EuroVis Workshop on Visual Analytics (2023)
M. Angelini and M. El-Assady (Editors)

A Methodology for Task-Driven Guidance Design

Ignacio Pérez-Messina, Davide Ceneda, Silvia Miksch

TU Wien, Institute of Visual Computing & Human-Centered Technology, Austria

Abstract

Mixed-initiative Visual Analytics (VA) systems are becoming increasingly important; however, the design of such systems still needs to be formulated. We present a methodology to aid and structure the design of guidance for mixed-initiative VA systems consisting of four steps: (1) defining the target of analysis, (2) identifying the user search tasks, (3) describing the system guidance tasks, and (4) specifying which and when guidance is provided. In summary, it specifies a space of possible user tasks and then maps it to the corresponding space of guidance tasks, using recent VA task typologies for guidance and visualization. We illustrate these steps through a case study in a real-world model-building task involving decision-making with unevenly-spaced time-oriented data. Our methodology's goal is to enrich existing VA systems with guidance, being its output a structured description of a complex guidance task schema.

CCS Concepts

• ***Human-centered computing*** → ***Visualization design and evaluation methods;*** *Visualization theory, concepts and paradigms;*

1. Introduction

Mixed-initiative Visual Analytics (VA) environments—interfaces for visual data analysis incorporating varying degrees of automated support during data analysis—are becoming more prominent as the complexity of analytic workflows and models increases. Guidance has been proposed as a conceptual framework for understanding how mixed-initiative VA can actively facilitate users in accomplishing analytical tasks [CGM*17]. Beyond user support, it also sheds light on the potential impacts, both beneficial and adverse, of system recommendations on the analytical discourse and their contributions towards knowledge generation in VA [PMCEA*22]. Systems featuring such proactive support or embracing a human-in-the-loop approach can be unified and characterized within the framework of guidance theory [CGM19a].

Visualization task taxonomies and typologies have been proposed for understanding the VA process from a fine-grained, often behavioral, perspective. Likewise, guidance tasks have been recently introduced to reason about and analyze systems incorporating guidance [PMCEA*22]. However, task theory has had only a limited impact on the conception of VA solutions, mainly due to the lack of methodological frameworks to structure and aid the design of VA systems. Furthermore, the design of guidance-enhanced VA systems adds another challenge to the process. Although there are attempts at defining guidelines for mixed-initiative systems (e.g., Horvitz [Hor99], Liu et al. [LDT*20], Ceneda et al. [CAA*20]), and a method for guidance design based on decision support has been recently proposed [HS22], we believe there is a largely unexploited potential in task theory to provide formal design tools for guidance.

In this paper, we introduce a systematic, step-by-step methodology to aid in the design of guidance for the enhancement of VA systems (Sec. 3). Our approach is based on fundamental definitions from recent VA literature, namely user visualization tasks [BM13] and system guidance tasks [PMCEA*22]. The proposed methodology consists of four steps: (1) defining the *target of analysis*, (2) describing the *user search tasks* for this target, (3) describing the *guidance tasks* in terms of all the guidance degrees that can be provided to the defined user tasks, and (4) defining what and when guidance degrees are provided. Our methodology is model- and technique-agnostic; hence, it is an instrument for formalizing guidance design. We illustrate this methodology through a real-world case study. Finally, we discuss its advantages and limitations (Sec. 4).

2. Background

We start by illustrating important terminology for our methodology, in particular, terms and concepts referring to model building, visualization and guidance tasks.

VA as model building — Andrienko et al. define the VA process as a "goal-oriented workflow producing a model as a result" [ALA*18], where *model* is "any representation of aspects of a subject and relationships between them" [ALA*18]. A model is thus the final outcome of the VA process and it strives for *appropriateness* (congruent with the reality it depicts and with the task at hand). Data itself is not the primary interest of the analyst, but

DOI: 10.2312/eurova.20231094

only insofar it is a source of *evidence* that supports the model of the real subject. The *target of analysis* (and of a visualization task) is such evidence, that for each problem has a different manifestation in the data. The role of guidance in model-generating activities has already been discussed by Collins et al. [CAS*18].

