

A Typology of Guidance Tasks in Mixed-Initiative Visual Analytics Environments

I. Pérez-Messina¹, D. Ceneda¹, M. El-Assady², S. Miksch¹ and F. Sperrle³¹TU Wien, Austria²ETH AI Center, Zürich, Switzerland³University of Konstanz, Germany

Abstract

Guidance has been proposed as a conceptual framework to understand how mixed-initiative visual analytics approaches can actively support users as they solve analytical tasks. While user tasks received a fair share of attention, it is still not completely clear how they could be supported with guidance and how such support could influence the progress of the task itself. Our observation is that there is a research gap in understanding the effect of guidance on the analytical discourse, in particular, for the knowledge generation in mixed-initiative approaches. As a consequence, guidance in a visual analytics environment is usually indistinguishable from common visualization features, making user responses challenging to predict and measure. To address these issues, we take a system perspective to propose the notion of guidance tasks and we present it as a typology closely aligned to established user task typologies. We derived the proposed typology directly from a model of guidance in the knowledge generation process and illustrate its implications for guidance design. By discussing three case studies, we show how our typology can be applied to analyze existing guidance systems. We argue that without a clear consideration of the system perspective, the analysis of tasks in mixed-initiative approaches is incomplete. Finally, by analyzing matchings of user and guidance tasks, we describe how guidance tasks could either help the user conclude the analysis or change its course.

1. Introduction

In the traditional visual analytics (VA) process, knowledge is generated by users from data by exploiting visualizations, interaction, and the modelling capabilities of VA environments [SSS*14]. Users' domain knowledge is distilled into goals and hypotheses that are fed to the interactive system as actions. The result of the machine processing is shown back to the user, who interprets it into useful insights and integrates it as new knowledge that can be used for decision-making [KKEM10]. The conventional information visualization process has begun to be expanded into what are known as mixed-initiative approaches [CCI*15], where the system plays an active role in the analytical discourse through the incorporation of different degrees of agency. Guidance [CAS*18; CGM*16; CGM19a; SJB*21] has been developed as a theory to understand and encapsulate this phenomenon which goes beyond information visualization and is at the core of the VA promise.

Guidance in VA is characterized as an active process addressing “knowledge gaps” of the users that hinder their analytical progress by identifying them and providing orienting, directing, and prescriptive guidance [CGM*16]. This dimension, namely the “guidance degree”, has been identified over single- and mixed-initiative systems [CGM19a], proving to be an effective model to analyze systems with active user-supporting (i.e., guiding) capabilities.

Additionally, taxonomies of interaction tasks have had an important role in the design and understanding of user-initiative interactive visualization systems [BM13], as user tasks are considered the

building blocks of higher-order intellectual processes in VA. However, interaction tasks describe only one part of the story, as they only depict user intentions and interactions. System-side intelligent agents are left out of this narrative as well as their supporting role in human decision-making [DS21] and the proactive guidance they provide to the user. This makes it difficult to understand the complex behavior that arises when both human and system interact [SJB*20]. Thus, we are left with the following questions: (1) *How can we classify the system's intentions and tasks?* (2) *How does guidance contribute to the knowledge generation process?* (3) *How are the analytical discourse and the user's tasks affected by the guidance?*

In this paper, our aim is to answer these questions from a theoretical perspective. To accomplish this, we first extend the Knowledge Generation Model by Sacha et al. [SSS*14] by taking into account the contribution of the guidance system to the generation of insights and knowledge. We show that at the crossroad of human and machine agency, the analytic discourse can take different paths. We further analyze the interaction between human and machine intentions using the model provided by Brehmer and Munzner [BM13], as it succinctly captures the intentions, operations, and input/outputs of user action. We then introduce the notion of *guidance tasks* and derive our own typology of guidance actions. We illustrate in several case studies the application of our typology to show how it can be used to decompose the analysis into a series of user and system tasks by making explicit the role of guidance in VA.

We contribute the following: (1) An expansion of the Knowledge Generation Model in guidance-enriched VA environments showing

the interaction between user and guidance systems, from which we derive (2) a model of guidance degrees and how they relate to user tasks, which we call perspective change dynamics; (3) a typology of system guidance tasks covering the why, how, what, and when of guided interactions, whose use is demonstrated through (4) three case studies.

2. Related Work

To arrive at a better understanding of the role of guidance in the VA discourse, we first review the literature on guidance, analytical models, and analytical tasks.

Guidance. The term guidance was first introduced by Schulz et al. to unify under a common framework terms as “recommender systems”, “user support” and “assistance” within VA [SSMT13]. Ceneda et al. define guidance as “a computer-assisted *process* that aims to actively resolve a *knowledge gap* encountered by users during an *interactive* visual analytics session” [CGM*16, p. 2]. Several aspects of guidance have been characterized and used to classify the existing literature [CGM*18], and to describe mixed-initiative approaches, in which both the user and the VA system are considered to have an active role in analysis [CGM19b] and adapt to each other [SJB*20; SJB*21]. The study of guidance has led to novel VA techniques [SBS*18] and guidelines for design [CAA*20]. Different types of knowledge and their importance for guidance have been described [CAS*18; FWR*17]. What is still missing, though, is a deeper understanding of the role of guidance in the way insights are gained from the data and in analytical processes.

Models of Analytic Discourse. Models of analytic discourse (knowledge generation, sense-making, information retrieval, etc.) have up to now dealt only with user-initiative systems, i.e., systems where the computer plays no role apart from executing the user’s explicit actions. To arrive at an understanding of the interactions between user and system, we must extend these models to consider a higher degree of freedom in the computer, i.e., a system initiative.

Our model, which will be presented in the following sections, extends the Knowledge Generation Model proposed by Sacha et al. [SSS*14]. We chose it because it captures many preceding models and is, to the best of our knowledge, the most VA-specific. This model shows data-driven knowledge acquisition by users as a structured process composed by the computer (with a visualization system) and three cognitive loops that build upon each other. This model is not a stand-alone piece and we can trace its elements back to previous models. The coarse structural foundation can be found already in Norman’s model for cognitive engineering that pictures how the human actor interacts with the computer (or any physical system) by leaping twice through the gulf that separates them: first by translating goals into actions (the gulf of execution) and then back by interpreting the feedback of the system into something meaningful (the gulf of evaluation) [Nor86]. Norman’s model is not specific to data-driven research, but it captures the challenges any human undergoes when becoming a “user”.

The idea that there is more than one process at work in the cognitive effort of the user, an interaction-intensive low-level loop and a more intellectual high-level loop, can be found 20 years later in Pirolli and Card’s Sensemaking Process consisting of the Foraging

Loop and the Sensemaking Loop [PC05]. The computer part of the Knowledge Generation Model (see center part of Fig. 1) was first introduced by Keim et al. [KKEM10] showing how data, visualization, and model connect to human knowledge. Extending this simple model with the ideas from Norman’s model and the Sensemaking model, we retrieve the main structure of the Knowledge Generation Model, where user interaction is found in the Exploration Loop, which is controlled by the hypotheses and insights gained in the Verification Loop. The outermost loop, the Knowledge Generation Loop, includes the internalization and socialization of new knowledge [SSS*14]. In section 3, we describe a similar nested structure for what we call the “Guidance Process”.

Analytical Tasks. The concept of analytical tasks is of utmost importance for mixed-initiative VA [CCI*15], as VA has been described as a task-driven process [KAF*08; MA14]. Amar et al. argue that, for the design of effective systems, tasks must prime over representation [AES05]. The analytical discourse has been modelled as a hierarchical structure, where low-level actions are derived from high-level goals [GZ09; RAW*16], hinting that the opposite (deriving goals from actions) is possible, although this might be very challenging [BWD*19]. Furthermore, open-ended exploration in visual analysis has been empirically characterized as task-driven [BH19].

Work describing visualization tasks is abundant and specialized (e.g., for biological pathways [MMF17], network evolution [APS13], genomic data [NHG19]), however, the first to traverse the gap between high- and low-level tasks was Brehmer and Munzner’s multi-level task typology [BM13]. This general typology also allows the construction of complex task structures. Similarly, a simplified schema of Norman’s cognitive engineering model [Nor86] has been used to elicit complex skill chains in computer games [Coo21; HCD17]. This model consists of four steps (Thinking, Action, System, Feedback) that form a “skill atom”, which is a single piece of knowledge about the game mechanics acquired through the successful completion of an interaction loop.

