

Visual Analytics for Complex Concepts Using a Human Cognition Model

Tera Marie Green,¹ William Ribarsky,¹ & Brian Fisher²

¹Charlotte Visualization Center
University of North Carolina
at Charlotte

²School of Interactive Arts and Technology
Simon Fraser University

ABSTRACT

As the information being visualized and the process of understanding it both become increasingly complex, it is necessary to develop new visualization approaches that facilitate the flow of human reasoning. In this paper, we endeavor to push visualization design a step beyond current user models by discussing a modeling framework of human “higher cognition.” Based on this cognition model, we present design guidelines for the development of visual interfaces designed to maximize the complementary cognitive strengths of both human and computer. Additionally, this paper introduces cognition-based principles of visualization design. Some of these principles are already being reflected in the better visual analytics designs, while others have not yet been applied or fully applied. But none have explained the deeper rationale that the model provides. Lastly, we discuss and assess these visual analytics guidelines through the evaluation of several visualization examples.

KEYWORDS: visual analytics, cognition and perception theory, embodied cognition, visualization taxonomies and models

INDEX TERMS: visual analytics, cognition and perception theory, embodied cognition, visualization taxonomies and models

1 INTRODUCTION

In a previous paper [1], we discussed the necessity of considering principles of human reasoning, problem-solving, and decision-making, in addition to the already familiar areas of sensation and perception, to develop information and knowledge visualizations capable of attacking today’s complex, important problems, which require both reasoning and analysis. In most visualization development up to this point, “higher cognition” processes have been considered something of a human black box, into which information is inputted, and from which appropriate responsive behaviour is somehow obtained. (We will discuss this metaphorical view of human-visualization interaction during our consideration of the van Wijk operational model of visualization in Section 4.) But as both interactive visualizations and their tasks become semantically complex, it becomes apparent that we must peek into the black box in an attempt to design a human-computer system and its interactions.

In truth, the black box analogy is an imperfect one for human

higher cognition. Psychology and other behavioral sciences have been researching reasoning and other thought processes for decades. One reason that much of this research has, as yet, been unused in the construction of interactive visualizations is the lack of a unifying theory of human reason. It is, as Newell once wrote, as if “science advances by playing twenty questions with nature.” [2]. The study of higher cognition is not studied holistically; it is usually broken down into bite-sized subprocesses, and competing theories of small, often binary, aspects of reasoning have dominated the research.

In addition, holistic higher cognition, unlike human sensation and perception, employs combinatorial use of multiple heuristics, which is rarely binary, almost never perfectly sequential, and has, as yet, defied traditional model-based prediction.

But while complex, higher cognition is still predictable to some degree. The number of available heuristics is finite, and is therefore theoretically knowable. And while extant research may disagree on the fine points, it does confirm one essential fact: humans use the simplest heuristic possible to accomplish the task at hand. This by itself is no small starting point. But humans are also predictable in other ways, which we will soon explore.

This paper endeavours to lay the framework of a human cognition model, whose guidelines would guide the development of visual interfaces more able to attack the complex problems now at hand. Additionally, this paper shows how this model contributes new, cognition-based principles of visualization design. Some of these principles are already being used in the better visual analytics designs, but without the deeper rationale that the model provides. We will discuss and evaluate these visual analytics methods. Other principles from the model have not been applied or not fully applied, and we will discuss how their implementation and use will be of benefit.

2 COMPLEMENTARY STRENGTHS

Both human and computer bring strengths to a mixed-initiative visualization, in which both the computer and the human initiate processes and with which both collaborate in the exploration and creation of knowledge. Several of the obvious strengths are complementary, which further strengthens the potential of the collaboration.

