Interactivity in Visual Analytics: Use of Conceptual Frameworks to Support Human-Centered Design of a Decision-Support Tool

Paul Parsons
Western University
pparsons@uwo.ca

Kamran Sedig
Western University
sedig@uwo.ca

Arman Didandeh
Western University
arman.didandeh@uwo.ca

Arash Khosravi
Western University
akhosra3@uwo.ca

Abstract

Visual analytics (VA) combines the strengths of humans and computers such that joint cognitive systems are formed. To be effective, a VA tool should be designed such that the component parts of the whole system are strongly coupled and function in a harmonious fashion. These components include cognitive and perceptual issues, tasks, algorithms, data models, and other aspects of the systems that contribute to its overall efficacy. The quality of interaction among all of these components can be referred to as interactivity. In the existing visualization literature, not enough focus has been placed on developing our understanding of human-centered aspects of interactivity. We have recently developed some conceptual frameworks to inform and guide the design of visual analytics tools in a systematic, human-centered fashion. In this paper, we describe the design of a tool that supports decision-making and other complex cognitive activities. We discuss how the conceptual frameworks supported systematic design and coherent thinking about the interactivity of the tool. We also discuss some extensions of interactivity into important areas of concern for visual analytics tools.

1. Introduction

Visual Analytics (VA) tools combine the strengths of computers with those of humans to create joint cognitive systems. For VA tools to work well, however, there must be a strong coupling and harmonious functioning among all components [1], [2]. Such components include, among others, human characteristics (e.g., cognitive, perceptual), tasks, visual representations (henceforth 'VRs'), and algorithms and data models. One manner in which we can view the strength of the coupling in any VA system is through the lens of interaction—in particular, by examining the *quality of interaction* between the various components. This quality of interaction can accurately be referred to as 'interactivity' [2]. Because the design of interactivity is second-order design—in other words, we cannot design the quality of interaction directly—it is imperative that designers of VA systems have forms of support that enable systematic and principled design of this aspect of VA systems. For the past few years, we have been developing some comprehensive conceptual frameworks that can support critical, systematic design thinking about interactivity of visualization tools [2]–[5]. In this paper, we will provide a brief description and synthesis of some of this work. Furthermore, we will present a VA tool named VARSITY (Visual Analytics of university Research networkS and IndusTrY collaborations) that has been developed to support decision-making and other complex activities, and will demonstrate how the previously developed conceptual frameworks were used in the design process. We will also discuss the need for more attention on human issues in VA (i.e., more human-centered research), design issues, and future directions for extending interactivity research into specific issues with which VA research is currently concerned.

2. Human-centered visual analytics

Many researchers in recent years have suggested that more focus should be put on human issues in VA (e.g., [6]–[10]). These issues are described and characterized in various ways, and often include terms such as cognition, perception, reasoning, decision-making, sensemaking, judgment, insight, and others that are essentially tied to humans. Although similar sentiments have been expressed at various points throughout the past decade (see, e.g., [10]–[12]), it seems that only a minority of the research is truly human-centered—i.e., it places human issues at the center of consideration. The majority of VA research is focused on tools, visualization and interaction techniques, and data-related issues. Even though many research contributions may involve users and/or discuss human-related issues (such as data encoding issues related to visual perception), they are not

1530-1605/15 $31.00 © 2015 Crown Copyright
DOI 10.1109/HICSS.2015.138

IEEE computer society

'human-centered' unless the human is at the center of concern. Most VA papers and research programs are probably more accurately characterized as technology- or data-centered—with human concerns necessarily involved but not central components. This is not to suggest that such research is not important; rather, it is stated simply to point out that truly human-centered research has not received a significant focus in VA.

The one aspect of human beings that has received a reasonable amount of attention in the general visualization literature is visual perception. This is perhaps the most obvious starting point for systematic research when dealing with visual phenomena. However, considering the fact that VA tools are used nowadays for complex activities that involve many interconnected processes of higher-order cognition, focusing on perception alone is not adequate [13]. In recent years, researchers have suggested that more attention must be paid to higher-order, complex cognition [5], [6], [13], [14]. Moreover, perhaps the most consequential area of concern—the one that ties everything else together—is interaction [2], [5], [9], [15]. It is through interaction that a human-information discourse is enabled. Considering that interaction with external information representations (i.e., outside of the head) is a fundamental aspect of most tasks and complex cognitive activities that humans perform, a focus on interaction should form a core component of VA research. Furthermore, previous research has demonstrated that the provision of interaction mechanisms is not sufficient for ensuring that tools are effective; rather, the quality of interaction—that is, *interactivity*—is of critical importance (see [3]). Although this is true of all components of a VA tool [2], a human-centered approach focuses on the components that have a direct relationship with users. While the scope of directly related components may change in various contexts, they are usually those that are 'external' to the tool (e.g., VRs), and not those that are 'internal' (e.g., databases). One exception is the notion of users guiding or steering the underlying algorithms or models from which the VRs are generated. For instance, users can adjust the parameters of underlying clustering algorithms, which changes the VR of the data, and may thus have an effect on the user's goals and mental models.