Until now, analytical tasks have been reserved for users and analysts, as only humans are involved in the analysis process as decision-makers [DS21]. We do not intend to challenge this, however, the role of guidance is to *support* analysis and, consequently, it takes part in the decision-making process. Thus, it is necessary to complement user’s tasks models with guidance.

3. Model of Knowledge Generation in Guided VA

As we have seen, guidance and its role in the analysis are not typically considered when describing how knowledge is generated. Hence, in this section we provide an expanded version of the knowledge generation process, with the inclusion of guidance. We performed this expansion after analyzing existing guidance approaches [CGM19a]. We chose as a base for our discussion the Knowledge Generation Model described by Sacha et al. [SSS*14] because it provides a fine-grained view of the analysis process, matching well with the visualization tasks perspective, and for it being a fundamental inspiration for this work. The User Side is kept the same as in the original model (see top portion of Fig. 1). Our expansion considers the addition of a “Guide Side” (bottom portion of Fig. 1), opposed to the User Side, which interacts with the Computer

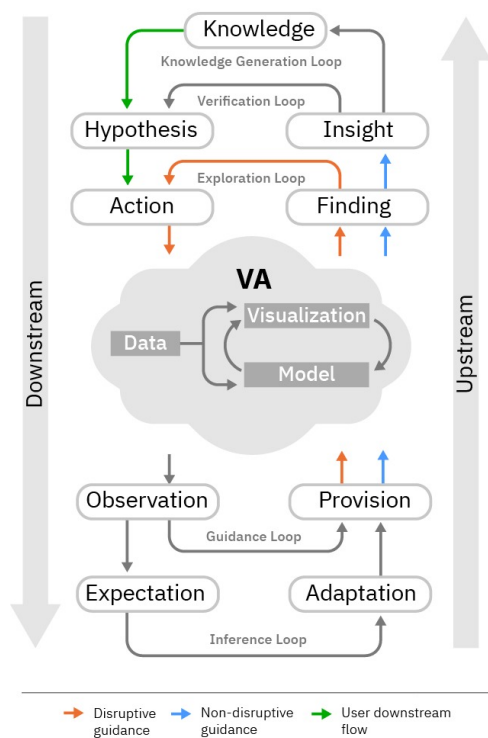


Figure 1: The Model of Knowledge Generation in Guided VA showing how guidance contributes to the progress of the analysis. The different arrows model the interactions between User (top) and Guide (bottom). Downstream (User-to-Guide) and Upstream (Guide-to-User) arrows signal the two directions in which information can flow. The model is an expansion of the well-known Knowledge Generation Model. [SSS*14].

through a Guidance Loop (mirroring the User Exploration Loop), which is controlled by an Inference Loop (mirroring the User Verification Loop). Then, we show how the information flowing from the Guide to the User can affect user actions and the progress of the analysis. In the following, we start introducing the well-known model for knowledge generation and then show how we expanded it to include guidance.

The Knowledge Generation Model. In the Knowledge Generation Model, as described by Sacha et al. [SSS*14], information flows *downstream* from the user to the system and *upstream* from the system to the user. Information can also move sideways from the upstream side to the downstream side, but not the other way around (as in Fig. 1, no arrow flows from left to right). This allows information to circulate in loops, but only the innermost, the Exploration Loop, has a direct input/output relation with the system. The Verification Loop has a similar relation to the Exploration Loop: **Findings**, i.e., trivially verifiable patterns in the data [ALA*18], serve as input for the Verification Loop, which can be interpreted as **Insights** and inform the generation of a **Hypothesis** that shapes the next **Actions** to be taken.

The Guidance Loop. We now show how guidance is generated and provided to the user by introducing two additional information

loops, represented in Fig. 1. The first one is the Guidance Loop which is the main loop of any guidance system. It directly connects with the VA system by an input/output relation, handling the downstream information coming from the user that is useful for guidance generation (see the **Observations** block in Fig. 1) and enacting consequent guidance actions (the **Provision** block). **Observations** are user interactions seen from the system perspective (downstream information, also called soft data [EFN12]) and can be used to drive the guidance that is forwarded to the user. For instance, a system can observe the user to determine when to provide guidance, what path or target to suggest, or with which guidance degree. An observation can be anything that stems from a user decision, e.g., a change in the current view, a movement of the mouse, and also the lack of interaction. However, a guidance system can also provide guidance without considering user-input, i.e., without making adaptations to the guidance (e.g., in Ip and Varshney [IV11]). These latter systems provide guidance based mostly on the dataset under analysis, and make their suggestions available from the beginning of the session, suggestions which will not undergo any changes unless the data under analysis does. The **Provision** block encompasses the actual actions and modalities of suggestion (degrees) that the guidance system can provide and how they are provided. Having the innermost Guidance Loop is the minimal requirement for a system to be considered guidance-enriched.

The Inference Loop. The Inference Loop is the second-order loop of the Guidance process. It steers the Guidance Loop similarly to how Insights and Hypotheses steer the Exploration Loop [SSS*14] on the user side; but in this case, it is information coming from the user that triggers changes on the guidance loop. Guidance inference has been defined by Ceneda et al. as the process by which guidance is derived [CGM19a]; the Inference Loop steers the Guidance Loop by capturing user *intentionality* [PSCO09]. On the downstream side of the Inference Loop we find the **Expectation** block, which receives information from the **Observations** block. Expectations have been defined as previous notions the system could possess about the user or the tasks [SJB*21]. By comparing the expectations of the system with the observations, i.e., with what the user is actually doing, the system adapts what it knows about the tasks and the user in the **Adaptation** block and derives new intentions and goals for the next round of guidance. For example, a system can start a session by assuming every data point to be equally interesting to the user, and then update this expectation based on what the user interacts with the most. When expectations are met (or not), this information is used to trigger an **Adaptation** process that will affect how guidance is provided in the future, and also how future (or even past) observations are weighted to steer guidance.

Guidance Information Flowing to the User. As the user feeds actions to the system, guidance is provided, and it is perceived by the user as the rest of the information coming from the visualization. The guidance information flows upstream (from the system to the user) similarly to the visual information about the data (and other abstract forms of information as uncertainty [SSK*15]). The guidance suggestions reach at some point, unless unnoticed or immediately discarded, the status of a **Finding**. From there, suggestions can follow two different paths (as illustrated in Fig. 1 by the blue and red paths): it can either be interpreted along with other findings into **Insight**, which will naturally have an effect on the **Hypothesis**

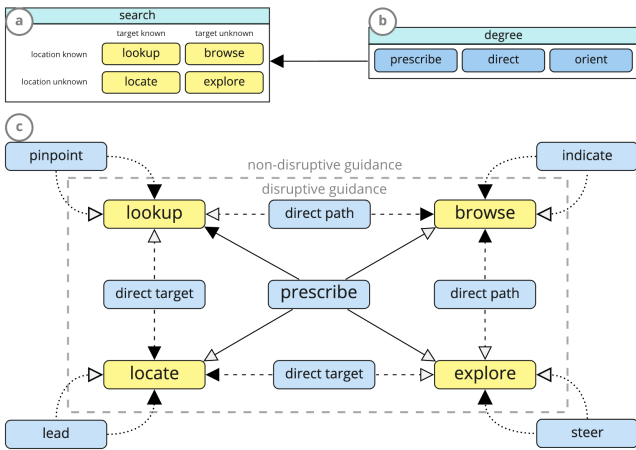


Figure 2: Perspective change dynamics. Users performing a task of a certain search type (a) can be supported by providing guidance of a certain degree (b). Our observation is that guidance can make a search type become a different one (c). White arrows indicate targeted search type and black arrows resulting search type. Directing (dashed line) and Prescribing (solid line) guidance induce perspective change on the user, pushing their current search towards another one, while Orienting (dotted line) does not (bottom).

generated or confirmed (effectively having an influence on the following analytic discourse) or it can be “automatically” transformed into actions, by moving sideways through the Exploration Loop (as it most usually happens with findings). When guidance is provided, usually the next action of the user is either accept or choose one of the suggestions, a one-click decision. Therefore, the guidance information about whether or not it was accepted is fed back to the system almost unchanged. The result is that the Exploration Loop can be short circuited by the guidance system, as the user willingly delegates agency to the guidance. This is illustrated in Fig. 1 by the red path that goes all around the Exploration Loop (information from the guidance system) and the green line that stops at the Action block (information from the user). Thus, guidance can be double-edged: it can make analysis more fluid and enrich it, but it can also disrupt downstream flow (i.e., information-rich interaction from the user), which carries information about the users’ goals, and constrain their decisions. We formalize this effect of guidance in the next section.