2.1 Human strengths

Some of the earliest reasoning skills humans develop are those of adaptation and accommodation [3]. Adaptation is the ability to incorporate newly perceived information into extant knowledge schema and it relies heavily on an ability to categorize sensory stimuli at an “instantaneous” rate. Even when what is perceived is so novel it will not fit existing knowledge schema, accommodation allows a human to temporarily place a marker in a closely similar schema or create a new one [4]. This

-
- Tera Marie Green: grepmon@gmail.com
 - William Ribarsky: ribarsky@uncc.edu

fast and frugal reasoning ability [5] enables humans to more effectively deal with rapidly-changing situations. Biederman's 1987 "recognition by components" model provides a mechanism by which basic level categorization of objects takes place rapidly and accurately regardless of viewpoint or changes in non-essential characteristics of those objects. [6]. This all seems effortless to a human, and allows the reasoning process to move forward, even when the information is incomplete. It is also vastly superior to computer recognition, which is rigid; when presented with what it has not been programmed to recognize, the computer is stymied, often aborting the process.

Human perceptual abilities are also well adapted to complex and rapidly changing scenes, defined here as complex sets of objects and events distributed in space that interact with each other in often novel ways. This takes place through a combination of a low-level "gist" mechanism that recognizes important scene characteristics and relationships [7]. This preattentive visual process guides the allocation of multiple attentional tokens i.e. FINSTS, see [8,9] that support the automatic calculation of a set of operations [10] on their properties and relations to each other. All of this occurs prior to, and in support of, endogenous attention (attention to the task at hand). This resulting cascade of processes frees cognition for consideration of higher order characteristics of the information contained in the display rather than the display itself-- object properties and spatial relations with other objects and causal relations to events taking place. Thus, cognitive operations can proceed using more parsimonious representations that are well-suited for the task at hand. This can be thought of as a two-step process by which unconscious inference the "logic of perception", [11] works hand-in-hand with cognitive processes to support reasoning. In this view, expertise is not only a characteristic of higher-order cognitive logic, but also of perceptual logic, which can be trained to better support cognitive operations through "perceptual expertise" [12].

Partially due to a lifetime of experience in adaptation and accommodation, humans are also vastly superior reasoners to computers. Humans have a compendium of reasoning and problem-solving heuristics, which can be used singly or concomitantly to accomplish the task at hand. The simplest of these, elimination heuristics such as satisficing [13], eliminate available choices that do not possess an identified important attribute. Elimination heuristics require little cognitive effort, and so are often what a human will use first in an attempt to narrow down available choices.

Of course, if the problem becomes semantically complex, more effort is required. Our model assumes a mental model to inferred rules mechanization [14], wherein the human uses all available information to create a mental model of the concept being considered. From this model, the human infers generalizable rules – sometimes in a matter of seconds – that are used in later instantiations of a similar concept or problem. Because these models are based entirely on available (including previously held) information, it is imperative that all pertinent information is available to avoid the creation of incomplete mental models, which are, in turn, likely to be the basis of invalid rules.

2.2 Computer strengths

A computer is capable of two distinct processes that complement human reasoning strengths well: superior working memory and information processing without cognitive biases. Humans depend on their working memory as they reason, but

are, at best, able to remember 7 ± 2 chunks of information [15]. The computer, on the other hand, has a "working memory" limited only by hardware. The computer's ability to keep all pertinent information present before the human aids in complete mental modelling, among other things.

The other computer strength is the lack of inherent biases. Unlike humans, computers do not filter out pertinent information in accordance with a perceived mindset or due to the way a problem is presented [16,17]. By presenting all relevant information, the computer can aid not only in mental modelling, but also in the analysis of competing hypotheses.

3 USE OF A HUMAN COGNITION MODEL

This section will discuss the ways in which a knowledge of higher cognition focuses a multi-modal human cognition model (HCM) design model, as well as several guidelines which can be derived from the model's use. On each of the HCM's submodels, please see [1] for more discussion.

