finite search space to find those elements that fit their preferences.

Our proposed concept’s interaction flow is designed according to Shneiderman’s visual information seeking mantra: overview first, zoom and filter, then details-on-demand [22]. In our concept, the user starts with picking the first data attribute which he or she would like to visualize, as Wongsuphasawat et al. [27] proposed. After the first visualization is shown, the user has to decide whether the visualization fits the requirements. The user can adjust a visualization by using vertical swipes for changing the information (or level of detail) of a visualization (including/excluding data attributes) or using horizontal swipes for (slightly) changing the appearance. We aimed at achieving a clear and playful interaction design to foster

potential immersion [2]. In the following, we describe how the interaction works in detail.

4.1 Interaction

In principle, the FLUID [14] interface offers four different swipe gestures (up, down, left, and right) mapped to four different kinds of refinement with an additional undo button for covering errors and a horizontal space including buttons for each data attribute of the data set (see Figure 1). All swipe gestures are starting in the screen’s center. Unlike BINGO [1], we decided to collect just the high level decisions rather than asking the user to enter the reason for changing the visualization. Visual exploration by nature is vague and foggy. The user often only has a faint idea of what a visualization should look like.

When swiping right, the visualization changes slightly, e.g., the orientation changes or a line chart gets interpolated. When swiping left, the visualization type changes, e.g., a bar chart may be transformed into a line chart (if appropriate). Swiping up adds information (increasing the level of detail) to the visualization, whereas swiping down removes the information (decreasing the level of detail). On each side there is a label that describes the action of the corresponding swipe.

4.2 Adding and Reducing Information

The swipe-up gesture is used to add information to the current visualization. Information is added to the visualization by including an unused attribute. From the perspective of a user with little or no knowledge about the data set it is hard to decide which additional data attribute helps to get a deeper understanding of the currently visualized data. In addition, users are likely not aware how they should visualize the additional attribute in the best way, but this is crucial. The decision on the use of visual variables is one of the most important steps of the data visualization process [3]. The proposed concept uses a two-step procedure to handle the decision of which attribute to add and how include it in the visualization.

In the first step, the algorithm decides which attribute of the data set should be added to the current visualization. One potential strategy is choosing the attribute which offers the highest information gain to the user. Several measures (e.g., variance) could be thought of to augment this information. A useful concept requires a measure which is applicable to every data type (e.g., variance only works on quantitative data) as well as a way to evaluate the impact of a potential attribute. Finally, since important conclusions are based on

²<https://www.gotinder.com>

relations between variables, the user’s reasoning will be supported best by combining those variables that have a strong relationship.

Based on these requirements we choose *multivariate mutual information (MMI)* as the measure for picking data attributes. We define the *information gain* of an attribute X of a data set \mathcal{D} by its future MMI [23]: $\text{MMI}(X; \mathcal{Y})$ where $\mathcal{Y} = \{Y_1, \dots, Y_n\}$ is the set of currently visualized data attributes. The MMI is based on the entropy H , a measure of the uncertainty of a data attribute. The MMI is the generalization of mutual information, which describes for two random variables S and T the reduction of uncertainty of S given the knowledge of T and vice versa. For each attribute in $\mathcal{D} \setminus \mathcal{Y}$ (all attributes that are not yet included in the visualization) the MMI will be computed and the one which offers the highest information gain (maximizes $\|\text{MMI}\|$) will be chosen. In this way the system adds the attribute with the highest possible information gain to the visualization in each interaction step. The user can observe the impact of variables prone to be strongly related to each other and immediately derive conclusions.

$$H(X) = - \sum_{x \in X} p_X(x) \log(p_X(x)) \quad (1)$$

$$\text{MMI}(X; \mathcal{Y}) = \sum_{k=1}^{|\mathcal{X} \cup \mathcal{Y}|} (-1)^{k-1} \sum_{Z \subset \{X, Y_1, \dots, Y_n\}, |Z|=k} H(Z) \quad (2)$$

The second step is based on the order of the visual variables regarding the effectiveness for each level of measurement (nominal, ordinal, and quantitative) [16]. For the chosen attribute we know its level of measurement, which also indicates the most effective available visual variable, e.g. if a categorical attribute has the maximum MMI, it will be encoded by the color instead of the size.