3. Design of interactivity

Although comprehensive analyses of interactivity do not exist in the VA literature, a brief examination of existing literature shows that visualization researchers certainly do have concerns about some aspects of interactivity. For example, one issue that is commonly discussed is the response time of user actions and its effect on the performance of a VA activity (e.g., [11], [16]–[18]). This concern is not about an interaction *per se*, but is about the quality of the interaction in reference to human characteristics and expectations. In other words, the tool itself is not concerned whether a response takes 100 milliseconds or 5 seconds; rather, it is the human user that is affected by the response time, thus affecting the overall efficacy of the VA activity. This is clearly a concern about interactivity that is more human-centered than technology-centered. One of the primary motivators for developing our interactivity framework was the lack of general frameworks, taxonomies, theories, and/or other comprehensive structures that can support systematic research and design of interactivity in visualization tools.

Without any type of conceptual support structure for thinking about human-centered issues, it is unlikely that VA tools will be designed in a systematic and principled manner. Based on the work of a number of design researchers, Stolterman [19, p. 63] suggests that designers appreciate and are inclined to use 4 forms of design support: (i) precise and simple tools or techniques (e.g., prototypes); (ii) frameworks that do not prescribe but that support reflection and decision-making (e.g., design patterns); (iii) individual concepts that are intriguing and open for interpretation and reflection on how they can be used (e.g., affordance); and (iv) high-level theoretical and/ or philosophical ideas and approaches that expand design thinking but do not prescribe design action (e.g., human-centered design). Our previous research supports design in 3 of the 4 ways (forms ii, iii, and iv). First, by providing a broad and coherent conceptual foundation, our research supports reflection on the design situation (e.g., what are the main tasks and activities in which users may engage? How does the provision of different interactions affect the overall quality?), and also facilitates principled decision-making (e.g., which actions should my tool support? when should the user be able to adjust the properties of VRs?). Second, our research provides a number of individual (although not isolated) concepts that are not commonly discussed by designers. For example, interaction patterns and their cognitive utility, interactivity elements and factors, and others, are all amenable to interpretation and use within a contextual design setting. Fourth, our research supports design by offering high-level, unified theoretical and philosophical ideas for design. Our research unifies a number of seemingly disparate concepts and research findings into a coherent over-arching framework. The overall framework can act as a conceptual support structure to facilitate the development of mental models that expand design

thinking, opening up new avenues of design and stimulating creativity in the mind of the designer. Effective design depends on designers being "fully aware of internal thinking processes and mental models, for they influence communication processes and content as ideas take shape" [20, p. 195]. Our research is intended to both inform designers and promote meta-cognitive awareness of the design process.

4. Conceptual frameworks

A major focus of our work in recent years has been on the development of a number of conceptual frameworks in an effort to bring more systematicity to research and design of visualization tools. We have developed the frameworks at a general level, independent of specific domains, tools, technologies, users, and data types. These individual frameworks are components of a larger, over-arching framework, which we have called EDIFICE (Epistemology and Design of human-InFormation Interaction in complex Cognitive activitiEs). Three components that have been fully developed deal with: interactivity [3], interaction design patterns [5], and adjustable properties of VRs [4]. A fourth component, dealing with a pattern language for designing VRs, is currently under development. For the purposes of this paper, we will focus mainly on the first and second components. We will describe them briefly in the following sub-sections.

4.1 Interactivity

Because interactivity is a broad and general construct, it encompasses many issues. To help deal with these and bring clarity to the discussion, we divide interactivity into smaller conceptual units: first, we categorize interactivity into two broad sub-components: external and internal interactivity. External interactivity is concerned with the quality of interaction between the user's mental space, perception, VRs, interaction space, and other considerations 'external' to the tool. Internal interactivity is concerned with the quality of interaction between the VRs, algorithms, data-centric operations, storage, and other considerations 'internal' to the tool. Second, we categorize interactivity into two levels: micro and macro. Interactivity at the micro level is concerned with the structural or anatomical elements of individual interactions. Since any interaction logically consists of an action and a reaction, each of these can be analyzed into structural components. For instance, consider a single filtering interaction. To perform the action component, the user can: articulate it in different ways (e.g., verbally, manually); act directly or indirectly upon a VR of interest; and carry it out instantaneously or over a period of time. These constitute some of the elements of the action component. Similarly, the reaction component has a number of elements, each of which can be operationalized in a number of different ways. Although the basic interaction stays the same (e.g., acting upon a VR to show or hide a subset of encoded information), there are numerous ways in which the different structural elements can be operationalized. Empirical studies have shown that many of the elements and their various operational forms have a significant influence on the overall interactivity and efficacy of a tool (see [3]). Interactivity at the macro level is concerned with how different interactions are combined and put together in the context of performing tasks and activities. Research at this level deals with the factors that affect the overall quality of interaction, in the context of the properties and relationships of a tool's aggregated interactions and how interactivity emerges from these. For the sake of brevity, we discuss only macro-level interactivity in this paper. Table 1 lists and briefly describes some macro-level interactivity factors. Although it is important to consider these macro-level interactivity factors, they are not entirely sufficient to guide interactivity design on their own. Designers must also have a good understanding of what interaction patterns are available, so that they can reason about the interactivity factors in a principled manner.