4. Perspective Change Dynamics

The model we introduced depicts the normal interaction between users and guidance. However, what happens when the user does not know how to continue the analysis, i.e., there is a knowledge gap? How can guidance close the user’s knowledge gap is an open-ended question. In the following, we show how providing different guidance degrees can affect the way a knowledge gap is solved. We do this by describing how guidance influences the way a user searches for information during the analysis, i.e., the path/target known/unknown dimension that is common to both user search type [BM13] and knowledge gap [CGM*16].

Each analysis can be seen as the efforts users make while searching for specific information. Our assumption is that guidance can

affect users’ tasks by constraining or re-targeting their search space: users are pushed towards a narrower set of options or pulled to consider something different. We call the effect that guidance can have on the user **Perspective Change**, as it occurs in the theoretical transition space where the user’s goals are mapped to actions known as the Perspective Dimension [RAW*16].

Search Types. There are four types of search operations that can be performed by a user during visualization tasks [BM13]. The search type depends on the level of knowledge that users hold about the object of the search and the path to reach it (see yellow boxes in Fig. 2a): if users do not hold any hypothesis yet (if the analysis is, for instance, an open-ended search) they are performing an *explore* task; if users know the characteristics of their target but not where to find it, they are performing a *locate* task; if users know their target is within a set of elements, they are performing a *browse* task; finally, if the users know both what their target is and where to find it, they perform a *lookup*.

Thanks to the model of guidance we introduced earlier, we show how the provision of guidance can change and influence these search tasks. In particular, how they can change and evolve when supported with different guidance degrees (Fig. 2b). Our approach can be imagined as a function that takes a guidance degree and a search type and maps them to a search type resulting from the perspective change incurred:

$$\text{perspectiveChange}(\text{search}_a, \text{degree}) \rightarrow \text{search}_b.$$

Some guidance degrees do not change the search type (when $\text{search}_a \equiv \text{search}_b$), and some degrees change the search type (when $\text{search}_a \neq \text{search}_b$). We call the former *non-disruptive* guidance, and we differentiate them from *disruptive* guidance degrees, in the latter case. Moreover, to differentiate guidance degrees which starting from a search type can lead to more than one search type, we introduce *second-order* degrees ($(\text{degree}, \text{search}_a, \text{search}_b) \rightarrow \text{degree}'$). Second-order degrees are an additional specification of a guidance degree that specify the role that the degree has when applied to a specific search task. The contribution presented here is thus twofold: we produce the mapping function of user search types when supported by guidance and a finer-grained set of guidance degrees (Fig. 2c). In the following, we describe the second-order degrees, grouped by disruptive (deriving from prescribing and directing guidance) and non-disruptive (derived from orienting guidance).


4.1. Disruptive Guidance

Among all guidance degrees, we identify those that are disruptive, i.e., directing and prescribing guidance. Such guidance degrees can have a direct influence on the user’s task. In fact, they can make the user’s task change as a consequence of the Exploration Loop “short-circuit”, which was described in section 3. These higher degrees of guidance are typically used to directly address a user’s knowledge gap. However, since the knowledge gap is elusive in practice, it is hardly the case that there is a perfect match between the provided guidance and the identified gap. This mismatch, in conjunction with the fact that such degrees limit the user’s freedom, could lead to unwanted behaviors and hence to their disruptiveness. Considering this is important for guidance design because it affects the use of the system and could lead to a series of undesired effects, such as

the system’s tools getting trampled by guidance (as the user task for which they were designed may get “overwritten” by disruptive guidance) and a decrease in user interaction (as users simply select the “guide’s pick”) to feed the guidance and inference loops.

We use the notation $degree \rightarrow search_a \Rightarrow search_b$ to say that search type $search_a$, targeted by a guidance degree $degree$, is changed to $search_b$. *Degree* represents the second-order guidance degree that is meant to support the user’s search task.

Prescribing guidance. When “the system establishes a set of mandatory actions, or specifies a step-by-step instruction the user should take to proceed” [CGM19a, p. 3] we say the system is providing prescribing guidance. When actions are prescribed, the user’s freedom is restricted to only accepting or declining the system’s suggestion, thus making any other kind of search impossible at that moment, as the system has decided on both target and path. Thus we say that the user’s search is reduced to a lookup:




Prescribe→Explore/Locate/Browse ⇒ Lookup. Prescribing guidance makes explore, locate, browse tasks become a lookup task. Consequently, there is only one type of prescribing guidance, with no second-order operators. Prescribing guidance is the degree with the highest chance of incurring in information loss, i.e., of nullifying possible analytic paths stemming from the user. We can find prescribing guidance in Horvitz et al. [HBH*13] and Ip and Varshney [IV11].

Directing guidance. Directing guidance can be considered as a “partial prescribing” since it can address only one type of knowledge gap at a time (target or path), thus leaving some options for the user to choose from. When directing guidance is provided, a perspective change happens because the users’ search is pushed towards a browse or locate search task. In fact, the user can locate or browse within the options suggested by the directing guidance, as directing actively reduces the user’s options to an ordered few, e.g., when using a map application, we are provided with a few path options towards our destiny, ranked by trip-duration. However, directing guidance can also be used to *expand* the search space [PSCO09] when the user is performing a lookup. To differentiate between directing guidance provided to constrain the search space from guidance provided to broaden the search we use the terms *converging* and *diverging*, respectively. The former is the most common use of directing guidance, as the purpose of guidance is generally for user and system to arrive together to a solution, reducing at each step the result set. In spite of this, (diverging) guidance can also be used as a technique to make the users consider different options when their search has reached its bottom, e.g., to promote a breadth-first search.


As mentioned, directing guidance can push user tasks to evolve towards browsing or locating. This entails directing guidance can produce two different effects (perspective changes) over the user search. Which of these is the case depends on the knowledge gap to overcome: a *path unknown* or a *target unknown* gap. The former is solved by the guidance suggesting a *direct path*, which pushes user’s search task to become a browse task (see Fig. 2c). The latter is solved by suggesting a *direct target*, which makes the user search become a locate task. This is independent of the converging/diverging nature of the provided guidance, as the type of answer (path/target) is

specific to each resulting user search (locate/browse).

Direct target→Explore/Lookup ⇒ Locate. An *explore* or a *lookup* task when supported with *target-directing* guidance becomes a *locate* task. When enough information about the characteristics of the users’ search is available to the system (e.g., after the user has done some exploration and focused on certain elements/areas), the system can provide *target-directing* guidance. For instance, Wongsuphasawat et al. [WMA*15] present a guidance system to support exploratory data analysis of a relational table that recommends appropriate views and encodings for the data, thus closing the *target unknown* gap and making the user’s task change from an *explore* to a *locate* search.



Direct path→ Explore/Lookup ⇒ Browse. An *explore* or *lookup* task provided with *path-directing* guidance becomes a *browse* task. Path-directing guidance narrows the search space to a few ranked options, i.e., it provides a *direct answer* to the users’ task by automatically choosing a path to possible targets to their search. Examples of this are providing lists of suggestions based on previous users’ searches [DLB13] or automatic layout suggestions in graphic design tools [OAH15].




4.2. Non-disruptive Guidance

Non-disruptive guidance preserves the user search type, as it does not constrain the original freedom of the user within the VA interface. However, orienting guidance provides the user with some new information of assistance, and this information can vary depending on the kind of search type the orienting guidance is designed to support. Because of the freedom it leaves to the user, non-disruptive guidance is mainly associated to this search type. We use the notation $orienting \rightarrow search \Leftrightarrow degree'$ to say that when orienting guidance is applied to support a search task *search*, we call $degree'$ the second-order complement to that task.