The swipe-down gesture decreases the information of the visualization. In contrast to swipe up it picks for removal the attribute with the lowest information with respect to the current visualization. When n is the number of used data attributes in the current visualization, the MMI is computed for every subset of size $n - 1$. After removing an attribute, if the mapping [16] is not optimal, a remapping of the data attributes on the visual variables is performed.

4.3 Implementation

We implemented our design concept as a Web app using the JavaScript library React. For the visualization, we used the high-level visualization grammar Vega-lite [21]. Vega-lite uses a JSON-like structure to formally describe the layout of a visualization. Because Vega-Lite supports only basic visualization types (e.g. bar chart, line chart, and the visual encodings color, opacity, size/volume, and shape), the resulting explorable visualization space is limited. Nevertheless, it is a very elegant and simple methodology for systematically generating visualizations.

5 EVALUATION

We decided to focus on a qualitative evaluation on the usability and usefulness of our concept. We recruited 12 participants within an industry company. The participants’ background was either in data science (6 participants) or business administration (6 participants), but none of them were educated in data visualization. This mix of participants allowed us to analyze potential differences in behavior. For creating visualizations the majority (10 participants) uses Microsoft Excel. Each participant got a short introduction to the user interface. As test data we chose a high-dimensional data set of student performance in mathematics in Portugal [8]. It consists of 33 dimensions of all levels of measurement (nominal, ordinal, and quantitative) and contains 395 data points. We provided the participants with a brief written description of the data set’s attributes.

We conducted a think-aloud user study where we tracked the observations and statements. Regarding the estimation of the usefulness of our concept, every participant was asked to explore the

data set without any time limit (open exploration). The participants were told to stop at the point where they were certain to have a solid comprehension of the data. We asked the participants to tell us valid statements about the data they discovered as they explored the data. Afterwards we assessed these statements in terms of their complexity and checked for correctness. We defined the complexity of a statement as proportional to the number of data attributes included in the statement.

6 RESULTS AND DISCUSSION

The proposed interface received a SUS score of 74 on average (see Figure 2(b)), which indicates good usability. Furthermore, our concept was well received by the participants. They liked the easy way to check the data attributes and to get some initial ideas on additional attributes. One participant said: *“I really like the four easy gestures for interaction.”* Nevertheless, there were some issues regarding the concept.




Figure 2: The three box plots show the participants’ (a) self-assessment in comprehending the data set, (b) SUS score rating, and (c) duration of using the system.

In the beginning, the most participants had difficulties to use the swipe gestures correctly. Surprisingly, they did the swipes in the opposite direction, e.g., instead of swiping up for more information, they swiped down (the gesture started at the block labeled “show me more”). Another misconception was that some participants attempted a left swipe for undoing a previous right swipe, instead of using the undo button. After a short time the participants used the interactions correctly. However, the errors point at a design issue and the underlying reason have to be analyzed.

All participants started observing the different data attributes one by one, but some of them stopped after four attributes were included within a visualization. They referred to the difficulty of interpreting high-dimensional visualizations. However, 89.2 % of the participants’ statements were correct during the exploration. The most complex statements of the participants included three data attributes. At the end, the participants rated their subjective estimate of comprehension of the data used in the experiment (see Figure 2a) on a scale from 1 (no knowledge) to 10 (expert).

We observed a clear difference in the statement formulation depending on the participant’s data analysis experience. Participants with a background in data science formulated not just statements regarding the shown data, but already discussed if the statement makes sense, e.g., *“there is a strong linear correlation between G1 and G2, as I expected. But this is not good, because it shows us that the students did not improve.”* In addition, these participants were more interested in the approach in general. As Figure 2(c) illustrates, they used the system significantly longer ($p = 0.023$, Kruskal-Wallis H-test) which indicates a higher degree of immersion [12].