Table 1. Macro-level interactivity factors

Pattern	Description
diversity	number and diversity of interactions that are available to the user
complementarity	harmonious and reciprocal relationships among interactions, and how well they work with and supplement each other
fitness	appropriateness of interactions for the given VRs, the tasks and the activity, and the user's needs and characteristics
flexibility	range and availability of adjustability options
genre	types of transactions that are available to the user

4.2 Interaction

For many years, researchers have recognized the need for explicating the design space of interaction in VA [11], as well as the relationship between interaction and higher-order cognitive processes such as reasoning, decision-making, and others [7], [10], [14]. While some work has been done in this regard,

most of the literature is concerned with data-, tool-, task-, or user-specific interaction techniques, and not with general, comprehensive patterns (see [5] for more discussion). It is not feasible to engage in systematic design of interactivity with only a series of scattered interaction techniques with which to work. One of the main components of our framework is a comprehensive catalog of fundamental action patterns that users perform on VRs during a human-information discourse with visualization tools [5]. Each action pattern is characterized at a general level in terms of its epistemic and cognitive utility, and is discussed in the context of performing complex cognitive activities with visualization tools. The development of the catalog was motivated by the lack of existing comprehensive catalogs or taxonomies concerned with interaction in the visualization literature. In most realistic VA situations, users perform a series of different actions to carry out even a single task. Having a comprehensive catalog of possible actions and their cognitive utility can greatly aid designers in systematic design of their VA tools. Table 2 lists and briefly describes 10 of the 32 patterns.

4.3 Distribution of analytic processing

One extant issue in VA research is how to best distribute analytic processing across the different components of VA tools. Some questions that designers should ask include the following: How much of the analytic processing should the computer be responsible for? How much should be left to the user? Should some be done in a joint fashion, involving both computational and mental processing? Although VA is fundamentally concerned with joint processing, it is not always clear how to best distribute the processing load. In previous work [21], we have identified a number of issues on which the ideal distribution is dependent.

One essential issue to consider is the intended users' characteristics. Whether the user is a novice or an expert, for example, influences the degree to which the processing load should be placed on him or her. In many cases, experts want to have more control and involvement in the trajectory of a VA activity than do novices, due to their high degree of domain knowledge and expert judgment. Furthermore, an issue that is particular to VA tools is the degree of control users have over the underlying models or algorithms that constitute the computational aspect of the analytics. In the case of the tool we are presenting in this paper, we will discuss in the following section how expert users were given a reasonable degree of control over the parameters of the underlying analytics algorithms.

Table 2. Some action patterns from [5]

Pattern	Description
annotating	augmenting VRs with personal meta-information
arranging	changing the ordering of component parts of VRs
assigning	binding a feature or value to a VR or one of its components
blending	fusing multiple VRs together into one
cloning	creating multiple identical copies of VRs
comparing	determining the degree of similarity or difference between VRs
measuring	quantifying some aspect(s) of a VR
scoping	dynamically working forwards or backwards to view the compositional development or growth of the underlying data/information
transforming	altering the geometric form of a VR
translating	converting VRs into alternative informationally- or conceptually-equivalent forms

5. VARSITY

In this section we present our tool, named VARSITY (Visual Analytics of university Research networkS and IndusTrY collaborations). We give a brief overview and describe the context in which it was designed. We will give a quick description of some of the technical aspects of the tool and the data analytics component. The subsequent section will discuss the role of the aforementioned frameworks in its design.

5.1 Overview and context

University administrators (e.g., deans, department chairs, and other stakeholders) require a good understanding of the research networks that exist within their own faculties and departments in order to identify areas of strength and areas of potential growth, to make effective decisions about funding strategies, and to perform numerous other tasks and activities. Furthermore, being able to identify and make sense of industry collaborations facilitates analysis of the effectiveness of the university's knowledge transfer and exchange strategies. The goal of VARSITY is to enable such users to gain insight into the nature of the relevant data, to ultimately make strategic administrative decisions.

5.2 Tasks and activities

We have previously characterized VARSITY as a decision-support tool. Although the ultimate goal of using the tool is making strategic decisions, the overall activity is complex and involves many other high-level

cognitive activities such as sensemaking, problem solving, and planning. While each of these activities has its own characteristics, we will not consider them in any detail here. It is very typical that numerous activities are embedded within one another in any complex VA activity [2], [5], [22]. For example, it is entirely conceivable that an activity in which an analyst is engaged may involve sensemaking, problem solving, analytical reasoning, planning, and learning—all before making some sort of final decision. Although this is the case, researchers often describe a VA tool or research endeavor as if it supports only one specific activity: e.g., decision-making or sensemaking. As a discussion of this nature is outside the scope of this paper, we simply characterize VARSITY as a decision-support tool.

