Erschienen in: EuroVis Workshop on Visual Analytics (EuroVA) / Turkay, Cagatay; Vrotsou, Katerina (Hrsg.). - Goslar : Eurographics Association, 2020. - S. 61-65. - ISBN 978-3-03868-116-8
https://dx.doi.org/10.2312/eurova.20201088

Learning and Teaching in Co-Adaptive Guidance for Mixed-Initiative Visual Analytics

F. Sperrle[1], A. Jeitler[1], J. Bernard[2] D. Keim[1], and M. El-Assady[1]

[1]University of Konstanz
[2]University of British Columbia, Canada

Figure 1: *In the co-adaptive guidance process, both the system and the user initiate guidance with the goal of* ***learning*** *(adapting their own models data, task and system/user) or* ***teaching*** *(adapting the models of the other), in order to improve the shared analysis process.*

Abstract

Guidance processes in visual analytics applications often lack adaptivity. In this position paper, we contribute the concept of co-adaptive guidance, building on the principles of initiation and adaptation. We argue that both the user and the system adapt their data-, task- *and* user/system-models *over time. Based on these principles, we propose reasoning about the guidance design space through introducing the concepts of* learning *and* teaching *that complement the existing dimension of implicit and explicit guidance, thus, deriving the four guidance dynamics* user-teaching, system-teaching, user-learning, *and* system-learning. *Finally, we classify current guidance approaches according to the dynamics, demonstrating their applicability to co-adaptive guidance.*

1. Introduction

Guiding users in their analysis process is an essential part of visual analytics (VA) systems. Many VA systems provide such guidance in the form of assistance that helps users to overcome knowledge gaps. The guiding elements are often fixed parts of the user interface and typically shown during the entire analysis session. In recent years, guidance has been re-defined to mean an active, mixed-initiative process [CGM*17, CGM19] that provides "*just-in-time*" facilitation [CAS*18]. This definition as an active process means that guidance should be provided as a reaction to previous user actions.

Additionally, Collins et al. specify that guidance should be contextualized and able to adapt to different scenarios dynamically [CAS*18]. While the current definition of guidance captures the mixed-initiative nature of the process, it does not shed light on how users and systems adapt over time. To that end, we propose the concept of *co-adaptive guidance*, which is characterized by *initiation* of guidance and *adaptation* of (mental) models. First, it highlights how users and systems converge towards a common understanding of the task and a shared analysis process to reach their goals. Second, it characterizes typical guidance interactions as *learning* or *teaching*, structuring the guidance process. This paper contributes the concept of co-adaptive guidance that takes into account which actor *initiates* guidance, and which (mental) models change and *adapt* over time. The co-adaptive guidance provides a new perspective on the design space for mixed-initiative guidance, as it cross-cuts the axis of *initiation* and *adaptation* with the axis of ***learning*** and ***teaching***.

2. Background and Related Work

Both *adaptive and mixed-initiative systems* have long been studied in human-computer interaction [Opp94, Hor99]. Early approaches describe the generation of "*knowledge bases*" for controlling adaptive dialog-based systems [Tiß93] and state that systems should model the user, the task, the domain, and themselves [KT94]. More recently, *guidance* has been identified as a promising attempt to enable a better collaboration of the human and the computer [CGM19]. For decision support systems, Morana et al. provide design features for guidance [MSSM17]. In visual analytics, Ceneda et al. define guidance in terms of a knowledge gap, available in- and outputs, and the degree of guidance [CGM*17] as an extension to van Wijk's visualization model [vW06]. Collins et al. criticize the model as "*too abstract to use practically*." [CAS*18] and propose a more process-oriented model based on high-level VA tasks [ALA*18]. They identify "*just-in-time facilitation*" as an important goal of guidance" and state that the knowledge of an "*intelligent guide*" can be categorized as *prior knowledge*, *session-specific knowledge*, and