Nearly each participant was curious about the way the system adds data attributes to the visualization. After a short explanation of the basic idea the participants found the approach reasonable. Increasing the transparency on the used measure or giving the user the opportunity to directly choose a particular measure would further

mitigate this issue. One data scientist argued that the top three attributes should be highlighted and the choice should stay with the user. Especially data scientists asked if they can combine attributes by themselves. To cover this essential remark, potential shortcuts for experts are needed. This can be done by using more multitouch gestures, e.g., two fingers swipe up for highlighting the top three data attributes for potential inclusion. In terms of visualizing the data, no participant found the suggested visualizations inappropriate. Instead, they were satisfied with the offered mappings and did not intend to change them.

7 LIMITATIONS AND FUTURE WORK

Our preliminary work is limited in the way the information (choice of attributes) is added to the visualization. As discussed above, we will evaluate additional measures for supporting various analysis purposes, e.g., to find independent attributes. Machine learning approaches for predicting interesting and helpful data attribute combinations could be an alternative, e.g., the feature importance of a random forest. The user study has shown a mismatch between the physical action (swipes) and the movement through the visualization space, in particular regarding swipes in opposite directions. In order to prevent the user from disorientation, we are going to add a small overview representing the user's position within the visualization space. This could also allow marking and jumping to specific visualizations.

In a next step, we will evaluate the potential of using interaction sequences to compute recommendations for visualizations. Heer et al. have shown that predicting interactions can help in case of data transformation tasks [10]. By logging the user's interactions, we know for every visualization how it was changed, and the dead ends at which the user stopped. Those explored sequences are very promising for computing recommendations under the assumption that users aim to navigate to visualizations that fit their preferences. In order to achieve a user-independent approach, we will investigate a potential generalization based only on the data characteristics.

8 CONCLUSION

In this paper, we proposed a novel concept for immersive exploratory data analysis on a tablet device that enables user navigation through a visualization space. The concept offers four simple swipe gestures for manipulating a visualization that is located in the display's center. We introduced the information-theoretic measure of the multivariate mutual information as an opportunity to adjust the information of a visualization. The conducted preliminary qualitative user study showed the concept's usability (74 SUS score) and usefulness (89.2% correct statements). In the present study we did not compare the concept against a baseline, which is planned as future work. The gesture-based navigation concept has been judged positively by the participants. Finally, we strongly believe in this promising concept as a first but necessary step to future preference-oriented visualization recommender systems.

ACKNOWLEDGMENTS

We thank Justin Bayer, Marian Harbach, Daniel Weimer, Tim Düntje and all participants for their input. The results, opinions, and conclusions are not necessarily those of the Volkswagen AG.

REFERENCES

- [1] J. F. Beltran, Z. Huang, A. Abouzied, and A. Nandi. Don't just swipe left, tell me why: Enhancing gesture-based feedback with reason bins. In *Proc. IUI '17*, pp. 469–480. ACM, 2017.
- [2] E. Brown and P. Cairns. A grounded investigation of game immersion. In *CHI EA '04*, pp. 1297–1300. ACM, 2004.
- [3] S. K. Card and J. D. Mackinlay. The structure of the information visualization design space. In *Proc. InfoVis '97*, pp. 92–. IEEE Computer Society, 1997.