Through a participatory design approach, 2 main areas of focus for users were identified: (1) researchers' collaborations and (2) funding and grants. In the context of collaborations, users were interested in understanding questions such as: which researchers are collaborating? What proportion of collaboration is happening within departments? Are there clusters of collaboration? Do some types of collaboration (e.g., grants) tend to lead to others (e.g., publications)? How do these change over time? In the context of funding and grants, users were interested in understanding questions such as: What is the distribution of funding agencies across departments? Which programs within funding agencies are most prevalent, and what are their connections to various departments? What proportion of grant applications are accepted versus those that are rejected? The examples in Section 6 will demonstrate how some of these tasks were supported.

5.3 Technical aspects

We do not intend to go into detail about the technical aspects of VARSITY, as it is not part of the focus of this paper. However, we will provide a very brief description for readers who are interested. Briefly, we are running a Node.js web server. We have some scripts that run and automatically retrieve publication data from two APIs: SciVerse Scopus (http://www.scopus.com) and Thomson Reuters Web of Science (http://wokinfo.com). The relevant data is stored in an Apache CouchDB database. Subsequently, a script is run to extract the relevant data and determine relevant relationships (e.g., connect a publication author to a department and to a grant). This data is then stored in a relational MySQL database. When the client requests data, it is sent in JSON format. On the client side, aside from HTML5 and CSS3, two main JavaScript libraries are used for visual representation and interaction. The first is D3.js, a library for visually encoding data as SVG elements. This is done by binding data to elements within the Document Object Model (DOM). The other is jQuery, a library that facilitates client-side scripting. Both libraries offer a level of abstraction for incorporating interaction into the tool, mostly in the form of determining responses to user actions.

5.4. Data analytics

The data analytics component of VARSITY provides our expert users with functionalities that allow them to perform decision-making and sensemaking procedures— in tandem with the analytic processing of the tool—at multiple levels of granularity. The analytics component currently allows for two main procedures: (1) network clustering and community detection, and (2) suggestive analysis of causal relationships between grants and publications. We will briefly discuss this component of VARSITY in the next 2 sub-sections, and revisit and expand on the second procedure in Section 6.2.

5.4.1. Network clustering and community detection. The community detection algorithm in VARSITY gives the user the ability to visually discover and navigate communities (i.e. clusters) within the network of collaborations (e.g., collaborations on grants, publications, and co-supervisions) [23], [24]. This functionality allows the user to discover patterns in the collaboration network and use them for further sensemaking and decision-making tasks. Although there are many available methods, we are currently using the Louvain method [25], which tries to optimize a modularity-based objective function to detect communities in the network of collaborations. The algorithm initially considers every node as a community, and then merges the nodes in a neighborhood in such a way that maximizes the local modularity. We are also planning to integrate a stability method (e.g., [26]), as well as a Markov random walk method (e.g., [27]) in near future. The latter is built on the assumptions that a random walk on the nodes of a network is much more likely to remain in a community of nodes (a.k.a. cluster or component) for a long period, rather than moving to an adjacent one. In other words, the nature of a community imposes more walks within and fewer among the clusters. Hence, the denser a community is, the more likelihood of remaining in the same community. In addition, in networks that are conceptually linked to the behavior of human beings (such as research collaboration networks), there often exists a very large subset of nodes that contains a majority of all nodes, known as the giant component. The methods that are to

be added allow expert users to adjust, visually perceive, and compare these characteristics of the network.

5.4.2 Analysis of relationships between grants and publications. VARSITY currently provides the expert user with a list of possibly related publications with respect to a certain selected grant. Based on the placement of these publications and the likelihood of them relating to the grant, our users can examine the effect of publications being used to receive a grant, as well as those that are a result of an accepted grant, which can then inform strategic decision-making. This aspect of VARSITY will be described more fully in Section 6.2 through the use of some examples.

6. Design considerations

In this section, we describe the design process of VARSITY, with a particular focus on the role of the aforementioned conceptual frameworks in guiding the design. There are a number of considerations that we do not cover here. For instance, the choice of different VRs and the interaction techniques will not be covered, nor all of the interactivity-related design decisions. Considering that this tool was developed over a period of more than one year, it would be impossible to completely describe a significant portion of the design process. We simply provide a sampling of some various aspects that illustrate some of the rationale and systematicity of the design process. As a side-note, due to the sensitivity and confidentiality of some of the data, labels and values in some of the screenshots in this section have been altered to hide confidential information, and none of their contents should be assumed to be accurate.

In our previous work, analysis of macro-level interactivity was a preliminary attempt to give some general structure to interactivity at this level. Therefore, at this point in the research, we are far from having a comprehensive understanding of macro-level interactivity that could lead to any sort of principled and comprehensive prescriptive guidance. However, as a descriptive tool, the framework can still support systematic design thinking at a high level.