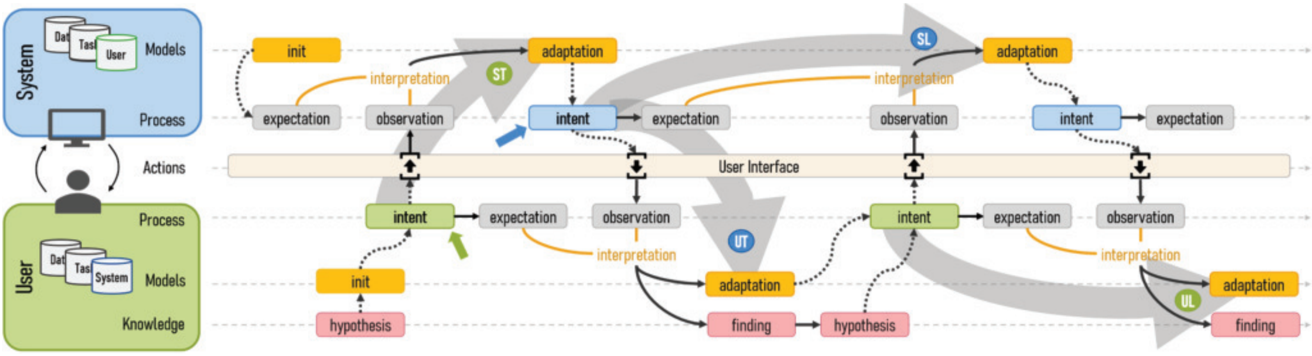

Figure 2: *Action-reaction pairs form the foundation of the co-adaptive guidance process. Reactions are observed and compared to an expectation, leading to the adaptation of the data, task or user/system models, and the derivation of new intents. Here, the user initiates the process (green arrow), and both user and system adapt. The system can also initiate the process, which would then start at the blue arrow. The grey arrows indicate the guidance dynamics system teaching (ST), user teaching (UT), system learning (SL), and user learning (UL).*

situation knowledge [CAS*18]. Federico et al. present a framework that incorporates "*the function and role of tacit and explicit knowledge in the analytical reasoning process*." [FWR*17] More recently, Ceneda et al. explicitly state that guidance is a mixed-initiative process that includes user- and system-guidance [CGM19]. Here, it is interesting to consider who initiated the guidance and who is adapting as a result. In human-machine collaboration, such adaptation process have been studied [Saw05, GBDL15] and modelled game-theoretically [NNPS17]. In this paper, we provide an alternative view on guidance by considering learning and teaching processes. These processes are linked to the provision of explanations and should follow principles from pedagogy, such as clarity, elicitation of responses from learners, and relevance to the learner [Odo14].

Recently, Ceneda et al. (re)defined guidance as an active process [CGM19], rendering several (established) approaches "*not guidance*" and effectively questioning their ability to support users in the analysis process. In this paper, we call approaches matching this new definition ***active*** *guidance* to avoid confusion with earlier systems that employ *guidance*. Further, the terms "system guidance" and "user guidance" are ambiguous. They could each describe both directions of guidance, which can (and has already) led to mix-ups. We thus propose to disambiguate the terms through including the *target* of the guidance in the name: as we further argue in this paper, the success of guidance can be measured through the *adaptation* it induces in the actor being guided. Hence, the *source* of the guidance can be considered interchangeable when comparing different guidance schemes. In this paper, we thus define *system guidance* as support to the system (by the user or another actor), and *user guidance* as support to the user (by the system or another actor).

3. Co-Adaptive Guidance Process

While the target of visual analytics is the generation of hypotheses and the extraction of knowledge [SSS*14], the goal of guidance is to support the analysis tasks at hand. Consequently, active guidance needs to adapt over time, taking the progression of the analysis into account. To illustrate this, we provide a detailed process of this adaptation in the *co-adaptive guidance process*, depicted in Figure 2. The interaction process consists of action-reaction pairs → that are exchanged between the user and the system. Building on the model by Gotz et al. [GZ09], we define *actions* as aggregations of individual events. Actions are, in turn, aggregated into higher-level user or system *intents*. Each intent is associated with an *expectation* that captures if, and how much, the (mental) models of the recipient should adapt as a result of the performed action(s). Users and systems *interpret observations* and expectations with respect to the available models and, in the case of the user, knowledge. The result of this interpretation may lead to an *adaptation* of the recipient, as well as the generation of new *findings*.