Visualization tasks — Visualization tasks in the literature share a triadic structure: they have an input, an output, and a target. The target is "[the] part of the data in which [a visualization task] is carried out" [SNHS13], or the part of the data that is *queried* within the search task's input [BM13]. The input is the whole of the representation where the target is searched for. The output is a new piece of information, and action to take, or a representation (i.e., the original target, a new target, or a path). For the output to be a new piece of data or representation acknowledged by the system, the search task must be accompanied by a *produce* task.

Search tasks — Brehmer & Munzner introduced the four search types to characterize tasks according to *what* is the object of search (i.e., the target) [BM13]. The characterization into *lookup*, *browse*, *locate*, and *explore* search types depends on whether the identity (i.e., reference) of the target is known a priori and (independently) whether the location of it is known a priori. The target of analysis, as the evidence that the users search for to create and support their model, corresponds with the target of the user search tasks. In *lookup* and *locate* type searches this target is known while *browse* and *explore* searched it is unknown, also meaning that in the latter types, the appearance of a particular piece of evidence is not forced and its absence not meaningful.

Guidance — Guidance is defined as the active process of resolving "knowledge gaps" or providing useful assistance even when there is none. Similar to visualization tasks, the guidance function $guidance(gap, input) \rightarrow answer$ receives a guidance input, targets a knowledge gap and outputs an answer. The knowledge gap is classified into *target unknown* and *path unknown*, thus relating the guidance answer directly to the user search tasks. The guidance answer can come in different levels of strength, or user agency limitation, called guidance *degrees* [CGM*17]: (from highest to lowest) prescribe, direct, and orient. These range from providing a complete solution to the knowledge gap, a set of options or only giving cues to the user.

User-guidance task interaction — Guidance degrees are classified as *disruptive* if they impose a change in the target user task by limiting user agency (prescribing and directing) and *non-disruptive* if they conserve the user task (orienting) [PMCEA*22]. This "target user task → guidance task → outcome user task" flow is encapsulated in the guidance task function

$$guidanceTask(userTask_{target}, degree) \rightarrow userTask_{outcome}. \quad (1)$$

Moreover, guidance degrees can also be classified according to their relation with specific user search types, i.e., to the different knowledge gaps they aim to solve. There are in total seven guidance *second-order degrees*, also considering *prescribing* guidance, which does not have a second-order degree. Our proposed methodology consists, in a few words, in defining all user tasks (using the above mentioned taxonomy) and matching them with the appropriate guidance tasks, as shown in Fig. 1.

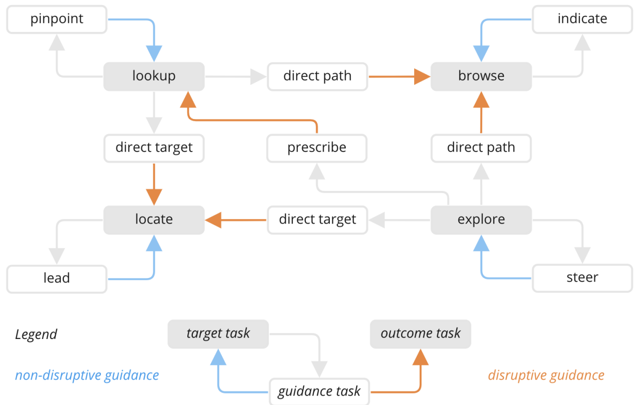

Figure 1: *Guidance and user tasks and their interactions in terms of guidance second-order degrees (white) and user search types (grey). A user task is always targeted by a guidance task (grey arrow) and returned to itself (non-disruptive guidance, represented by blue arrows) or a different user task (disruptive guidance, represented by orange arrows). [PMCEA*22]*

3. Methodology for Guidance Design

Our proposed methodology consists of four steps:

1. **Defining the target of analysis**
   Define what the target of the analysis is, i.e., what type of data artifacts and possible insights users are looking for, as this will be the target of the search tasks. If more than a target is present (usually when more than one analytical objective is present in the analysis workflow), follow the next steps independently for each target.
2. **Describing the four user search tasks for the target**
   Once a target is defined, structure the analysis workflow according to the user search tasks. The search types *lookup*, *browse*, *locate*, and *explore* each stand for a class of active approaches the users will have to reach the target of analysis (a class of user tasks). Each of these tasks produces a different output that is or leads to a piece of evidence for the analysis outcome.
3. **Describing the seven guidance second-order degrees for the user search tasks**
   Afterward, having defined the user tasks, we have to match them with the guidance tasks to support and enhance their execution with guidance. The guidance tasks are *pinpoint*, *indicate*, *lead*, *steer*, *direct target*, *direct path*, and *prescribe* in terms of how they contribute to or assist the execution of the correspondent target user tasks.
4. **Defining what and when guidance is provided**
   Verify that guidance tasks are provided timely by correctly associating them with visualization states and decide excluding certain guidance degrees if appropriate.