Consider the two interactivity factors of diversity and complementarity. We know from previous research that complex cognitive activities do not follow a single, linear trajectory, and that it is beneficial to provide users with a number of complementary interactions [5]. At this stage of the research, we do not have a precise understanding of which are best. However, with an awareness of these two factors, and a catalog of possible interaction patterns along with their cognitive utilities, experienced designers can rely on their domain expertise to systematically determine which interactions should be provided. Assuming that the interaction catalog is comprehensive, designers can use it in a generative fashion to combine interactions in interesting ways. Figure 1 shows the result of a series of interactions with VARSITY: the user has first *transformed* the VR by adjusting parameters of the force layout algorithm to more easily detect clusters of interest and then zooming into a specific region; has then *selected* a group of nodes (people); has then *filtered* them according to a specific criterion (department); and will then *drill* into those selected items to see much more detail and to conduct systematic comparisons with other VRs (not shown). These are diverse and complementary actions that help the user to perform typical sensemaking tasks such as identifying information items of interest, ranking and categorizing information items, and comparing information items.

Because the catalog deals with general patterns of interaction and not with specific techniques, designers can also determine which techniques to use to implement instances of the interaction patterns. Using such conceptual frameworks should promote systematic thinking about the general, high-level aspects of design, while still allowing for creativity and flexibility at the level of implementation. For example, although we determined that a selecting interaction should be made available, there were many potential techniques that could be used to instantiate the selecting pattern. The screen captures in Figure 1 show how the user can select a group of items using a lasso selection technique (top). The user can also select or deselect individual nodes by clicking on them directly. These are two different techniques for carrying out the same general interaction.

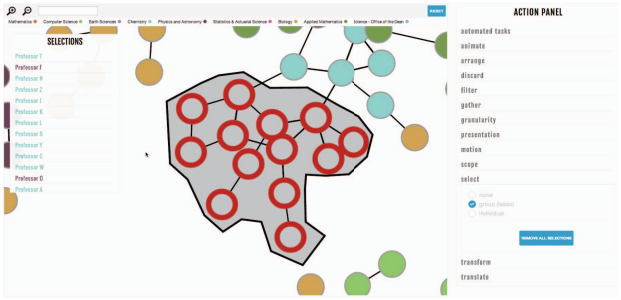

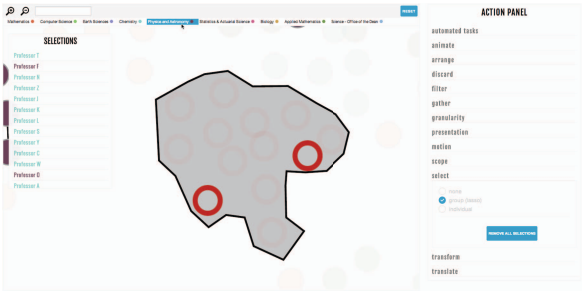

Figure 1. Combining a number of different interactions to carry out a task.

Another macro-level interactivity factor is *fitness*, which is concerned with the appropriateness of interactions for given VRs and tasks. This factor is multi-faceted, and one of its facets is task-fitness. We know that one of the tasks that users want to perform with VARSITY is determining the proportion of intra- versus inter-departmental collaborations over a specific period of time. Furthermore, we know that the network VR and its available interactions is not very good at supporting this task—in other words, it is not task-fit. By consulting the interaction catalog, we see that one of the actions patterns with utility in this regard is *translating*. By providing this interaction possibility, users can translate to a VR that is better fit for this specific task, and can translate back to the node-link VR when moving to other tasks for which it is more appropriate. Three screenshots that show progressive stages of carrying out such a task are shown in Figure 2: first, the user is working with the node-link diagram—it is not task-fit, so she performs a translating action to convert it to a matrix VR (middle). Subsequently, she performs another action to *arrange* the VR in such a way that inter- and intra-departmental collaborations become perceptually evident (bottom). From this VR, it is readily apparent that most collaborations, across all types, happen within and not between departments.

Since the interaction catalog is comprised of both well-known interaction patterns (e.g., filtering, drilling, searching) and lesser-known patterns (e.g., scoping, blending, assigning), another benefit of using the catalog is that it may encourage the designer to implement interactions with which he or she is not familiar and/or experienced. For example, while most designers are familiar with filtering and searching patterns, many are not familiar with the scoping pattern (i.e., dynamically working forwards or backwards to view the compositional development or growth of the underlying data). During the design of VARSITY, considering that an explicit goal of the users was to be able to identify changes in the information space over time, and to reason about temporal patterns, we determined that the scoping pattern had clear utility for achieving such a goal.

Figure 2. Three screenshots show progressive stages of carrying out a task.