As expectations are derived from intent, correctly identifying said intent is an important first step of guidance. It is particularly important that the provided guidance targets adaptation that matches the recipients expected model (data, task, user/system). A mismatch between model types might lead to the guidance being interpreted as less successful or lead to an undesired adaptation. Figure 2 shows an interaction in which all opportunities for interpretation and adaptation have been realized. In practice, many actions will not be interpreted, e.g., because most current systems lack support for intent identification, and users might choose to focus on their task at hand rather than analyzing every system action. Figure 2 also shows grey arrows that indicate the four co-adaptive guidance dynamics that will be introduced in detail in section 4. In the remainder of this section, we introduce the concepts *initiation* and *adaptation* that characterize said dynamics.

3.1. Initiation of Guidance at Different Degrees

Mixed-initiative analysis processes are characterized by a range of interaction possibilities, from suggestions by the different actors, to direct analysis actions. The suggestion operations follow the guidance degrees *orienting*, *directing* and *prescribing* introduced by Ceneda et al. [CGM*17]. In this paper, we propose to situate these guidance degrees on a spectrum that encompasses other unguided analysis operations, going beyond providing

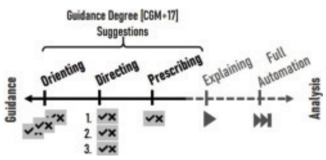

suggestions. We argue that actors can also demonstrate their analysis through explaining and teaching their process and rationale to their counterpart. This explainable interaction is a step towards understanding problems to fully specify and further automate them.

We define the *initiation* of guidance as the action that starts a co-adaptive guidance iteration at a given degree. Initiation can be triggered by both users and systems (see green and blue arrows in Figure 2). Ceneda et al. identified the degrees *orienting*, *directing* and *prescribing* for system-initiated guidance [CGM*17]. Making these degrees available to users as well enables applications to exploit the full potential of mixed-initiative guidance. Most existing systems rely on prescribing user guidance, while some utilize the increased information content of directing guidance [WDC*18]. Eye-tracking devices, among other techniques, would provide the information necessary for orienting guidance. In current applications, the associated guidance processes, however, are typically initiated by the system, and not the user.

User-initiated guidance remains an under-explored field that holds potential for investigating efficient human-machine collaboration. As users convey more information about their tasks, preferences and needs, systems should become better at providing suitable, bespoke reactions. Consequently, they could transition from offering orienting guidance to *dictating* (partial) analysis results, raising questions of how to trade off agency versus automation [Hee19]. To summarize, considering guidance degrees for user-initiated guidance opens up an interesting design space that goes beyond considering *feedback* and *feedforward* [CGM19] for steering the guidance process.

3.2. Adaptation of Knowledge Representation Models

Krogsæter and Thomas state that knowledge-based systems require a model of the user, of the task, of the domain and of themselves (system model) [KT94]. According to their definition, the system model contains knowledge that the system has about its functionality and limitations. As this information is unlikely to change during the guidance process, it is not considered here. Instead, Figure 1 shows that users maintain such a system model. Additionally, users also have a task and domain model. We define *adaptation* as the summary of changes to those models during the guidance process.

While all four models store different information, which will be elaborated in more detail below, the respective adaptation processes are the same and thus combined into one in Figure 2. Taking their current knowledge and the derived *expectation* into account, agents *interpret* and *observe* reactions. The result of the interpretation can then be used as a basis to adapt one or multiple models. For example, users might become more aware of unexplored regions of the data or additional system functionality that could be beneficial to solving the current task. Systems may capture the task users are trying to solve more accurately. Additionally, the interpretation of expected and observed reactions is precisely what fuels the human knowledge generation process, potentially turning hypotheses into findings. As described by Andrienko et al., a "*model*" is an appropriate representation of a subject under study [ALA*18]. We describe the four models based on this definition of the term.

Data Model — The data model contains information such as data distributions, descriptive statistics, identified outliers, and relations and similarities between data points. Typically, systems are expected to have a more complete data model due to their increased computational abilities.

User Model — The system stores a specific user model for each user. This model contains all knowledge that the system has explicitly or implicitly gathered about the user. The user model aims to capture, among others, the users' knowledge, their level of expertise, potential biases, personal preferences, and personality traits. Beyond knowledge, user models should also consider cognitive abilities, such as perceptual speed, visual working memory, and verbal working memory, as personalization can counteract these inter-user differences in performance [CCTL15].