3.1 Information Discovery & Knowledge Building

The central process of the HCM is *discovery*, during which the computer presents information in an ontological-like structure within a relevant, possibly human-defined, context. Presenting information with a relevant context is one method of mitigating human cognitive overload in the midst of overwhelming semantic data points. The human directly interacts with the visualized information, focusing the attention of discovery.

An intuitive multi-model visualization also encourages knowledge building through *new knowledge creation*. Throughout the process of discovery, the human may uncover a relationship between two currently unrelated concepts or ideas. By creating a new relationship between the two concepts and perhaps annotating the relationship, the human collaborator can extend the knowledge base of the visualization, not only for what is to be accomplished in that particular session, but for every session by every human who uses the visualization thereafter.

The computer can augment the discovery of relevant information through *computer-aided discovery*. Through observation of what interests the human collaborator, the computer can suggest information that is semantically related, but up to this point, has not been considered. This also would include relational data which has been added by other human collaborators, which allows one person to learn from another. The human is free to explore or to reject the suggestion. But by making the effort in ensuring that nothing important is overlooked, the computer works to counteract human cognitive biases which can interfere with complete mental modelling.

3.2 Guidelines for Discovery and Knowledge Building

We will now briefly discuss several guidelines based upon the HCM discovery and knowledge submodels.

3.2.1 Multiple views

When the information being explored is semantically rich, and could be visualized through a variety of categorization levels, it is often left to the discretion of the visualization developer as to which level merits the primary view. It is important to categorize information to aid the human in directing attention, but we would argue that a visualization that utilizes multiple organizational views of the same information can be a powerful

aid. As the human interacts with information in any view, the relational changes are visualized in all views.

While the concept of multiple views is not a new one [18], what we would highlight is how multiple views are informed by human cognition. First, as humans perceive information in a variety of ways including through the filter of their own assumptions, patterns are more likely to be discovered if represented multiple ways.

Secondly, as we have discussed previously, humans prefer to narrow down the field of choices by eliminating those that do not possess desired attributes. This is usually done before utilizing more complicated heuristics. Multiple views make the process easier; multiple layers of relational attributes are readily knowable without additional search.

Thirdly, multiple views enable more intuitive manipulation. Humans themselves do not interact with information in one dimension; humans are capable of multi-layered processing: perceptual, emotional, and higher-cognitive. Indeed, of all the guidelines we will discuss, use of multiple views is the one most likely to lead to spontaneous insight.

3.2.2 Direct interaction

By definition, a well-designed information visualization allows the user to directly interact with information. But we would take direct interaction one step forward. In computer-aided discovery, for example, the guideline of direct interaction would

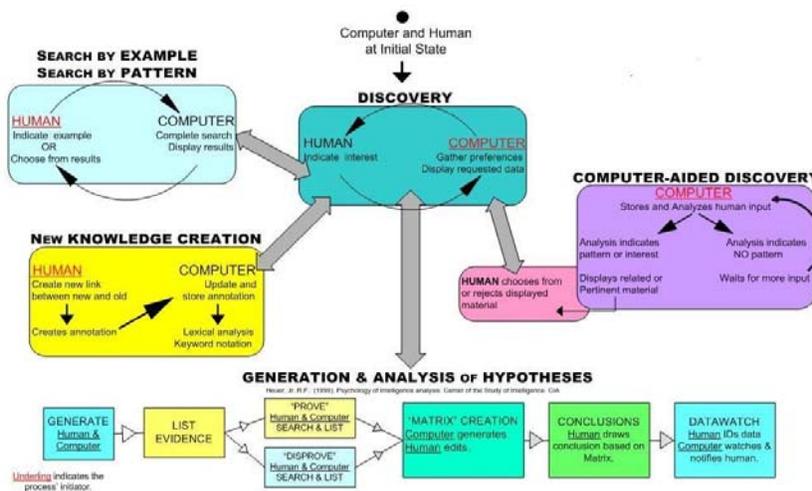


Figure 1. The Human Cognition Model.

propose that whatever tactic the computer uses to suggest relational information to the human must be done without interfering with a human's train of thought or flow of reasoning.