Orienting guidance. Orienting guidance has been described as the lowest guidance degree, whose main function is to preserve or enhance the users’ mental map, usually by adding an extra layer of information to the visualization [CGM*16]. This definition points out that orienting is a non-disruptive form of guidance, i.e., the user goals-to-action mapping is not directly affected by the provided guidance. Thus, we observe that when orienting guidance is chosen, the guidance task depends on the search type of the user task. This gives rise to four subtypes of orienting guidance: **pinpoint**, which is the kind of orienting guidance provided for *lookup* tasks; **indicate**, stemming from *browse* tasks; **lead**, associated to *locate*; and **steer**, to support *explore* tasks.

Orient→Lookup ⇔ Pinpoint. Pinpoint is the orienting guidance complement of *lookup* tasks, where target and path are known by the user. In *lookup* tasks, although little guidance may be needed, help can be provided still to pinpoint the information searched and avoid mistakes in the operation. For instance, pointer magnet, word autocomplete, automatic zooming,




slider marks of relevant values, etc., are all examples of pinpointing guidance. Pinpoint does not require complex knowledge from the user and should work for all targets alike.


Orient→**Browse** ⇔ **Indicate**. Indicate is the orienting guidance complement of *browse* tasks, where only the path is known by the user. Browsing is a limited kind of exploration, where usually the search is constrained to a small subset of the data, and so indications should be provided within this context to find the searched information. For instance, Koop et al. [KSC*08] provide indications in the form of previsualizations while users browse a list of automated suggestions of visualization pipelines.



Orient→**Locate** ⇔ **Lead**. Lead is the orienting guidance complement of *locate* tasks, where only the target is known by the user. While locating, users may search different areas of the visualization while inspecting non-spatial features of data, and so leads can be provided about which features may be more interesting. Examples of leading can be found in systems that suggest data dimensions and relations, such as Morpheus [MAK*08], Stratomex [LSS*12] and Domino [GGL*14].



Orient→**Explore** ⇔ **Steer**. Steer is the orienting guidance complement of *explore* tasks, when both path and target are unknown to user. Exploration is a common open-ended phase of analysis that usually leads to clearer goals and to locate and browse searches. Exploring allows users the highest degree of freedom, and orienting guidance has a very important role here, as it is able to steer the users' wandering without constraining their freedom but helping to maintain flow and propitiating findings. For instance, Gladisch et al. [GST13] provide steering guidance for the navigation of large graphs through glyphs embedded on the exploratory view.



5. Typology of System Guidance Tasks

We have shown how the provided guidance degree interacts with the users' search. To tie this in with abstract user visualization tasks, in this section we introduce the notion of guidance tasks.

We propose a typology of system guidance tasks to describe and analyze guidance systems in VA and their interaction with users and their tasks. With this last objective in mind, we designed our typology to be compatible with the multi-level visualization task typology [BM13], but it also draws elements from previous work on guidance [CGM*16; CGM*18]. Our typology – portrayed in Fig. 3 – spans the three main dimensions of user tasks: *Why* (describing the intent of the guidance), *How* (showing how an intent is translated into actions) and *What* (input/output). An extra dimension is added as context: the *When*, capturing the analytic objective that frames the user-guidance task complex.

5.1. When

When is the dimension that represents the analytical context during which the task takes place and thus affects both user and system tasks. Its only category, **Objective**, was introduced as *analytical objective* by Ceneda et al. [CGM19a] to categorize guidance according to which phase of analysis it supported. It is divided into: *Data*

Transformation, *Visual Mapping*, *Parameter Setting*, *Model Visualization*, *Model Building*, *Exploration* and *Knowledge Generation* (for a detailed description, see Ceneda et al. [CGM19a]). Users and system must share the same *when* (e.g., transforming the data or performing data exploration) for the provided guidance to be effective.

5.2. Why

The second dimension we examine when analyzing guidance tasks is the *why* which considers the intentions of the system when it decides to support user's tasks. The *why* dimension can be considered as a reflection of the *intent* with which the system was designed, which is manifested in the high-level activities it performs, i.e., the guidance tasks. The *why* goes from high-level (process vs. explain) to mid-level (guidance degree) to low-level (second-order degrees).

Guide. The main function of guidance is to provide active support to users. It achieves this with the four processes described in the Guided Knowledge Generation Model (section 3): *provide*, *observe*, *expect*, *adapt*, corresponding to the four blocks in the Guide Side (Provision, Observation, Expectation, Adaptation; Fig. 1). Although all of these processes have influence on the guidance provided at a given time, we focus this typology only on the *provide* function, i.e., the part of the loop that produces the concrete actions taken to support the user's analysis. As described, we imagine the intent of the system (ideally) as a mirror of the user's intent, and the *why* dimension of the visualization task typology [BM13]. In other words, the role of the guidance, as we have seen, is supporting search tasks.

Explain. A guidance system can explain its decisions to the user or prompt the user for feedback. *Explain* is a different function from support and not a main element in the Guided Knowledge Generation model, since the information it communicates is not about the data being analyzed, but about the guidance itself. Explain tasks, in other words, serve as a complement to the guidance tasks and concur to their effectiveness. In Fig. 3 we can see how the Explain block is connected to the whole Guide block since an *explain* task can take place at any moment, i.e., either during the main Guidance Loop or the Inference Loop. Specifically, this dimension considers the case in which the system *informs* the user about its current internal state, the analysis state, or other relevant information about the guidance and inference process. If the user has the possibility to ask, the system task can also be to *clarify* the system's knowledge to the user and explicit guidance decisions, which otherwise would remain hidden, not to disturb the analytical discourse. Also, when the system's intent is to *ask* the user, or prompt it for feedback. *Correct* is a modality of explaining where the system provides negative feedback to the users, i.e., the system signals the users that they have at some point taken a path considered misleading.

Degree. Depending on the number of targets or paths returned to users by the guidance function, the guidance task is said to be one of three degrees (previously defined in section 4) [CGM*16], which correlate with the *query* dimension of user tasks [BM13] in the following way: *prescribing* gives users only one option to follow next, and is meant to complement users' *identify* task; *directing* returns two or more ranked suggestion, in response to a *compare* task; and finally *orienting* returns multiple points (e.g., by highlighting or aggregating into a new encoding), as to support a *summarize* task, on the user side.

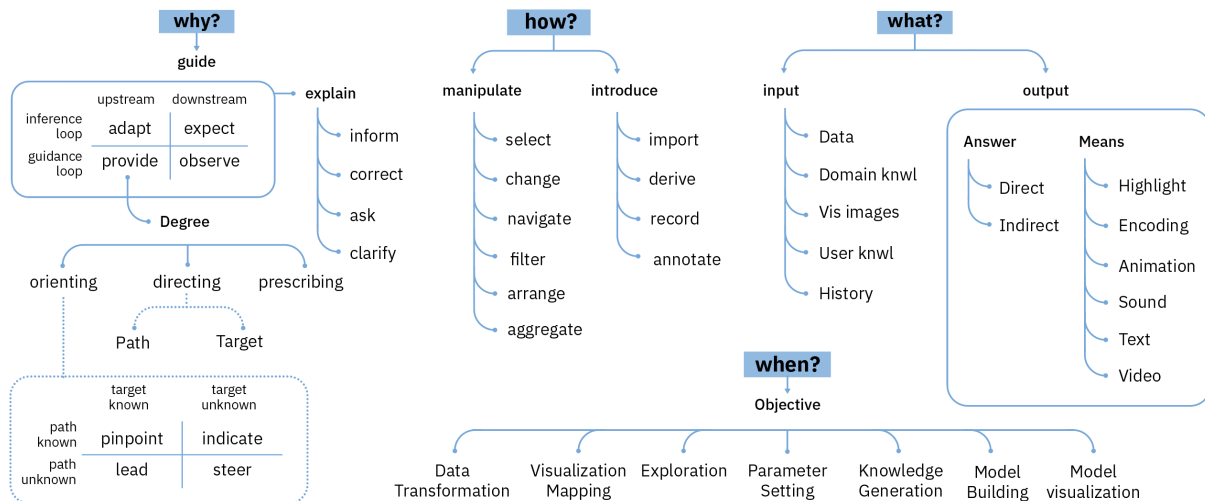


Figure 3: Typology of system guidance tasks. It spans the three dimensions of the multi-level visualization task typology [BM13] plus a new dimension that captures the analytical objective of an analysis phase (when?). It allows describing: system task intent (why?) by different detail levels (aim, first- and second-order degree), also with an accompanying explanation task (explain); the suggestion method (how?) in terms of data manipulations and means of communication; and the information inputs and type of output relative to the targeted user task (what?).