System Model — The system model is the mental model of the system that users create during the analysis. It includes knowledge about the implemented algorithms with their strengths and weaknesses, available visualizations, and guidance operations that the system offers. The system model is created over time through interaction with the system, but also influenced by previous knowledge of similar systems. The system model, therefore, fundamentally influences the expectations the user has about each task outcome.

Task Model — The task model contains all knowledge that is necessary to solve the tasks along the analysis process: an order to execute tasks in, the (hypothesized) solutions, relations and similarities between tasks, and the analysis context.

4. Learn or Teach: Co-Adaptive Guidance Design Space

Initiation and adaptation introduced in the previous section form the foundation of two central concepts in the co-adaptive guidance process: *learning* and *teaching*. In this context, we define the actors' learning intent as the aim to adapt themselves, with the help of knowledge provided by other actors. Conversely, we define teaching intent as the aim to induce adaptation in the other actor. *Learning guidance* and *teaching guidance* are initiated with learning intent or teaching intent, respectively. Learning and teaching are related to implicit and explicit guidance input [CAS*18]: intuitively implicit input can lead to learning, and explicit input could be considered as teaching. However, in this paper, we place the focus on which actor initiated or requested guidance, and with what intent. This is especially interesting as approaches utilizing implicit input for teaching guidance exist [ITB17]. It is important to note that, after requesting learning guidance, neither actor learns in isolation. Instead, the feedback from the other party is fundamental in resolving the encountered knowledge gap. Consequently, *system learning* in guidance is different from general *machine learning*. The combination of initiation and adaptation results in four different guidance dynamics that provide a process-oriented view on guidance in visual analytics: *user teaching*, *system teaching*, *system learning*, and *user learning*. During the analysis process, these dynamics often do not appear in isolation but can be interleaved, as Figure 2 illustrates. Ultimately, systems should aim to enable multiple, if not all, dynamics if they are to be mixed-initiative systems. In the following sections, we introduce each of the dynamics in more detail, followed by real-world examples that represent these principles particularly well.

4.1. Guidance with Teaching Intent

Teaching guidance is initiated by an actor that aims to adapt the models of the other actor. Goals for teaching include providing help in a given situation in order to facilitate the analysis, informing about alternative analysis options, suggesting potential corrections, explaining the current model, or providing a tour as guided exploration. Typically, systems provide teaching guidance targeting the data and task models of users. In contrast, users typically teach systems about the task and their subjective preferences.

User Teaching — *User Teaching* is the most common form of guidance in modern systems, where systems aim to teach users. It directly translates to the original goal of guidance, which is resolving encountered knowledge gaps. To that end, systems, e.g., highlight data points to consider [SGL09] or present alternative analysis pathways [KPHH11].

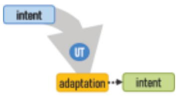

Shao et al. [SSES17] support users in exploring large scatter plot matrices with teaching guidance: based on eye-gaze data, the system shows plots that are visually dissimilar from those already explored. This guidance aims to teach users an unbiased data model that considers all data regions and maximizes the amount of information analyzed per time interval. A similar approach has been used for gaze-based pattern recommendation [SSV*18]. LightGuider is a VA application for creating lighting designs [WSL*19]. Here, teaching guidance supports users in efficient exploration of the large model parameter space, enabling faster task completion by providing alternative model parametrizations, while still supporting "manual intervention and artistic freedom." [WSL*19]

System Teaching — In *system teaching* guidance, the user aims to teach the system their understanding of the task or data. As such, it is closely related to the concept of machine teaching [SAC*17]. However, while machine teaching is typically concerned with providing systems with "*labels, features [or] structure*" [SAC*17], system teaching in guidance also allows systems to update their user model with, e.g., observed preferences and biases.

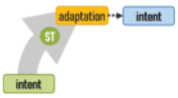

In current applications, system teaching is typically realized via explicit, prescribing guidance: users adapt target sliders [WSL*19] or introduce entity relations [EFN12]. Podium, a system for ranking multi-variate data points, includes directing guidance from the user [WDC*18]. Users teach the system their understanding of the data by reordering the rows of a data table. From this guidance, the system infers a feature weighting model that captures "*which attributes contribute to a user's subjective preference for data.*" [WDC*18] As this model is transparently made available to users, they can compare expectation and observation to make changes.