Additionally, direct interaction facilitates the goals of other HCM guidelines by facilitating rich, fast, and effective interaction. The human thinks in terms of the analysis task, which is closely aligned with the interaction, and then looks at the visualized results. As a result, the user is more able to stay in the cognitive zone, as we will discuss shortly, even with multiple windows.

With this in mind, visualization design should avoid, as much as possible, menus or other actions that take the user outside of the frame of the task. Interactions should be direct manipulation or closely attached (e.g., with buttons in or on the window). This would include pull-down menus, which require the human to sort through and think about menu items.

3.2.3 Central Role of Interaction

Human-computer interaction is not a series of disjointed behaviors. And while the visual process is an important part of visualization, it is not the central role. Interaction has a rhythm, a cycle of give and take. Interaction is the process by which the computer and the human give and take and create knowledge. We will see an example of this when we will consider the van Wijk operational model in the next section.

3.2.4 Insulation of Reasoning Flow

One goal of intuitive visualization should be the facilitation of the flow of human reasoning. Once the human has focused cognitive resources in an area of interest, the visualization should not hamper the rhythm of reasoning until the human chooses to refocus resources elsewhere. This insulation can be achieved partially through direct interaction within context and intuitive computer-aided information discovery, as is discussed elsewhere in this section. But insulation would also encompass any interface process or rendering feature that has the potential to interfere rather than inform the reasoning flow.

Concomitant with the concept of reasoning flow is the idea of being "in the zone." This metaphor has been used in a myriad of ways, but we will use it to discuss the ebb and flow of reasoning combined with what Vygotsky called the "zone of proximal development." [19] (ZPD). The ZPD, simply defined, is a zone in which what the human wants to accomplish is just out of reach but can be attained with the right support -- in this case, an intuitive interface. A human in the zone is focused on a problem to solve, a question to answer, or an area of interest to explore. These tasks absorb energy and cognitive bandwidth, requiring the human to focus available resources in the process of the interaction. Any unrelated task -- finding a menu item, closing a superfluous pop-up window, trial-and-error text searches -- can penetrate the zone and interrupt the reasoning flow, which slows, changes the direction of, or halts the progression of the task at hand.

In the terminology of the HCM, being in the zone allows the human collaborator to reason without attentional or cognitive impediment. And further, when the task becomes crushingly complex or when the human exhibits functional fixedness in thought, the computer provides a scaffolding of support, presenting the information within relevant context, suggesting what may have been overlooked, and keeping relevant information present.

3.2.5 Intimate interaction

It is important that the interaction is so translucent to the human collaborator that the give and take which occurs in a successful collaboration is seamless. Entering the interaction should seem natural and obvious. The use of on-screen tools should not require additional cognitive focus – i.e. be usable without the human having to “think about it.” Intimate interaction deters attentional interference during the cognitive flow, and enables the reasoning process to move forward unabated.

When interaction is intimate, what the human should see and cognitively manipulate is not the tool being used or the method of interaction, but only the interaction itself. Intimate interaction is an important asset to flow insulation, and is supportive of the central role of interaction.

3.3 Search by Example & Search by Pattern

Searching for information has traditionally been done by stipulating a search term in a text box. But text boxes require humans to know exactly what to look for, as well as to stop where they are in the reasoning process to look for information. We would propose that search be done by graphically indicating an example to search by, or by drawing a bounding box around a pattern or organization of relational information. In either case, search is done more intuitively, and reasoning moves forward.

3.4 Creation and Analysis of Hypotheses

One extension of the knowledge building process that holds great potential for multi-modal visualization is in the creation of hypotheses. Hypothesis generation can be fraught with human cognitive bias, as humans are wont to seek out evidence that proves what they already believe or want to believe. Getting past these biases can be time-expensive and destructive to the process.