We illustrate our methodology, which can be swiftly reviewed in Fig. 2, with a case study motivated by a domain-specific task which can be described as the design of a VA system to analyze unevenly-spaced time-oriented data. As the actual task is complex and requires ample domain knowledge, many details regarding relationships and attributes of the data are omitted, for sake of brevity. The case study corresponds to the process of designing a guidance-

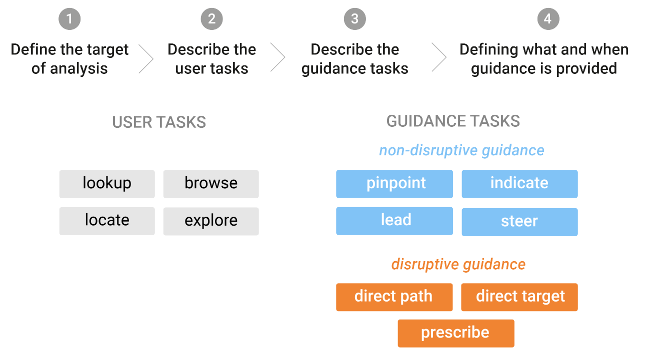

Figure 2: *Schema of our proposed 4-step methodology for guidance design. It consist of defining the interdependent user tasks and guidance tasks in a systematic way.*

enhanced VA system, currently under development in collaboration with domain expert, to support the analysis of historical images of unexploded ordnance in time. Next, we provide additional details of our scenario and then show how we applied our design methodology.

Image Selection for UXO detection — UneXploded Ordnance (UXO) detection is the task of mapping historical bomb droppings (such as during wars) within an Area Of Interest (AOI). The goal is to make sure that such bombs do not pose a risk to people or workers on construction sites. The final output of the analysis process is a risk map of all geographical points where undetonated bombs (UXOs) possibly (or almost certainly) can be found. This is achieved through several stages of analysis where the main object is archival aerial photographs. The first of these stages is Image Selection, where a relatively small subset from all available images that have coverage over the AOI (rounding the hundreds even for small AOIs) must be selected before proceeding to the next phases of analysis (georeferencing and crater detection). Another important input for this task is a record of the dates of aerial attacks where the bombings took place, as images are selected in order to have geographical coverage over the AOI *and* temporal coverage over the attacks.

Image Selection as VA model building — The task of Image Selection for UXO detection conforms to a VA task as its outcome is effectively a *model* of physical territory and a description of all damage suffered *in time* by buildings and infrastructures. This model, a historical 4D reconstruction, is made of aerial images which contain evidence of the UXOs (i.e., the bombs). Theoretically, there are 2^n possible selections (subsets) in a set of images, where n is the number of images. Each of them constitutes a different model (although some may be very similar). However, only a very small fraction of these subsets make for an *appropriate* model. According to Andrienko et al. [ALA*18] an appropriate model is one that complies with the requirements of *correct*, *fit to the purpose*, *comprehensive*, *sufficient scope*, *generalization*, *parsimony*, *specificity* and *resource efficiency*. The purpose of the analysis is to find a subset of the images that serve as evidence for such an appropriate model. It is a wicked problem, as many variables must be optimized and not all of them can be apprehended by a computer model, a Pareto front of solutions may exist and the decision

Figure 3: *Detail of the VA prototype interface that serves as input for our example case, consisting of two views: the timeline visualization (above) shows the available images organized in structured corresponding to flights; the map (below) shows the same images in their geospatial dimension.*

over a final set has ethical consequences. As can be observed, the task fits into many known problem categories: multivariate interactive optimization [LDT*20], combinatorial optimization, discrete parameter space exploration [SHB*], etc.