4.2. Guidance with Learning Intent

Actors request *learning guidance* with the intent of verifying or adapting their own models. Beyond asking for help with the analysis, the goals of learning guidance include probing the other actor's models, verifying hypotheses and understanding the current situation.

User Learning — Users initiate *user learning* with the goal of learning about the data, the system, or its understanding of tasks. This operation can be considered a "*probe*", providing users with additional knowledge without necessarily advancing the analysis.Clustrophile 2, a system for interactive cluster analysis [CD18], offers various algorithms and settings. When users request support with feature selection or algorithm parametrization by toggling the *Help me decide* menu, the system provides, e.g., feature relevance scores or silhouette coefficients for selecting the number of clusters.

System Learning — *System learning* describes guidance in which the system requests user feedback with the aim of improving its user, task or data model. While this operation may or may not have an immediate benefit to the analysis process, the gathered information can be used to improve further guidance as it helps systems to better understand users and tasks.

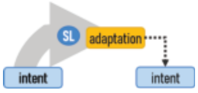

Micallef et al. [MSM*17] developed an application that supports users during the generation of machine learning models with small data sets. The system employs a user model and asks users to refine features in a subset of the overall features by assigning user relevance for the overall prediction task. This step is initiated by the system to learn the user's domain knowledge, repeating the knowledge elicitation step as many times as necessary until the prediction model returns improved predictions. Further approaches include feedback-driven view exploration [BKSS14] and DataTone [GDA*15].

5. Conclusion & Research Opportunities

We have introduced the process of co-adaptive guidance in visual analytics, building on the principles of initiation and adaptation. The concept motivates four guidance dynamics, characterized in terms of *learning* and *teaching*. This characterization emphasizes the mixed-initiative nature of guidance for effective human-machine collaboration, and operationalizes how systems and users converge towards a common analysis process. Furthermore, we have argued for the extension of guidance degrees to cover system guidance. To conclude, we discuss open challenges through providing an overview of promising research opportunities for co-adaptive guidance.

Selecting Appropriate Dynamics — In a co-adaptive guidance process, both users and systems need to employ appropriate guidance dynamics. For novice users, more system-initiated teaching might be appropriate, while expert users can initiate teaching guidance themselves or respond more faithfully during system learning. As with selected guidance degree, applications should support changing dynamics during the progression of the analysis.

Degree of Adaptation — As users provide orienting, directing, and prescribing guidance, systems have to decide to what extent they adapt their models and incorporate the available information. Setting the correct learning rate determines not only both the stability and the adaptability of the system, but also to what extent users might regard the guidance as being successful. More generally, future work should investigate the effects of rejecting provided (teaching) guidance, reacting by initiating "corrective" guidance instead.

Communication of Intent — As intent plays a central role in the co-adaptive guidance process, researchers should not only investigate how it is best communicated, but also how to deal with failures in intent identification. One possible solution could be the introduction of explicit learning or teaching modes.

Quality of (Teaching) Guidance — The goal of teaching guidance is to adapt the (mental) models of the recipient. Especially in *user teaching*, further research should determine how to capture the perceived quality of the provided guidance, e.g., in relation to contained explanations and the amount of adaptation induced.

References

[ALA*18] ANDRIENKO N., LAMMARSCH T., ANDRIENKO G., FUCHS G., KEIM D., MIKSCH S., RIND A.: Viewing Visual Analytics as Model Building. *Computer Graphics Forum 37*, 6 (2018), 275–299. doi:10.1111/cgf.13324. 1, 3

[BKSS14] BEHRISCH M., KORKMAZ F., SHAO L., SCHRECK T.: Feedback-driven interactive exploration of large multidimensional data supported by visual classifier. In *IEEE Conf. on Visual Analytics Science and Technology (VAST)* (2014), pp. 43–52. doi:10.1109/VAST.2014.7042480. 4

[CAS*18] COLLINS C., ANDRIENKO N., SCHRECK T., YANG J., CHOO J., ENGELKE U., JENA A., DWYER T.: Guidance in the human-machine analytics process. *Visual Informatics 2* (2018), 166–180. doi:10.1016/j.visinf.2018.09.003. 1, 2, 3

[CCTL15] CONATI C., CARENINI G., TOKER D., LALLÉ S.: Towards user-adaptive information visualization. In *Proc. AAAI Conf. on Artificial Intelligence* (2015), pp. 4100–4106. 3