As Heuer described the process [20], hypothesis analysis starts with a list of hypotheses and a list of evidence that proves and disproves each one. As the human creates a list of hypotheses, the computer can aid in finding relevant evidence. From there, the computer, with its superior working memory, creates a weighted matrix which the human edits with superior reasoning. Using the edited matrix, the human draws conclusions about which hypotheses are correct, and if desired, sets up a data watch in the visualization which will notify the user of data changes.

Hypotheses generation is initiated by the human, but the computer plays a significant role in shortening the process, as contributing to more solid conclusion through use of its strengths.

4 CONSIDERING THE VAN WIJK MODEL

We will now discuss what the implications of our model would look like if integrated into the van Wijk operational model of the visualization process [21]. We do this primarily as another way to envision how a human cognition model interrelates with other aspects of visualization theory.

Van Wijk models the “user” as P (perception and cognition), K (knowledge), and E (interactive exploration), as is demonstrated in Figure 2. The user perceives the image and utilizes the specifications in exploration.

It is difficult, however, to separate “knowledge” from the reasoning process that created it. A person’s knowledge is not simply a compendium of declarative facts; it is also the

relational or inferential semantic meaning a person gives the facts, patterns of facts and their relationships, the perceived worth of those facts, and the ways in which facts are used to reason about the encounter of future novel information. Indeed, facts are useless without the reasoning power to manipulate them, and so we must believe that the ‘K’ submodel, must, by definition, include the cognition processes that created it.

Knowledge defines the methods used when new knowledge is integrated with the old. Van Wijk also seems to imply this in how he models his “user;” his model pictures Knowledge feeding Exploration. But K cannot inform E without the guiding focus of a reasoning process. Indeed, exploration itself is cognition in action.

With these thoughts in mind, it might truly be more representative if Perception, Knowledge, and Exploration were all modelled as cognitive process informing each other. We would see P as the early cognitive processes of selective attention, categorization, and accommodation, K as knowledge with the use of reasoning, problem-solving and other thought processes which allow the human to create knowledge, and E as a focused cognitive process utilizing both P and K.

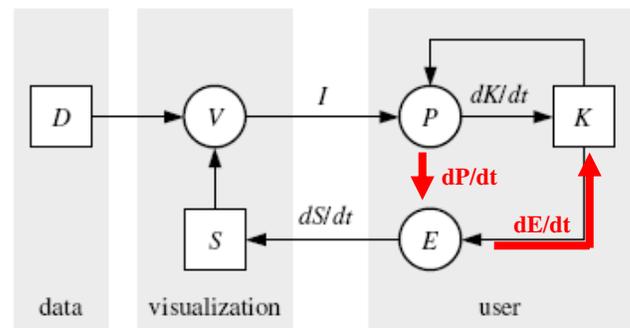


Figure 2. Van Wijk’s model with our integrations in red

When viewing the model this way, it’s easy to see that two additional directional arrows need to be added to the model: from P to E, and from E to K. The first arrow indicates the important role that “lower” cognition plays in active exploration

The second arrow signifies how a rhythm of interaction feeds knowledge reasoning. As the human explores and learns, that learning directs and focuses the attention of further exploration.

Additionally, van Wijk expressed Knowledge in this way:

$$K(t) = K_0 + \int_0^t P(I, K, t) dt \quad [21]$$

While this expression does encapsulate the idea that Perception is a vital part of the process, our integration would express the creation of knowledge over time more like the following:

$$K(t) = K_0 + \int_0^t E(P, K, t) dt$$

where Knowledge is the extension of currently held knowledge through the integrated perceptual and reasoning cognitive processes of Exploration.

5 EXAMPLES

In this section, we will demonstrate the model guidelines by using them to evaluate and/or illustrate several visual analytics designs. The model, which was not used as a basis for these designs, provides a deeper understanding of the choices made

and how the designs might be improved. Because we can discuss the rationale behind them more fully, we present predominantly designs in which we were developmentally involved. However, the arguments we make here would also apply to many other designs.