VA system — As starting point for the application of our method, we consider the VA system shown in Fig. 3. The system is composed of multiple coordinated views of a timeline and a map where the images are visualized as glyphs (grouped according to the flights they belong to), and a heuristic model that can evaluate a selection's appropriateness and thus can be used for guidance. The model captures two opposing aspects of a selection: temporal coverage and economy, effectively modeling it as an optimization problem (similar to the knapsack problem [PRP10]). The supported interactions on the timeline visualization are image hover (to reveal a tooltip with metadata), image de/selection, temporal zooming-in, temporal filtering, and placing flags on the timeline. In the following, we show how we applied our design methodology to enhance the system with guidance.

3.1. Step 1: Define the target of analysis

The goal of the user is to find a set of images (a subset of the whole dataset) that contributes to defining an appropriate reconstruction of the damages suffered by buildings and possible UXOs lost in the terrain of the area of interest. The target, thus, is *an image or a group of images*. The user usually targets a group of aerial images at a time (e.g., a whole flight) in order to find one or two to add to the selection. A single image may also be the target (e.g., to get its metadata) and, of course, the whole current selection of images can be considered and refined. When we speak of the path, in opposition to the target, we usually refer to the temporal position of an image, as time is the main ordering principle of the timeline, and to the actions taken in order to find an image.

3.2. Step 2: Define the search tasks

The second step of the methodology is to describe the four search tasks for the target defined in the first step. Here, we provide a *general definition* of each search task, followed by its description in the context of the case study, also describing *how*, in terms of views and actions, it is performed. With these first two steps, we will have arrived at a loose but good enough description of *why*, *how* and *what* of the user tasks [BM13].

Lookup Task — *Search for the target's attributes or find the target across different representations.*

Lookup an image's metadata or lookup an image from the map into the timeline or viceversa: display a tooltip with information about the image on hover and highlight image on map.

Browse Task — *Search for a target with the desired attributes within a promising subset of data or spatial locality.*

Browse candidate images within a temporal frame. This task occurs when the user is zoomed into a time frame to have a closer look at a subset of the images.

Locate Task — *Search for a subset of the data or spatial locality (path) where the target should be found.*

Locate a time frame where a good image should be selected. This task can be performed implicitly by zooming in on a period in the timeline or explicitly by placing (*producing*) a flag at a certain point in the timeline.

Explore Task — *Search for* possible *behavior in the data, whose reference or characteristics appears only in reference to the whole of the representation.*

Explore possible temporal partitions and image combinations in relation to their effect on the quality of the final selection. Explore tasks are the hardest to perform and detect [BH19]; however, as we have already confined the other tasks to specific analytical moments, we can leave the exploratory task as happening everywhere else, i.e., only when the user is at the overview level of the visualization.

3.3. Step 3: Describe the guidance tasks

The third step consists of making a description of the guidance in relation to the above-defined user tasks. In other words, in this step, we aim to match each user task described above with the most appropriate guidance task. We organize this section into the three guidance degrees (orienting, directing, and prescribing), which we briefly describe, and for each of its second-order degrees, we provide a *general definition* followed by their illustration within the case study, as before. This description should encapsulate the *why*, *how*, and *what* of system guidance tasks [PMCEA*22].

3.3.1. Orienting guidance tasks

The main function of orienting guidance is to enhance or help maintain the users' mental map [CGM*17]. Any visualization technique can be used to encode orienting guidance as this kind of guidance only adds a new layer of information and does not constrain user action. The difference between orienting guidance and the rest of the visual environment lies in what it communicates: as all guidance, it is a product of a *guidance model*, it is not the raw data or model under analysis, and so it is a different source of information as well as of uncertainty [SSK*15].

Pinpoint (Orient → Lookup) — *Provide cues about relevant aspects of the lookup target, or about similar or related cases.*

In a lookup task, the user does not have any knowledge gaps, however, relevant information which is otherwise not shown about the target (hovered) image can be encoded in the map and timeline, e.g., its temporal and geographical coverage, its rank on the overall distribution of images, its effect on the overall quality of the current selection, the dimensions in which it is better or worse than its possible replacements, etc.