[CD18] CAVALLO M., DEMIRALP C.: Clustrophile 2: Guided visual clustering analysis. *IEEE Trans. on Visualization and Computer Graphics* (2018). doi:10.1109/TVCG.2018.2864477. 4

[CGM*17] CENEDA D., GSCHWANDTNER T., MAY T., MIKSCH S., SCHULZ H.-J., STREIT M., TOMINSKI C.: Characterizing Guidance in Visual Analytics. *IEEE Trans. on Visualization and Computer Graphics 23*, 1 (2017), 111–120. doi:10.1109/TVCG.2016.2598468. 1, 2, 3

[CGM19] CENEDA D., GSCHWANDTNER T., MIKSCH S.: A Review of Guidance Approaches in Visual Data Analysis: A Multifocal Perspective. *Computer Graphics Forum 38*, 3 (2019), 861–879. doi:10.1111/cgf.13730. 1, 2, 3

[EFN12] ENDERT A., FIAUX P., NORTH C.: Semantic interaction for visual text analytics. In *Proc. SIGCHI Conf. Human Factors in Computing Systems* (2012), ACM, p. 473. doi:10.1145/2207676.2207741. 4

[FWR*17] FEDERICO P., WAGNER M., RIND A., AMOR-AMORÓS A., MIKSCH S., AIGNER W.: The Role of Explicit Knowledge: A Conceptual Model of Knowledge-Assisted Visual Analytics. In *IEEE Conf. on Visual Analytics Science and Technology* (2017), pp. 92–103. doi:10.1109/VAST.2017.8585498. 2

[GBDL15] GALLINA P., BELLOTTO N., DI LUCA M.: Progressive co-adaptation in human-machine interaction. In *Int. Conf. on Informatics in Control, Automation and Robotics* (2015), vol. 02, pp. 362–368. 2

[GDA*15] GAO T., DONTCHEVA M., ADAR E., LIU Z., KARAHALIOS K. G.: Datatone: Managing ambiguity in natural language interfaces for data visualization. In *Proc. Symp. on User Interface Software & Technology* (2015), ACM, pp. 489–500. doi:10.1145/2807442.2807478. 4

[GZ09] GOTZ D., ZHOU M. X.: Characterizing Users' Visual Analytic Activity for Insight Provenance. *Information Visualization 8*, 1 (2009), 42–55. doi:10.1057/ivs.2008.31. 2

[Hee19] HEER J.: Agency plus automation: Designing artificial intelligence into interactive systems. *Proc. National Academy of Sciences 116*, 6 (2019), 1844–1850. Publisher: National Academy of Sciences Section: Colloquium Paper. doi:10.1073/pnas.1807184115. 3

[Hor99] HORVITZ E.: Principles of mixed-initiative user interfaces. In *Proc. SIGCHI Conf. on Human Factors in Computing Systems* (1999), ACM, pp. 159–166. doi:10.1145/302979.303030. 1

[ITB17] INTHARAH T., TURMUKHAMBETOV D., BROSTOW G. J.: Help, It Looks Confusing: GUI Task Automation Through Demonstration and Follow-up Questions. In *Proc. Int. Conf. on Intelligent User Interfaces* (2017), ACM, pp. 233–243. doi:10.1145/3025171.3025176. 3

[KPHH11] KANDEL S., PAEPCKE A., HELLERSTEIN J., HEER J.: Wrangler: Interactive visual specification of data transformation scripts. In *Proc. Conf. Human Factors in Computing Systems (EA)* (2011), ACM, p. 3363. doi:10.1145/1978942.1979444. 4

[KT94] KROGSÆTER M., THOMAS C. G.: Adaptivity: system-initiated individualization. In *Adaptive user support: ergonomic design of manually and automatically adaptable software*. L. Erlbaum Associates Inc., 1994, pp. 67–96. 1, 3

[MSM*17] MICALLEF L., SUNDIN I., MARTTINEN P., AMMAD-UD DIN M., PELTOLA T., SOARE M., JACUCCI G., KASKI S.: Interactive Elicitation of Knowledge on Feature Relevance Improves Predictions in Small Data Sets. In *Proc. Int. Conf. on Intelligent User Interfaces* (2017), ACM, pp. 547–552. doi:10.1145/3025171.3025181. 4