Second-order degrees. *Orienting* and *directing* guidance can have different modalities depending on the user task they target. As this dimension depends on both guidance *degree* and the type of *search* the user is performing, they are linked by a dotted line in Fig. 3. *Orienting* can be further structured in four second-order degrees, which directly depend on the user task: *pinpoint*, *indicate*, *lead*, and *steer*. *Directing* has two second-orders: *path* and *target*. The full description of second-order degrees is found in section 4.

5.3. How

The *how* part of our guidance typology refers to the data interaction methods. Like a user, the guide can perform data *manipulation* tasks, as these are non-persistent, but it can also *introduce* persistent data (e.g., from web sources).

Manipulate and Introduce. As guidance provides suggestions for the user *to do*, guidance tasks are equivalent to user tasks in the dimension concerned with formal data operations that take place within visual analytical discourse. In other words, since the analytical discourse takes place through the same medium (a screen and a visualization), the way the guidance is provided to the user matches the way the user can interact with it. Thus, we keep *manipulate* and *introduce* as the same categories used to describe user tasks. For a description of available actions, see Brehmer and Munzner [BM13].

5.4. What

The *what* dimension is divided into input and output. It encapsulates the origin source of the information used to provide guidance and what is communicated to the user.

Input. Input is the (static) information sources from which the guidance system obtains its parameters for the guidance task. It is different from information coming from an *observation* related to a task, as observations are derived from real-time user interaction (downstream flow). For a detailed description of possible input sources, see Ceneda et al. [CGM*16].

Output. The Output is what guidance directly passes on to the user, and it is divided in *answer*, the solution to the knowledge gap, and *means*, through which media and channels it is conveyed. **Answer**, as described by Ceneda et al. [CGM*16], is the content of the guidance solution to the identified knowledge gap. It is classified by its relation to the knowledge gap: a *direct* answer is such that it contains a possible solution state of the user task (i.e., it is a target for the task), while an *indirect* answer is of a different nature than a solution state of the user task (i.e., it is a path for the task). **Means** are the visual (or otherwise) channels that will convey the answer to the user, as any *how* action taken by the guidance system must be communicated somehow to the user, and it could potentially be conveyed in different ways. It can turn the users' attention by *highlighting* visual variables; produce a new *encoding* of its own; use *animation* through the continuous motion of visual variables; use the *sound* channel to convey information; use *text* to deliver a message in natural language; and *video* to illustrate with a prerecorded sequence. For a detailed description of means, see Ceneda et al. [CGM*16].

6. Case Studies

To show the practical utility of our typology, we analyze in this section three different mixed-initiative VA systems found in the literature: Topic-Tree [ESD*18], ForceSPIRE [EFN12], and Design-Scape [OAH15]. We have chosen this particular set of systems as they provide us with a certain diversity of guidance behaviors (see The Guidance Spectrum, section 7). We first explain how the typology can be used to generate a user-guidance task decomposition.

How to use our Typology. Our guidance task typology was designed to be used in conjunction with a user task typology, in particular with the multi-level abstract visualization task typology [BM13], which was also designed for describing complex interconnected task sequences. This interconnection between typologies allows to succinctly describe guidance-enriched VA systems in a modular way.

The description of guidance tasks must be preceded by a descrip-

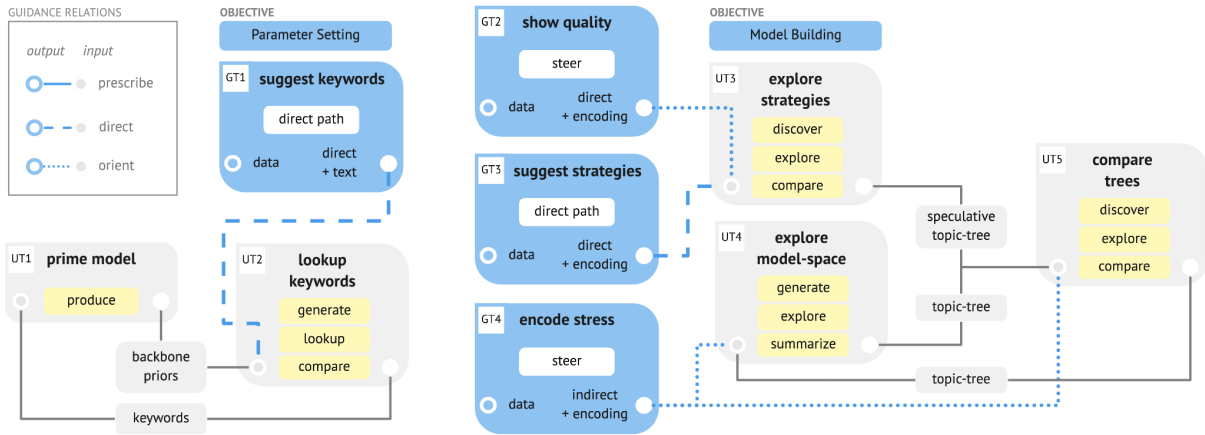


Figure 4: UT and GT decomposition of Topic-Tree [ESD*18]. The decomposition is split into two objectives (when), as the system uses guidance to support parameter setting (left connected component) and model building (right connected component).

tion of the user tasks within the VA system which is designed to support them. This method assumes that a VA system can be thought of as independent of the guidance solution and hints that tasks can be, thanks to their modular nature, recombined to produce varying guidance behaviors.

Once user tasks (UTs) and guidance tasks (GTs) are described, each task should be associated and target one or more UTs. Targeting means establishing one of the following relationships among tasks, which derive from the Guided VA Model (Fig. 1): *observing*, *providing*, or *co-adaptive* (which includes both). While *observing* the targeting GT is meant only to passively receive downstream information, but not to provide explicit guidance (UT output to GT input). A *providing* relation means that the targeting GT is meant only to provide guidance, but not to observe downstream information (GT output to UT input). Finally, a *co-adaptive* relation among tasks has both *observing* and *providing* links, and thus is guidance that is sensitive to the real-time information flow of the input UT (using the full guidance loop in Fig. 1).

The construction of user-guidance task descriptions can be summarized in the following rules, accompanied by their rationale:

1. UTs and GTs are first described independently of one another: the VA system and the guidance system are, at least abstractly, separable entities. The former affords UTs, the latter supports them.
2. UTs and GTs are assigned to analytical objectives: GTs can only target UTs within the same objective. This is important to consider because of the situational nature of tasks.
3. Each GT must target at least one UT: The purpose of guidance is to support the user, and so the idea of guidance as an analytical task in and by itself makes no sense. A GT must always be coupled to a UT, namely, its target.
4. A targeting relation can be *observing*, *providing*, or *co-adaptive*, indicating the direction of the information flow: A GT may feed (have an *observing* relation) from one UT while targeting (*providing*) to another one, or have both (*co-adaptive*).
5. *Produce* tasks may only be targeted by *observing* relations: *Produce* tasks [BM13] are common in VA, as analyses usually deal with creating or refining a certain artifact (e.g., a model). Unlike *consume* tasks, this kind of tasks does not possess a search type.

Guidance, in our model, is only understood through its effect on perspective (section 4), thus, it has no effect on tasks without a search type. A GT may get its input from a *produce* task and use it to provide guidance to a different UT.

In the visual diagrams presented next, UTs and GTs are encoded in gray and blue boxes, respectively. The same colors are used for lines denoting UT dependencies and guidance relations. The *why* for both kinds of tasks is specified, while the *how* has been ignored for clarity as it does not fulfill a structural role in the decomposition. The *what* is specified next to the input/output symbols for GTs (see Fig. 4 Guidance Relations) and on the dependency lines for UTs.