Indicate (Orient → Browse) — *Provide cues about relative interest for the task within the browsed elements.*

Encode in the color of the glyph the interest of an image relative to its surrounding images or possible replacements.

Lead (Orient → Locate) — *Provide cues about data subsets, spaces or analysis paths that are relevant for analysis.*

Highlight in the timeline the temporal locations that are most likely to hold images of interest.

Steer (Orient → Explore) — *Provide cues about important features in regards to the whole of the data or the whole progress of analysis.*

Highlight unvisited images that might be of particular interest, e.g., images that would have an important impact on quality of the solution.

3.3.2. Directing guidance tasks

Through directing, the system provides a list of options to the user, either of data cases (*direct path*) or actions to take (*direct target*) [PMCEA*22]. Directing guidance can be provided during explore and lookup tasks: in the former case, we say guidance is *convergent* as it aims at reaching a common ground with the user (i.e., providing confirmatory evidence), while in the latter we say it is *divergent* as its aim is to show a different, maybe even contradictory, option to the one the user has decided upon (i.e., *disconfirmatory evidence* [SSK*15]).

Direct Path (Explore/Lookup → Browse) — *Provide a ranked list of data cases which the user can browse.*

In the convergent case, suggest images for selection while on the overview level. In the divergent case, suggest image that could serve as replacements of a hovered image.

Direct target (Explore/Lookup → Locate) — *Provide a ranked list of actions that the user can take over the data, representation, or model. (e.g., zoom into certain area, add/delete element, change colormap, etc.)*

In the convergent case, provide suggestion of temporal regions to visit where a suitable target may be found while on the overview level. In the divergent case, suggest alternative time periods where an image of image may be found.

3.3.3. Prescribing guidance task

Prescribing guidance is only provided to explore tasks, as it answers both target and path unknown knowledge gaps. Thus, it does not possess any second-order degree. As in Fig. 1, we say it turns the exploratory task into a lookup task of verification. Prescriptive guidance can and should only be provided when it is actually feasible to do so (i.e., when the guidance model is capable of performing the user task by itself).

Prescribe — *Provide a unique and complete answer to the task, either immediately encoded, step by step, or by animation. Prescribing is only appropriate when the user task is sufficiently well defined or enough information is known about it at some point of analysis.*

Provide a full subset of images that satisfies all temporal coverage and quality requirements. After the solution is presented, the user is free to verify, accept, decline, and edit it (on acceptance).

3.4. Step 4: Defining what and when guidance is provided

Given the variety of guidance tasks and their specificity to user tasks, guidance tasks should not be active at all times all at once, otherwise resulting in *guidance overlap* and consequent user confusion. Thus, it must be defined *when* each degree will be provided. This question can already be framed as "when is the user performing a certain task?", thanks to the guidance tasks being directly coupled with user tasks (i.e., we can assume a specific guidance task should be active when the user is performing a task that the guidance task targets). Moreover, as search types are part of the user *visualization* tasks, they should be related to different representations of the data (which can be defined by means of different views or levels of semantic zoom [WLA17] within a view, e.g., at discrete levels of detail or abstraction in graph structures [BBT23]). Also, depending on the context and the information available about the user and task, not every guidance degree must and should be provided.

As per Fig. 1, the user task lookup has three guidance tasks that can target it (*pinpoint*, *direct target*, and *direct path*) and explore has four (*steer*, *direct path*, *direct target*, and *prescribe*). These cases need disambiguation. Non-disruptive guidance can coexist with disruptive guidance types, however, disruptive guidance is mutually exclusive as it forcefully sends the user to a different task, distinctive of each guidance task (*prescribe* $\rightarrow$ *lookup*, *direct path* $\rightarrow$ *browse*, *direct target* $\rightarrow$ *locate*). The appropriate guidance degree to use may not depend on time, but on the user's level of expertise (i.e., to the potential knowledge gaps of the user) [CGM19b]. The choice of which level of guidance to use thus can be left to the user who is informed on how to take an appropriate decision, for example, through onboarding [SCW*22]. The idea that the guidance degree of a guidance task can be automatically decided and adjusted by the guidance system is still an open research topic.