Without access to the interaction catalog, we may not have been able to think so systematically about implementing various interaction patterns to support the necessary tasks. The screen captures in Figure 3 depict successive stages of a scoping action in VARSITY, showing growth of the network from 2008-2009 (left), 2008-2010 (top right), and 2008-2011 (bottom right). By performing this action, the user can

identify important connections within the network. For instance, in Figure 3 (left), a number of small clusters are evident, four of which have been circled and labeled for clarity. As the user performs the scoping action to show growth in the VR until 2010, three of the clusters (A, B, and C) merge into one due to some collaborations that took place in that year. The user can clearly identify the collaborations that connect the three clusters. In a similar fashion, when scoping the growth further to the year 2011, a collaboration can easily be identified that incorporated the 'D' cluster into the larger cluster. Being able to identify such connections helps the user to make sense of the growth of the collaboration network, and to reason further about other aspects of the information space.

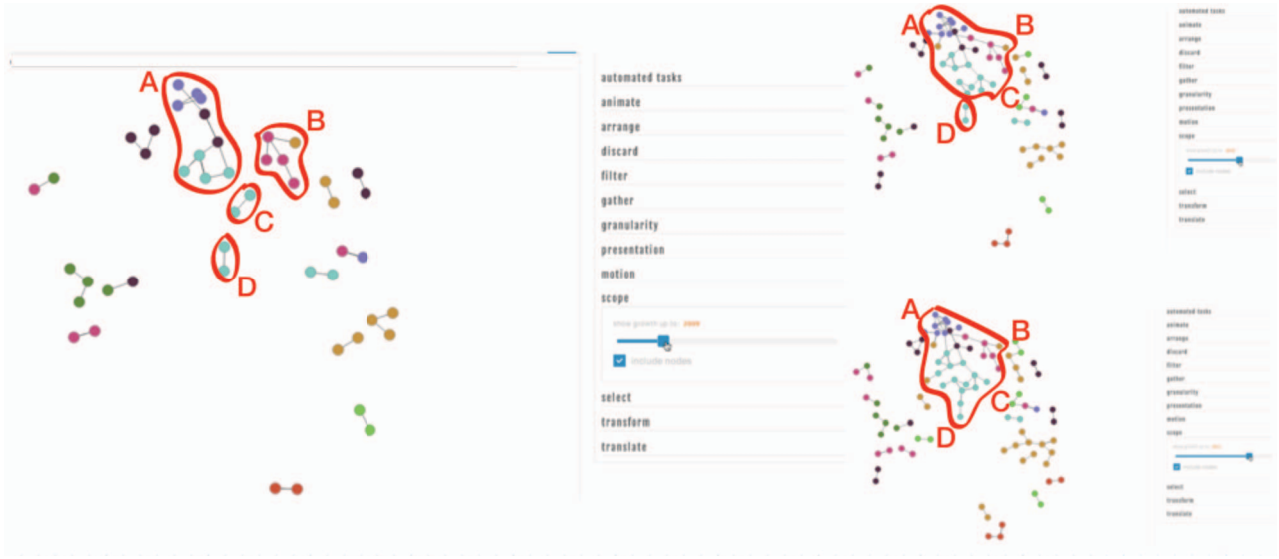

Figure 3. Successive stages of a scoping action, showing growth of the network until 2009 (L), 2010 (TR), and 2011 (BR). Areas of interest have been circled in red for explanatory purposes.

6.2 Distribution of analytic processing and user-guided analytics

As mentioned in Section 4.3, one important design consideration for VA tools is how to best distribute analytic processing across the joint cognitive system. Based on our analysis of VARSITY's users and their tasks, we decided to allow the users to adjust the parameters of the underlying algorithm if they desired. Because many of the intended users are domain experts, they can take on more of the processing load by determining how some of the parameters of the processing should be adjusted. If a user does not have the expertise, however, she can simply work with whatever information the default parameters provide. Figure 4 demonstrates this how this is so. In the initial state (top), the user is exploring the grants that have been awarded, which are encoded in the form of two linked treemaps that represent grants based on department and funding agency respectively. The user then selects a grant that is of interest (Figure 4, middle), and the algorithm outputs a weighted list of related publications for the grant, based on the following:

- time period within which the grant was active
- date of publications
- collaborators involved (as both investigators and authors)
- keyword matching

The algorithm assigns a score to each potentially related publication, to convey how likely it is related to the grant. The user then adjusts (increases) the score threshold of the algorithm, and the VR is updated (Figure 4, bottom). The procedure reveals not only potential causal relationships between grants and publications, but also encodes a threshold value. In other words, because two of the publications have values below the threshold, their opacity encoding is adjusted, and the user is made aware of the changes in the underlying data model. For the sake of space, other possibilities for adjusting the algorithm parameters are not shown here.