6.1. Case Study 1: Topic-Tree

El-Assady et al. present a mixed-initiative VA system for document collection segmentation [ESD*18]. In this approach, a topic model is presented to the user as a topic-tree, which supports a comparison of two models. In the main interaction loop, the system adds documents as leaves of the topic-tree one by one, so the user can follow and stop the hierarchical clustering algorithm at any time to manually change its decisions or instruct it to follow a different strategy. Several side panels inform the user on the details of documents, the overall quality of the model, and of the possible strategies to adopt, which the user can preview. It uses *speculative execution* (SpecEx) as a guidance mechanism (described abstractly by Sperrle et al. [SBS*18]). We describe the general workflow of this feature-rich system in terms of interrelated UTs and GTs, as shown in Fig. 4.

Before entering the topic-tree view, users can prime the model with their domain knowledge by *setting model priors* (Fig. 4 UT1): introducing representative keywords for each topic. Users are assisted in their search for keywords (Fig. 4 UT2) by an automatically ranked keyword suggestion list extracted from the data (Fig. 4 GT1). GT1 provides *directing* guidance to UT2 (an ordered list of possible words) so that the users can browse and select between words that they may not have contemplated before (diverging guidance) and so enrich the produced backbone priors. These tasks pertain to the Parameter Setting objective and thus are independent of the rest of the task decomposition.




Figure 5: UT and GT decomposition of the ForceSPIRE main interaction loop [EFN12].

Following the Parameter Setting, Model Building is the core objective of the system. Three main UTs are identified revolving around the refinement of a topic-tree: *explore strategies* (Fig. 4 UT3), where users can use speculative execution to observe the consequences of predetermined actions, *explore model-space* (Fig. 4 UT4), where the topic-tree is visualized and users can directly manipulate the model, and *compare trees* (Fig. 4 UT5), where they can compare a speculative tree to the current tree in a differential topic-tree view. These tasks depend on each other forming a loop (addressing the iterative, human-in-the-loop approach of the system). Then, we identify three GTs: *show quality* (Fig. 4 GT2), which encodes the quality metrics to orient the user about the progress of the model and of speculative models; *suggest strategy* (Fig. 4 GT3), which proposes ranked actions to undertake on a parallel model; and *encode stress* (Fig. 4 GT4), which shows the quality of each topic cluster, indirectly leading to certain paths and targets.

Observations. We can see that the SpecEx employs directing guidance (GT3) to constrain the users' exploration, the paths they can take through the model-space, to a browsing task. In this system's quantitative evaluation, users were asked to rank variations of the system based on their subjective appreciation of the results: (1) as it is, (2) with GT3 turned into prescriptive guidance (i.e., automatically selecting the best-ranked strategy) and (3) without GT3. They consistently chose directing guidance (1) as giving the best results, followed by no guidance at all (3) and lastly by prescriptive guidance (2), even though the calculated quality metrics on the system indicate that automatic SpecEx (2) yields the better results. Thanks to the abstraction made possible by our typology, this could point out to a generalizable result about guided interaction.

6.2. Case Study 2: ForceSPIRE

ForceSPIRE is a VA system for the refining of topic models in collection of documents presented by Endert et al. [EFN12]. It is based on a particular interaction technique called *spatialization*, which is a kind of *semantic interaction*, in turn an idea derived from the "human-is-the-loop" VA paradigm [EHR*14]. Spatialization allows the users to be "shielded from direct model steering" [EFN12], in the sense that they can iteratively refine a model by an intuitive interface that uses spatial interaction as its main language: users can freely move documents around and see the others react to this movement, revealing their underlying relations (an *exploratory move*), or they can force two documents closer together to signal the system their common entities should increase in weight, thus updating the relations in the whole corpus (an *expressive move*).

From a visualization tasks perspective, we can see that the users are involved in the *production* of a new model with each expressive move, and they do not do this directly but by rearranging the spatialization (Fig. 5 UT1). Users can then explore the result of the spatialization (Fig. 5 UT2). This analysis is still incomplete as it misses the agency of the guide that generates the reaction to the spatial interactions with the model and thus *steers* the users towards interesting features of the current model (Fig. 5 GT1). The guidance thus *observes* the output of the *produce* task and targets the *explore* task with a guidance degree. Considering this, UT2 actually corresponds with the *exploratory move* described in the system, where users learn about the model by observing the simulated reactions of the system to their actions that comes from the provided steering guidance in the form of the *animated rearrangement* of documents. GT1 feeds on both the model and the downstream user actions, and provides a direct answer to the user knowledge gaps "how do other documents relate to this particular document (i.e., their transitive relationships)?" and "how do these (transitive) relations change if I change the model?".

Observations. Both the iterative and mixed-initiative nature of the analytic discourse within ForceSPIRE, as seen in Fig. 5, are captured by our representation: the flow of information of the user-guidance task complex forms a loop where there is not an order of tasks imposed on users (they are free to start at any UT and to change tasks whenever they please). We can explain that the users of the conducted study "treated their investigation not as steering a model, but rather synthesizing information" [EFN12] through the fact that the steering is actually delegated to the *orienting* GT, so the users can focus on their upstream (synthesizing) flow.

6.3. Case Study 3: DesignScope

DesignScope is a non-VA tool which nonetheless introduces guidance components into its workflow [OAH15]. Its aim is to support users on the task of creating posters or graphical layouts in general. It has a similar interface to end-user programs such as Adobe InDesign, however, it adds two guided modes of use called the "suggestive interface" and the "adaptive interface". The former provides predominantly *prescribing* guidance while the latter *directing* guidance, although the algorithms behind both are the same.

To analyze this system, we first decompose the procedure of creating a graphical layout with a visual interface without guidance (Fig. 6 UT1-3). It can be described in the following way: the users *introduce elements* as images and text (Fig. 6 UT1), then *explore* possible layouts within a large "design space" (Fig. 6 UT2), to finally reach a more or less satisfying state where they will *fine-tune* the positions (lookup precise alignments, margins, etc.; Fig. 6 UT3).

UT3 is commonly supported in current editors by automatically snapping to grid the objects according to the bounding boxes of surrounding elements: this is a kind of pinpoint guidance, as a *lookup* is being targeted by *orienting* guidance (this is the only guidance offered in the baseline "direct interface", Fig. 6 GT4). The "suggestive interface" adds two panels to the normal interface where two kinds of layout proposals are presented (Fig. 6 GT2-3): design *tweaks* that improve the layout by making only small changes (which are reactive to the user's movements), and design *brainstorms*, where only possibilities that are far in the design space are shown (nonreactive to the

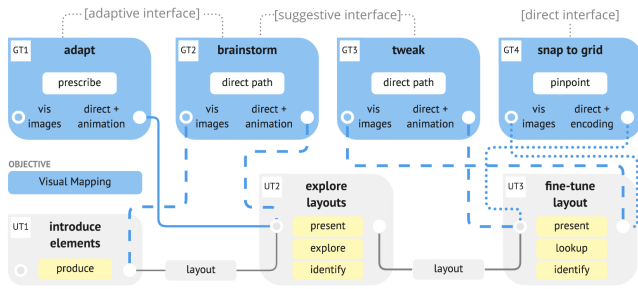


Figure 6: UT and GT decomposition of DesignScape [OAH15]. The three types of interface present in the system have each some guidance tasks associated.

user's movements, but only to introduced elements). These two GTs target UT3 and UT2, correspondingly. Both, however, offer *directing* guidance by signaling a few (shorter or longer) paths in the design space to reach a myriad of target states. Finally, the “adaptive interface” (Fig. 6 GT1-2) also offers the side-panel for *brainstorm* suggestions, but instead of *directing* to fine-tuning solutions, it automatically enacts them without asking the user (sometimes with an important effect on the current layout), thus *prescribing* guidance (GT1).

Observations. Being this a tool aimed at supporting a creative task, it is understandable from this analysis that in the reported results of O'Donovan et al. [OAH15], users felt in the end frustrated by their lack of agency in solutions they found nonetheless good (better than what they thought they could achieve by themselves), as, according to the dynamics of the user-guidance complex, their exploration is reduced in the “adaptive interface” to a *lookup* and in the “suggestive interface” to a *browsing*.