In our application example, each user task is associated with a certain part or semantic zoom level of the representation. As a consequence, orienting guidance tasks can be provided without overlap in the same way as user tasks are partitioned across analytical moments (as described in Steps 2-3). If the underlying guidance model is appropriate enough to provide a selection by itself, this should be prescribed from the start of the analysis so that users can decide to accept the suggestion (and then edit however they see fit) or to decline it and start from scratch, in which case the guidance system can switch to provide directing guidance.

4. Discussion & Conclusion

We have obtained, by the application of our method, seven guidance tasks spanning all guidance degrees and, according to the system guidance task typology, covered all possible manifestations of first-loop guidance (i.e., static guidance, non user-adaptive). We have also verified that there is no guidance overlap and that all user tasks are supported by guidance. In other words, all task-related knowledge gaps have been covered. Assuming a target user which is in no need for onboarding, the effectiveness and suitability of this guidance solution, its correspondence to the domain problem at hand and the user tasks, follows only from the appropriate application of each consecutive step.

Our approach facilitates not only the design but also the evaluation of guidance-enriched VA systems by isolating both user and guidance tasks, enabling a task-based evaluation where users can assess visualization and guidance components independently, by performing a task with and without guidance. This would also make our methodology compatible with Munzner's nested evaluation framework [Mun09], where a fifth *guidance layer* with clear dependencies to the inner layers could be added.

According to the guidance task typology and guided knowledge generation model [PMCEA*22], we have here only covered *provide (guidance)* tasks. The typology itself does not cover in detail *observe*, *expect*, and *adapt* tasks yet, and so this methodology is mainly for designing static guidance (i.e., not user-adaptive). How to include co-adaptive approaches [SJB*21,SSKEA21], which necessarily involve the inference loop of the guided knowledge generation model [PMCEA*22], is left as a future research direction. Many relevant aspects that impact the effectiveness of guidance, such as guidance generated uncertainty [SSK*15] and model trust building [LMDT23], are out of the scope as well. Another limitation of our work is that it constrains itself to user *search* tasks, while not considering how guidance could be defined in relation to other task spaces (e.g., *disseminative*, *observational*, *analytical* and *model-developmental* visualization tasks [CG15]).

Conclusion — In this paper, we introduced a systematic and structured methodology for designing guidance-enhanced VA environments. Our methodology takes an existing VA system and provides a framework for developing a comprehensive guidance space to enrich it. By incorporating concepts from visualization and guidance task literature, we demonstrated the utility of a task-driven framework in the design of VA solutions. A case study involving model building with unevenly-spaced time-oriented data showed the feasibility and efficacy of our methodology, instantiating its potential in enhancing the design of VA systems.

Acknowledgements. This work was funded by Vienna Science and Technology Fund (WWTF) under grant [10.47379/ICT19047], by the Austrian Science Fund (FWF) with the grant P31419-N31 (KnoVA) and by DoRIAH (FFG Grant #880883).

References

[ALA*18] Andrienko N., Lammarsch T., Andrienko G., Fuchs G., Keim D., Miksch S., Rind A.: Viewing visual analytics as model building. In *Computer Graphics Forum* (2018), vol. 37, Wiley Online Library, pp. 275–299. 1, 3

[BBT23] Berger P., Beleites S., Tominski C.: Comparing nodes of multivariate graphs through dynamic layout adaptations. *arXiv preprint arXiv:2303.00528* (2023). 5

[BH19] Battle L., Heer J.: Characterizing exploratory visual analysis: A literature review and evaluation of analytic provenance in tableau. In *Computer Graphics Forum* (2019), vol. 38, Wiley Online Library, pp. 145–159. 4

[BM13] Brehmer M., Munzner T.: A multi-level typology of abstract visualization tasks. *Transactions in Visualization and Computer Graphics 19*, 12 (2013), 2376–2385. 1, 2, 4

[CAA*20] Ceneda D., Andrienko N., Andrienko G., Gschwandtner T., Miksch S., Piccolotto N., Schreck T., Streit M., Suschnigg J., Tominski C.: Guide me in analysis: A framework for guidance designers. In *Computer Graphics Forum* (2020), vol. 39, Wiley Online Library, pp. 269–288. 1

[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*, 3 (2018), 166–180. 2

[CG15] Chen M., Golan A.: What may visualization processes optimize? *IEEE transactions on visualization and computer graphics 22*, 12 (2015), 2619–2632. 5