[MSSM17] MORANA S., SCHACHT S., SCHERP A., MAEDCHE A.: A review of the nature and effects of guidance design features. *Decision Support Systems 97* (2017), 31–42. doi:10.1016/j.dss.2017.03.003. 1

[NNPS17] NIKOLAIDIS S., NATH S., PROCACCIA A., SRINIVASA S.: Game-Theoretic Modeling of Human Adaptation in Human-Robot Collaboration. In *Proc. ACM/IEEE Int. Conf. on Human-Robot Interaction* (2017), ACM, pp. 323–331. doi:10.1145/2909824.3020253. 2

[Odo14] ODORA R. J.: Using explanation as a teaching method: How prepared are high school technology teachers in free state province, south africa? *Journal of Social Sciences 38*, 1 (2014), 71–81. doi:10.1080/09718923.2014.11893238. 2

[Opp94] OPPERMANN R. (Ed.): *Adaptive user support: ergonomic design of manually and automatically adaptable software*. L. Erlbaum Associates Inc., 1994. 1

[SAC*17] SIMARD P., AMERSHI S., CHICKERING M., EDELMAN PELTON A., GHORASHI S., MEEK C., RAMOS G., SUH J., VERWEY J., WANG M., WERNSING J.: *Machine Teaching: A New Paradigm for Building Machine Learning Systems*. Tech. Rep. MSR-TR-2017-26, 2017. 4

[Saw05] SAWARAGI T.: Dynamical and Complex Behaviors in Human-Machine Co-Adaptive Systems. *IFAC Proc. Volumes 38*, 1 (2005), 94–99. doi:10.3182/20050703-6-CZ-1902.01418. 2

[SGL09] SHRINIVASAN Y. B., GOTZ D., LU J.: Connecting the dots in visual analysis. In *IEEE Symp. on Visual Analytics Science and Technology* (2009), IEEE, pp. 123–130. doi:10.1109/VAST.2009.5333023. 4

[SSES17] SHAO L., SILVA N., EGGELING E., SCHRECK T.: Visual exploration of large scatter plot matrices by pattern recommendation based on eye tracking. In *Proc. ACM Ws. on Exploratory Search and Interactive Data Analytics* (2017), Glowacka D., Milios E., Soto A. J., Paulovich F., (Eds.), ACM, pp. 9–16. doi:10.1145/3038462.3038463. 4

[SSS*14] SACHA D., STOFFEL A., STOFFEL F., KWON B. C., ELLIS G., KEIM D. A.: Knowledge Generation Model for Visual Analytics. *IEEE Trans. on Visualization and Computer Graphics 20*, 12 (2014), 1604–1613. doi:10.1109/TVCG.2014.2346481. 2

[SSV*18] SILVA N., SCHRECK T., VEAS E., SABOL V., EGGELING E., FELLNER D. W.: Leveraging eye-gaze and time-series features to predict user interests and build a recommendation model for visual analysis. In *Proc. ACM Symp. on Eye Tracking Research & Applications* (2018), ACM, pp. 1–9. doi:10.1145/3204493.3204546. 4

[Tiß93] TISSEN A.: Knowledge bases for user guidance in information seeking dialogues. In *Proc. Int. Conf. on Intelligent User Interfaces* (1993), ACM, pp. 149–156. doi:10.1145/169891.169945. 1

[vW06] VAN WIJK J.: Views on Visualization. *IEEE Trans. on Visualization and Computer Graphics 12*, 4 (2006), 421–432. doi:10.1109/TVCG.2006.80. 1

[WDC*18] WALL E., DAS S., CHAWLA R., KALIDINDI B., BROWN E. T., ENDERT A.: Podium: Ranking Data Using Mixed-Initiative Visual Analytics. *IEEE Trans. on Visualization and Computer Graphics 24*, 1 (2018), 288–297. doi:10.1109/TVCG.2017.2745078. 3, 4

[WSL*19] WALCH A., SCHWARZLER M., LUKSCH C., EISEMANN E., GSCHWANDTNER T.: Lightguider: Guiding interactive lighting design using suggestions, provenance, and quality visualization. *IEEE Trans. on Visualization and Computer Graphics 26*, 1 (2019), 569–578. doi:10.1109/TVCG.2019.2934658. 4