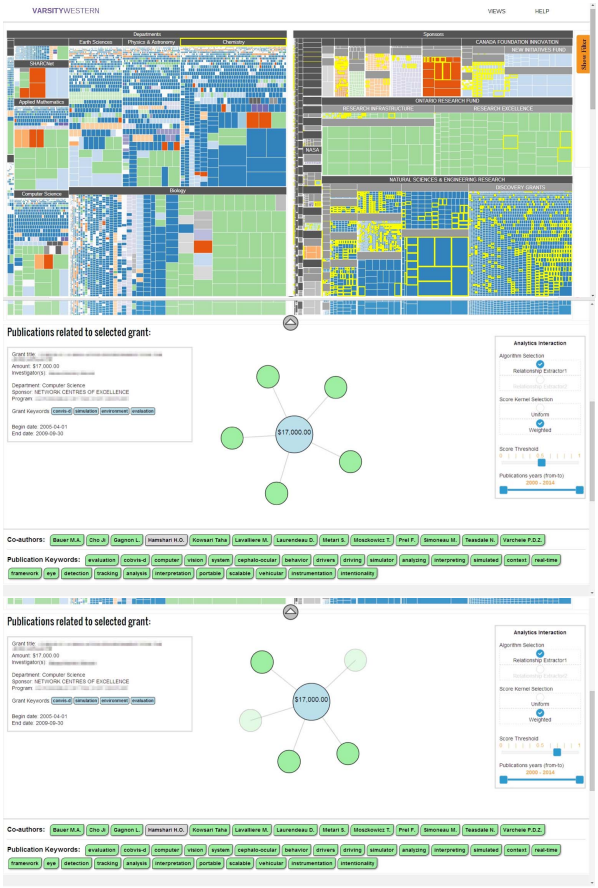

Figure 4. Successive stages of a task in which the user adjusts parameters of an underlying algorithm

7. Summary and future work

The initial premise put forward in this paper was that the quality of interaction among all components in a VA system is critical to the overall efficacy of the system. Previous work has characterized this quality of interaction as interactivity. We have argued that although not much focus has been placed on developing and analyzing the concept of interactivity in a comprehensive manner in existing VA literature, VA researchers are indeed aware and concerned about interactivity in VA tools. To engage in systematic design of VA tools, however, designers require more than isolated statements and research findings regarding aspects of interactivity in VA systems. Research on design thinking has shown that designers benefit from coherent conceptual frameworks that can guide design in a systematic, principled fashion. In recent years, we have developed a number of conceptual frameworks that shed light on the nature of interactivity in VA. In this paper, we have briefly described and synthesized some of the considerations necessary for designing VA systems with a high degree of interactivity. We have also presented VARSITY, a tool that has been developed over the past year using the aforementioned conceptual frameworks. We have attempted to demonstrate, in a limited amount of space, the role and importance of these frameworks in the design process, and to illustrate how and why the design of VARSITY was more human-centered as a result. Moreover, we have discussed the issue of how analytic processing should be distributed in a VA system, and briefly mentioned some of the considerations for determining an ideal distribution. Although we have not previously analyzed this aspect of VA in detail, it is an interesting and important area of current VA research, and designers should have access to conceptual frameworks that coherently address this issue. We thus envision further work on this issue as a logical next step for developing our frameworks further. As we described in Section 4.1, previous work considered only external interactivity in detail, and left concerns about integrating external and internal interactivity for future work [2]. We see this current paper, and the development of VARSITY, as a connecting point for moving in this direction.

With respect to VARSITY in particular, we will soon be providing a small set of real end-users with access to the tool, and will be soliciting feedback and collecting information on user behavior. We have already given the tool to a small number of visualization experts, who have given generally positive feedback. While we believe the informal feedback from our small pool of users will certainly be useful, to truly validate the efficacy of the tool we must conduct formal user studies. As we already have a set of real end-users lined up, we should be able to design and conduct such studies in the near future.

Finally, although mentioned only briefly in Section 4.1, another promising area of future work is in extending our general micro-level interactivity framework to deal with more specific VA concerns. Common concerns about action response time, for example, are fundamentally concerns about micro-level interactivity, and integrating a number of extant issues into one unifying framework could be very useful for systematic research and design. VA researchers and designers would surely benefit from a framework that incorporates common micro-level interactivity elements, and discusses how and why they contribute to the overall interactivity of a VA system. Such work would hopefully bring more clarity to the VA research and design landscape, perhaps a decade or more after initial calls for such research were made in the seminal VA agenda [11].

8. References

[1] W. Ribarsky, B. Fisher, and W. M. Pottenger, "Science of analytical reasoning," *Inf. Vis.*, vol. 8, no. 4, pp. 254–262, 2009.

[2] K. Sedig, P. Parsons, and A. Babanski, "Towards a characterization of interactivity in visual analytics," *J. Multimed. Process.*, vol. 3, no. 1, pp. 12–28, 2012.

[3] K. Sedig, P. Parsons, M. Dittmer, and R. Haworth, "Human-centered interactivity of visualization tools: Micro- and macro-level considerations," in *Handbook of Human-Centric Visualization*, W. Huang, Ed. New York: Springer, 2013, pp. 717–743.

[4] P. Parsons and K. Sedig, "Adjustable properties of visual representations: Improving the quality of human-information interaction," *J. Assoc. Inf. Sci. Technol.*, vol. 65, no. 3, pp. 455–482, Feb. 2014.