[CGM*17] Ceneda D., Gschwandtner T., May T., Miksch S., Schulz H.-J., Streit M., Tominski C.: Characterizing guidance in visual analytics. *Transactions in Visualization and Computer Graphics 23*, 1 (2017), 111–120. 1, 2, 4

[CGM19a] Ceneda D., Gschwandtner T., Miksch S.: A review of guidance approaches in visual data analysis: A multifocal perspective. In *Computer Graphics Forum* (2019), vol. 38, Wiley Online Library, pp. 861–879. 1

[CGM19b] Ceneda D., Gschwandtner T., Miksch S.: You get by with a little help: the effects of variable guidance degrees on performance and mental state. *Visual Informatics 3*, 4 (2019), 177–191. 5

[Hor99] Horvitz E.: Principles of mixed-initiative user interfaces. In *Proceedings of the SIGCHI conference on Human Factors in Computing Systems* (1999), pp. 159–166. 1

[HS22] Han W., Schulz H.-J.: Providing visual analytics guidance through decision support. *Information Visualization* (2022), 14738716221147289. 1

[LDT*20] Liu J., Dwyer T., Tack G., Gratzl S., Marriott K.: Supporting the problem-solving loop: Designing highly interactive optimisation systems. *IEEE Transactions on Visualization and Computer Graphics 27*, 2 (2020), 1764–1774. 1, 3

[LMDT23] Liu J., Marriott K., Dwyer T., Tack G.: Increasing user trust in optimisation through feedback and interaction. *ACM Transactions on Computer-Human Interaction 29*, 5 (2023), 1–34. 5

[Mun09] Munzner T.: A nested model for visualization design and validation. *Transactions in Visualization and Computer Graphics 15*, 6 (2009), 921–928. 5

[PMCEA*22] Pérez-Messina I., Ceneda D., El-Assady M., Miksch S., Sperrle F.: A typology of guidance tasks in mixed-initiative visual analytics environments. In *Computer Graphics Forum* (2022), vol. 41, Wiley Online Library, pp. 465–476. 1, 2, 4, 5

[PRP10] Puchinger J., Raidl G. R., Pferschy U.: The multidimensional knapsack problem: Structure and algorithms. *INFORMS Journal on Computing 22*, 2 (2010), 250–265. 3

[SCW*22] Stoiber C., Ceneda D., Wagner M., Schetinger V., Gschwandtner T., Streit M., Miksch S., Aigner W.: Perspectives of visualization onboarding and guidance in va. *Visual Informatics 6*, 1 (2022), 68–83. 5

[SHB*] Sedlmair M., Heinzl C., Bruckner S., Piringer H., Möller T.: Visual Parameter Space Analysis: A Conceptual Framework. 2161–2170. doi:10.1109/TVCG.2014.2346321. 3

[SJB*21] Sperrle F., Jeitler A., Bernard J., Keim D., El-Assady M.: Co-adaptive visual data analysis and guidance processes. *Computers & Graphics 100* (2021), 93–105. URL: https://www.sciencedirect.com/science/article/pii/S009784932100131X, doi:https://doi.org/10.1016/j.cag.2021.06.016. 5

[SNHS13] Schulz H.-J., Nocke T., Heitzler M., Schumann H.: A design space of visualization tasks. *IEEE Transactions on Visualization and Computer Graphics 19*, 12 (2013), 2366–2375. 2

[SSK*15] Sacha D., Senaratne H., Kwon B. C., Ellis G., Keim D. A.: The role of uncertainty, awareness, and trust in visual analytics. *IEEE transactions on visualization and computer graphics 22*, 1 (2015), 240–249. 4, 5

[SSKEA21] Sperrle F., Schäfer H., Keim D., El-Assady M.: Learning contextualized user preferences for co-adaptive guidance in mixed-initiative topic model refinement. In *Computer Graphics Forum* (2021), vol. 40, Wiley Online Library, pp. 215–226. 5

[WLA17] Wiens V., Lohmann S., Auer S.: Semantic zooming for ontology graph visualizations. In *Proceedings of the Knowledge Capture Conference* (2017), pp. 1–8. 5