7. Discussion

Our framework shows that guidance works in more varied ways than previously known. Its core value lies in providing a structured but expressive view of mixed-initiative approaches by considering both sides of the VA spectrum.

The Guidance Spectrum. Another relevant aspect that arises from our model is the difference between upstream (system-provided) and downstream (user-provided) guidance. The observation that in some systems, some types of guidance may appear more prevalent, or not appear at all, leads us to imagine of a two-dimensional guidance spectrum where systems can be placed. Each axis of this spectrum is the quantity or importance of each kind of guidance provided, and we can produce a phase space with four kinds of systems: *weak or no guidance* systems, *proactive* guidance systems (mostly providing “static” suggestions, as in Fig. 4), *passive* systems (mostly learning from the user but providing weak guidance features, as in Fig. 5), and *co-adaptive* systems (having both directions strongly represented, as in Fig. 6). This coarse-grained clustering of guided approaches leads us to ask *how do users respond to each kind of system, considering that the response of the system is different in each case?* We imagine that in co-adaptive systems, complex behaviors should be expected as users learn that they can tune the provided guidance by modifying *their* own behavior. If such a co-adaptive system also possesses an inference loop, this

would mean that the system could also respond dynamically to this kind of user behavior (e.g., by observing and adapting to user interaction while trying different local optima in a “guidance space”). We can think of an emergent “agonistic” behavior, where each agent plays its moves in consideration of eliciting an expected response from the other (as in partisan games).

Limitations. Our work can easily fall prey to its own reductionist approach. Its fundamental assumption, that guidance can be encapsulated as a phenomenon perfectly independent of the classical VA system, is certainly debatable (e.g., by asking, where do we draw the line between visualization features of a system and providing orienting guidance) and can be held as a reification of guidance [KK17]. Inheriting from Brehmer and Munzner's typology, our task decomposition schema is certain to produce different results, as at least some derivation variables depend on the analysts' tailoring (including: level-of-detail, which can even bring one task to be represented as many interdependent tasks with their own iteration loops; presumed user knowledge, which can change the search type of a UT; and presumed user intentionality, which can make the *when* structure of the task decomposition generate different connected components). The idea that variations in the decomposition can have an effect over the following analysis also calls the validity of our observations into question. Our case studies are limited but to some extent diverse, suggesting the applicability of our method to extant systems. Furthermore, appropriate guidelines would be beneficial to develop new approaches, which we are planing as future work.

8. Conclusion

We have presented a typology of system guidance tasks that enables a joint analysis of user and guidance task interdependence, illustrating it with different examples. We have also shown the effects of guidance over user tasks and vice versa, deriving finer-grained guidance degrees in the process. This is supported with a model of guidance within the VA knowledge generation process. Our typology appears to serve well the purpose of describing, abstracting, and generalizing VA systems with mixed-initiative approaches, providing succinct representations that we hope will enrich the incursion into guidance design and improve the communication of results, stimulating the production of guidelines that, with time and testing, may expand design considerations in VA.

Acknowledgements. This work was funded by Vienna Science and Technology Fund (WWTF) under grant ICT19-047, by the Austrian Science Fund (FWF) with the grant P31419-N31 (KnoVA), by the Deutsche Forschungsgemeinschaft (DFG) within the project Knowledge Generation in Visual Analytics (Project-ID 350399414), by DoRIAH (FFG Grant #880883) and by the ETH AI Center.

References

- [AES05] AMAR, R., EAGAN, J., and STASKO, J. “Low-level components of analytic activity in information visualization”. *IEEE Symposium on Information Visualization*. IEEE, 2005, 111–117 2.
- [ALA*18] ANDRIENKO, N., LAMMARSCH, T., ANDRIENKO, G., FUCHS, G., KEIM, D., MIKSCH, S., and RIND, A. “Viewing visual analytics as model building”. *Computer Graphics Forum*. Vol. 37. 6. Wiley Online Library, 2018, 275–299 3.
- [APS13] AHN, J.-W., PLAISANT, C., and SHNEIDERMAN, B. “A task taxonomy for network evolution analysis”. *Transactions in Visualization and Computer Graphics* 20.3 (2013), 365–376 2.
- [BH19] BATTLE, L. and HEER, J. “Characterizing exploratory visual analysis: A literature review and evaluation of analytic provenance in Tableau”. *Computer Graphics Forum*. Vol. 38. 3. Wiley Online Library, 2019, 145–159 2.
- [BM13] BREHMER, M. and MUNZNER, T. “A multi-level typology of abstract visualization tasks”. *Transactions in Visualization and Computer Graphics* 19.12 (2013), 2376–2385 1, 2, 4, 6–8.
- [BWD*19] BORS, C., WENSKOVITCH, J., DOWLING, M., ATTFIELD, S., BATTLE, L., ENDERT, A., KULYK, O., and LARAMEE, R. S. “A provenance task abstraction framework”. *IEEE Computer Graphics and Applications* 39.6 (2019), 46–60 2.
- [CAA*20] CENEDA, D., ANDRIENKO, N., ANDRIENKO, G., GSCHWANDTNER, T., MIKSCH, S., PICCOLOTTO, N., SCHRECK, T., STREIT, M., SUSCHNIGG, J., and TOMINSKI, C. “Guide me in analysis: A framework for guidance designers”. *Computer Graphics Forum*. Vol. 39. 6. Wiley Online Library, 2020, 269–288 2.
- [CAS*18] COLLINS, C., ANDRIENKO, N., SCHRECK, T., YANG, J., CHOO, J., ENGELKE, U., JENA, A., and DWYER, T. “Guidance in the human-machine analytics process”. *Visual Informatics* 2.3 (2018), 166–180 1, 2.
- [CCI*15] COOK, K., CRAMER, N., ISRAEL, D., WOLVERTON, M., BRUCE, J., BURTNER, R., and ENDERT, A. “Mixed-initiative visual analytics using task-driven recommendations”. *2015 IEEE Conference on Visual Analytics Science and Technology (VAST)*. IEEE, 2015, 9–16 1, 2.
- [CGM*16] CENEDA, D., GSCHWANDTNER, T., MAY, T., MIKSCH, S., SCHULZ, H.-J., STREIT, M., and TOMINSKI, C. “Characterizing guidance in visual analytics”. *Transactions in Visualization and Computer Graphics* 23.1 (2016), 111–120 1, 2, 4–7.
- [CGM*18] CENEDA, D., GSCHWANDTNER, T., MAY, T., MIKSCH, S., STREIT, M., and TOMINSKI, C. “Guidance or No Guidance? A Decision Tree Can Help.” *EuroVA@ EuroVis*. 2018, 19–23 2, 6.
- [CGM19a] CENEDA, D., GSCHWANDTNER, T., and MIKSCH, S. “A review of guidance approaches in visual data analysis: A multifocal perspective”. *Computer Graphics Forum*. Vol. 38. 3. Wiley Online Library, 2019, 861–879 1–3, 5, 6.
- [CGM19b] CENEDA, D., GSCHWANDTNER, T., and 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 2.
- [Coo21] COOK, D. *The Chemistry of Game Design*. <https://lostgarden.home.blog/2021/03/13/the-chemistry-of-game-design-2/>. 2021 2.
- [DLB13] DÖRK, M., LAM, H., and BENJELLOUN, O. “Accentuating visualization parameters to guide exploration”. *CHI’13 Extended Abstracts on Human Factors in Computing Systems*. 2013, 1755–1760 5.
- [DS21] DIMARA, E. and STASKO, J. “A Critical Reflection on Visualization Research: Where Do Decision Making Tasks Hide?”. *Transactions in Visualization and Computer Graphics* (2021) 1, 2.
- [EFN12] ENDERT, A., FIAUX, P., and NORTH, C. “Semantic interaction for visual text analytics”. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. 2012, 473–482 3, 7, 9.
- [EHR*14] ENDERT, A., HOSSAIN, M. S., RAMAKRISHNAN, N., NORTH, C., FIAUX, P., and ANDREWS, C. “The human is the loop: new directions for visual analytics”. *Journal of Intelligent Information systems* 43.3 (2014), 411–435 9.
- [ESD*18] EL-ASSADY, M., SPERRLE, F., DEUSSEN, O., KEIM, D., and COLLINS, C. “Visual analytics for topic model optimization based on user-steerable speculative execution”. *Transactions in Visualization and Computer Graphics* 25.1 (2018), 374–384 7, 8.
- [FWR*17] FEDERICO, P., WAGNER, M., RIND, A., AMOR-AMORÓS, A., MIKSCH, S., and AIGNER, W. “The role of explicit knowledge: A conceptual model of knowledge-assisted visual analytics”. *2017 IEEE Conference on Visual Analytics Science and Technology (VAST)*. IEEE, 2017, 92–103 2.
- [GGL*14] GRATZL, S., GEHLENBORG, N., LEX, A., PFISTER, H., and STREIT, M. “Domino: Extracting, comparing, and manipulating subsets across multiple tabular datasets”. *Transactions in Visualization and Computer Graphics* 20.12 (2014), 2023–2032 6.
- [GST13] GLADISCH, S., SCHUMANN, H., and TOMINSKI, C. “Navigation recommendations for exploring hierarchical graphs”. *International Symposium on Visual Computing*. Springer, 2013, 36–47 6.
- [GZ09] GOTZ, D. and ZHOU, M. X. “Characterizing users’ visual analytic activity for insight provenance”. *Information Visualization* 8.1 (2009), 42–55 2.
- [HBH*13] HORVITZ, E. J., BREESE, J. S., HECKERMAN, D., HOVEL, D., and ROMMELSE, K. “The Lumiere project: Bayesian user modeling for inferring the goals and needs of software users”. *arXiv preprint arXiv:1301.7385* (2013) 5.
- [HCD17] HORN, B., COOPER, S., and DETERDING, S. “Adapting cognitive task analysis to elicit the skill chain of a game”. *Proceedings of the Annual Symposium on Computer-Human Interaction in Play*. 2017, 277–289 2.
- [IV11] IP, C. Y. and VARSHNEY, A. “Saliency-assisted navigation of very large landscape images”. *Transactions in Visualization and Computer Graphics* 17.12 (2011), 1737–1746 3, 5.
- [KAF*08] KEIM, D., ANDRIENKO, G., FEKETE, J.-D., GÖRG, C., KOHLHAMMER, J., and MELANÇON, G. “Visual analytics: Definition, process, and challenges”. *Information Visualization*. Springer, 2008, 154–175 2.
- [KK17] KERRACHER, N. and KENNEDY, J. “Constructing and evaluating visualisation task classifications: Process and considerations”. *Computer Graphics Forum*. Vol. 36. 3. Wiley Online Library, 2017, 47–59 10.
- [KKEM10] KEIM, D., KOHLHAMMER, J., ELLIS, G., and MANSMANN, F., eds. *Mastering the information age : solving problems with visual analytics*. Goslar : Eurographics Association, 2010 1, 2.
- [KSC*08] KOOP, D., SCHEIDEGGER, C. E., CALLAHAN, S. P., FREIRE, J., and SILVA, C. T. “Viscomplete: Automating suggestions for visualization pipelines”. *Transactions in Visualization and Computer Graphics* 14.6 (2008), 1691–1698 6.
- [LSS*12] LEX, A., STREIT, M., SCHULZ, H.-J., PARTL, C., SCHMALSTIEG, D., PARK, P. J., and GEHLENBORG, N. “StratomeX: visual analysis of large-scale heterogeneous genomics data for cancer subtype characterization”. *Computer Graphics Forum*. Vol. 31. 3pt3. Wiley Online Library, 2012, 1175–1184 6.
- [MA14] MIKSCH, S. and AIGNER, W. “A matter of time: Applying a data-users-tasks design triangle to visual analytics of time-oriented data”. *Computers & Graphics* 38 (2014), 286–290 2.
- [MAK*08] MÜLLER, E., ASSENT, I., KRIEGER, R., JANSEN, T., and SEIDL, T. “Morpheus: interactive exploration of subspace clustering”. *Proceedings of the 14th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. 2008, 1089–1092 6.
- [MMF17] MURRAY, P., MCGEE, F., and FORBES, A. G. “A taxonomy of visualization tasks for the analysis of biological pathway data”. *BMC Bioinformatics* 18.2 (2017), 1–13 2.