[5] K. Sedig and P. Parsons, "Interaction Design for Complex Cognitive Activities with Visual Representations: A Pattern-Based Approach," *AIS Trans. Human-Computer Interact.*, vol. 5, no. 2, pp. 84–133, 2013.

[6] R. Arias-Hernández, T. M. Green, and B. Fisher, "From cognitive amplifiers to cognitive prostheses: Understandings of the material basis of cognition in visual analytics," *Interdiscip. Sci. Rev.*, vol. 37, no. 1, pp. 4–18, 2012.

[7] W. Aigner, "Understanding the role and value of interaction: First steps," in *International Workshop on Visual Analytics*, 2011, pp. 17–20.

[8] M. Hegarty, "The cognitive science of visual-spatial displays: Implications for design," *Top. Cogn. Sci.*, vol. 3, no. 3, pp. 446–474, Jul. 2011.

[9] A. Endert, M. S. Hossain, N. Ramakrishnan, C. North, P. Fiaux, and C. Andrews, "The human is the loop: new directions for visual analytics," *J. Intell. Inf. Syst.*, Jan. 2014.

[10] J. Meyer, J. Thomas, S. Diehl, B. Fisher, D. A. Keim, D. Laidlaw, S. Miksch, K. Mueller, W. Ribarsky, B. Preim, and others, "From Visualization to Visually Enabled Reasoning." 2007.

[11] J. Thomas and K. Cook, *Illuminating the path: The research and development agenda for visual analytics*. IEEE Press, 2005.

[12] B. Fisher, T. M. Green, and R. Arias-Hernández, "Visual Analytics as a Translational Cognitive Science," *Top. Cogn. Sci.*, vol. 3, no. 3, pp. 609–625, Jul. 2011.

[13] P. Parsons and K. Sedig, "Common visualizations: Their cognitive utility," in *Handbook of Human-Centric Visualization*, W. Huang, Ed. New York: Springer, 2013, pp. 671–691.

[14] T. M. Green and B. Fisher, "The personal equation of complex individual cognition during visual interface interaction," in *Human Aspects of Visualization*, A. Ebert, A. Dix, N. Gershon, and M. Pohl, Eds. 2011, pp. 38–57.

[15] K. Sedig, S. Rowhani, and H.-N. Liang, "Designing interfaces that support formation of cognitive maps of transitional processes: an empirical study," *Interact. Comput.*, vol. 17, no. 4, pp. 419–452, Jul. 2005.

[16] M. Worner and T. Ertl, "Retaining Interactivity in a Visual Analytics System for Massive Public Transportation Data Sets," *2014 47th Hawaii Int. Conf. Syst. Sci.*, pp. 1354–1363, Jan. 2014.

[17] D. Keim, G. Andrienko, J.-D. Fekete, C. Gorg, J. Kohlhammer, and G. Melancon, "Visual analytics: Definition, process, and challenges," in *Information Visualization, LNCS 4950*, A. Kerren, J. T. Stasko, J.-D. Fekete, and C. North, Eds. Berlin, Heidelberg: Springer-Verlag, 2008, pp. 154–175.

[18] S. K. Card, J. D. Mackinlay, and B. Shneiderman, *Readings in Information Visualization: Using Vision to Think*, vol. 1st, no. 0. Morgan Kaufmann, 1999, p. 686.

[19] E. Stolterman, "The Nature of Design Practice and Implications for Interaction Design Research," *Int. J. Des.*, vol. 2, no. 1, pp. 55–65, 2008.

[20] Y. M. Hansen, "Visualization for Thinking, Planning, and Problem Solving," in *Information Design*, R. Jacobson, Ed. Cambridge, MA: MIT Press, 2000, pp. 193–220.

[21] P. Parsons and K. Sedig, "Distribution of Information Processing while Performing Complex Cognitive Activities with Visualization Tools," in *Handbook of Human-Centric Visualization*, W. Huang, Ed. New York: Springer, 2013, pp. 693–715.

[22] T. M. Green and R. Maciejewski, "A Role for Reasoning in Visual Analytics," in *Hawaii International Conference on System Sciences*, 2013.

[23] S. Fortunato, "Community detection in graphs," *Phys. Rep.*, 2010.

[24] S. E. Schaeffer, "Graph clustering," *Comput. Sci. Rev.*, vol. 1, no. 1, pp. 27–64, Aug. 2007.

[25] V. Blondel and J. Guillaume, "Fast unfolding of communities in large networks," *J. Stat. ...*, pp. 1–12, 2008.

[26] J.-C. Delvenne, S. N. Yaliraki, and M. Barahona, "Stability of graph communities across time scales.," *Proc. Natl. Acad. Sci. U. S. A.*, vol. 107, no. 29, pp. 12755–60, Jul. 2010.

[27] P. Pons and M. Latapy, "Computing communities in large networks using random walks," *Comput. Inf. Sci. 2005*, 2005.