- [4] J. M. Chambers, W. S. Cleveland, B. Kleiner, and P. A. Tukey. Graphical methods for data analysis. *The Wadsworth Statistics/Probability Series*. Boston, MA: Duxury, 1983.
- [5] T. Chandler, M. Cordeil, T. Czauderna, T. Dwyer, J. Glowacki, C. Goncu, M. Klapperstueck, K. Klein, K. Marriott, F. Schreiber, and E. Wilson. Immersive analytics. In *2015 Big Data Visual Analytics (BDVA)*, pp. 1–8, Sept 2015.
- [6] H. Chen, R. H. L. Chiang, and V. C. Storey. Business intelligence and analytics: From big data to big impact. *MIS Q.*, 36(4):1165–1188, Dec. 2012.
- [7] E. K. Choe, N. B. Lee, B. Lee, W. Pratt, and J. A. Kientz. Understanding quantified-selfers' practices in collecting and exploring personal data. In *Proc. CHI '14*, pp. 1143–1152. ACM, 2014.
- [8] P. Cortez and A. M. G. Silva. Using data mining to predict secondary school student performance. 2008.
- [9] L. Grammel, M. Tory, and M.-A. Storey. How information visualization novices construct visualizations. *IEEE Trans. on Vis. and Comput. Graph.*, 16(6):943–952, Nov 2010.
- [10] J. Heer, J. M. Hellerstein, and S. Kandel. Predictive interaction for data transformation. In *CIDR 2015, Asilomar, CA, USA, January 4-7, 2015*.
- [11] P. Isenberg, T. Isenberg, T. Hesselmann, B. Lee, U. von Zadow, and A. Tang. Data visualization on interactive surfaces: A research agenda. *IEEE Comput. Graph. Appl.*, 33(2):16–24, Mar. 2013.
- [12] C. Jennett, A. L. Cox, P. Cairns, S. Dhoparee, A. Epps, T. Tijs, and A. Walton. Measuring and defining the experience of immersion in games. *Int. J. Hum.-Comput. Stud.*, 66(9):641–661, Sept. 2008.
- [13] U. Khurana, S. Parthasarathy, and D. S. Turaga. Read: Rapid data exploration, analysis and discovery. In S. Amer-Yahia, V. Christophides, A. Kementsietsidis, M. N. Garofalakis, S. Idreos, and V. Leroy, eds., *EDBT*, pp. 612–615. OpenProceedings.org, 2014.
- [14] B. Lee, P. Isenberg, N. H. Riche, and S. Carpendale. Beyond mouse and keyboard: Expanding design considerations for information visualization interactions. *IEEE Trans. on Vis. and Comput. Graph.*, 18(12):2689–2698, Dec 2012.
- [15] A. Lu, J. Huang, S. Zhang, C. Wang, and W. Wang. Towards mobile immersive analysis: A study of applications. In *2016 Workshop on Immersive Analytics (IA)*, pp. 25–30, March 2016.
- [16] J. Mackinlay. Automating the design of graphical presentations of relational information. *ACM Trans. Graph.*, 5(2):110–141, Apr. 1986.
- [17] G. Marchionini. Exploratory search: From finding to understanding. *Commun. ACM*, 49(4):41–46, Apr. 2006.
- [18] T. May, A. Bannach, J. Davey, T. Ruppert, and J. Kohlhammer. Guiding feature subset selection with an interactive visualization. In *IEEE VAST '11*, pp. 111–120, Oct 2011.
- [19] A. Nandi, L. Jiang, and M. Mandel. Gestural query specification. *Proc. VLDB Endow.*, 7(4):289–300, Dec. 2013.
- [20] J. C. Roberts, P. D. Ritsos, S. K. Badam, D. Brodbeck, J. Kennedy, and N. Elmqvist. Visualization beyond the desktop—the next big thing. *IEEE Comput. Graph. Appl.*, 34(6):26–34, Nov 2014.
- [21] A. Satyanarayan, D. Moritz, K. Wongsuphasawat, and J. Heer. Vegalite: A grammar of interactive graphics. *IEEE Trans. on Vis. and Comput. Graph.*, 23(1):341–350, Jan 2017.
- [22] B. Shneiderman. The eyes have it: A task by data type taxonomy for information visualizations. In *Proc. VL '96*, pp. 336–. IEEE Computer Society, 1996.
- [23] S. Srinivasa. A review on multivariate mutual information. *Univ. of Notre Dame, Notre Dame, Indiana*, 2:1–6, 2005.
- [24] J. W. Tukey. Exploratory data analysis. 1977.
- [25] J. J. van Wijk. The value of visualization. In *IEEE Visualization*, pp. 79–86, Oct 2005.
- [26] E. J. Wegman, Q. Luo, and J. X. Chen. Immersive methods for exploratory analysis. *Computing Science and Statistics*, pp. 206–214, 1998.
- [27] K. Wongsuphasawat, D. Moritz, A. Anand, J. Mackinlay, B. Howe, and J. Heer. Voyager: Exploratory analysis via faceted browsing of visualization recommendations. *IEEE Trans. on Vis. and Comput. Graph.*, 22(1):649–658, Jan 2016.
- [28] J. Xia, W. Chen, Y. Hou, W. Hu, X. Huang, and D. S. Ebertk. Dimscanner: A relation-based visual exploration approach towards data dimension inspection. In *IEEE VAST '16*, pp. 81–90, Oct 2016.