- [NHG19] NUSRAT, S., HARBIG, T., and GEHLENBORG, N. “Tasks, techniques, and tools for genomic data visualization”. *Computer Graphics Forum*. Vol. 38. 3. Wiley Online Library. 2019, 781–805 [2](#).
- [Nor86] NORMAN, D. A. “Cognitive engineering”. *User Centered System Design* 31 (1986), 61 [2](#).
- [OAH15] O’DONOVAN, P., AGARWALA, A., and HERTZMANN, A. “Designscape: Design with interactive layout suggestions”. *Proceedings of the 33rd annual ACM conference on human factors in computing systems*. 2015, 1221–1224 [5](#), [7](#), [9](#), [10](#).
- [PC05] PIROLLI, P. and CARD, S. “The sensemaking process and leverage points for analyst technology as identified through cognitive task analysis”. *Proceedings of International Conference on Intelligence Analysis*. Vol. 5. McLean, VA, USA. 2005, 2–4 [2](#).
- [PSCO09] PIKE, W. A., STASKO, J., CHANG, R., and O’CONNELL, T. A. “The science of interaction”. *Information Visualization* 8.4 (2009), 263–274 [3](#), [5](#).
- [RAW*16] RIND, A., AIGNER, W., WAGNER, M., MIKSCH, S., and LAMMARSCH, T. “Task cube: A three-dimensional conceptual space of user tasks in visualization design and evaluation”. *Information Visualization* 15.4 (2016), 288–300 [2](#), [4](#).
- [SBS*18] SPERRLE, F., BERNARD, J., SEDLMAIR, M., KEIM, D. A., and EL-ASSADY, M. “Speculative Execution for Guided Visual Analytics”. *Workshop for Machine Learning from User Interaction for Visualization and Analytics at IEEE VIS*. 2018, 1–8 [2](#), [8](#).
- [SJB*20] SPERRLE, F., JEITLER, A., BERNARD, J., KEIM, D. A., and EL-ASSADY, M. “Learning and Teaching in Co-Adaptive Guidance for Mixed-Initiative Visual Analytics”. *EuroVis Workshop on Visual Analytics (EuroVA)*. Ed. by TURKAY, C. and VROTSOU, K. The Eurographics Association, 2020, 61–65. DOI: [10.2312/eurova.20201088](https://doi.org/10.2312/eurova.20201088) [1](#), [2](#).
- [SJB*21] SPERRLE, F., JEITLER, A., BERNARD, J., KEIM, D., and EL-ASSADY, M. “Co-adaptive visual data analysis and guidance processes”. *Computers & Graphics* 100 (2021), 93–105. DOI: <https://doi.org/10.1016/j.cag.2021.06.016> [1–3](#).
- [SSK*15] SACHA, D., SENARATNE, H., KWON, B. C., ELLIS, G., and KEIM, D. A. “The role of uncertainty, awareness, and trust in visual analytics”. *Transactions in Visualization and Computer Graphics* 22.1 (2015), 240–249 [3](#).
- [SSMT13] SCHULZ, H.-J., STREIT, M., MAY, T., and TOMINSKI, C. “Towards a characterization of guidance in visualization”. *Poster at IEEE Conference on Information Visualization (InfoVis)*. 2013 [2](#).
- [SSS*14] SACHA, D., STOFFEL, A., STOFFEL, F., KWON, B. C., ELLIS, G., and KEIM, D. A. “Knowledge generation model for visual analytics”. *Transactions in Visualization and Computer Graphics* 20.12 (2014), 1604–1613 [1–3](#).
- [WMA*15] WONGSUPHASAWAT, K., MORITZ, D., ANAND, A., MACKINLAY, J., HOWE, B., and HEER, J. “Voyager: Exploratory analysis via faceted browsing of visualization recommendations”. *Transactions in Visualization and Computer Graphics* 22.1 (2015), 649–658 [5](#).