5.1 WireVis and Jigsaw

Discovering financial fraud in the great stream of transactions at a large bank is a difficult, time-consuming, and expensive process since it must employ expert analysts and uncover ever-changing modes of operation by criminals. The *WireVis* system was designed to combine the art of analysis with the science of visual analytics [22]. *WireVis* is an expert system, enhancing insight with what, in the terms of our Knowledge

expression in the last section, is presumed to be the human’s already sizeable K_0 ; it provides intuitive visual representations of wire transactions, enhancing P ; different views within the system allow the analysts to gain an overview of all transactions in one glance, while the ability to drill down into specific details of individual transactions enables finer examination and scrutiny. The main interface is shown in Figure 3. A time-based visualization allows the analysts to detect abnormal temporal patterns. Search by example permits selection of a particular keyword or time distribution pattern; the system then finds patterns that are similar to (or quite different from) the selected pattern. Finally, a keyword network shows keyword links for the selected group of transactions (where linked keywords appear in the same transaction), uncovering relationships that significantly aid the analyst in picking out suspicious transactions. This process highlights the importance of E in extending K . Results of a user evaluation found *WireVis* to be an efficient, effective tool, which can discover suspicious patterns in a great mass of transaction data [22]; the tool is also generally applicable to other types of transactional analysis.

WireVis has a number of capabilities that conform to the above cognitive model and highlights some of the design choices that must be made. Four windows are tailored to specific, important views and tasks. Though having a single window to focus the user’s attention may be ideal in some conceptual sense, and there presumably a cognitive and perceptual load during task switching, multiple windows seem necessary for many complex analytical problems [23, 24]. The key is to minimize the load in order to mitigate the interference to the human’s reasoning flow

In *WireVis*, the views were carefully chosen so that overviews of main aspects of financial analysis were maintained. Linking and brushing between all views was

enacted and immediate update to any interaction was enforced. (There are not only perceptual but cognitive aspects to maintaining high interactivity.) In addition, *WireVis* uses “balanced interaction”, signifying that interaction in one window is balanced by conceptually similar interactions the others. Thus

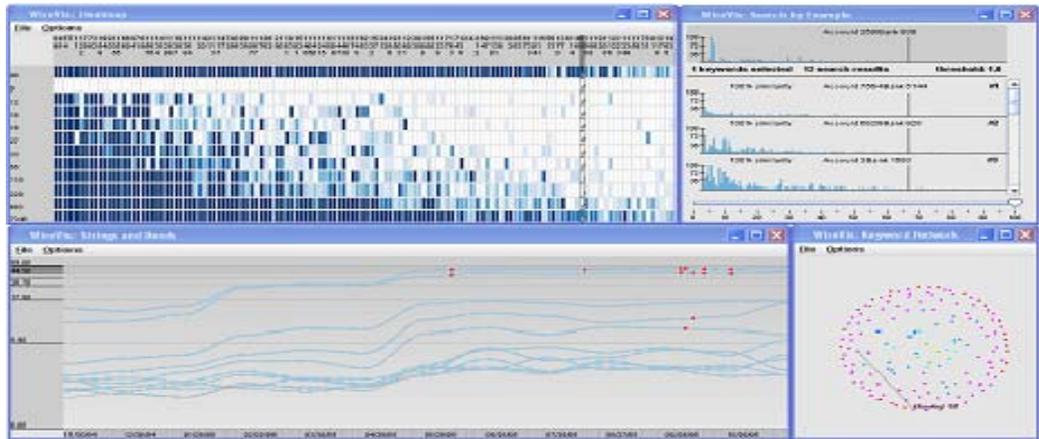


Figure 3. *WireVis* interface: keyword window (upper left), transaction temporal window (lower left), search by example

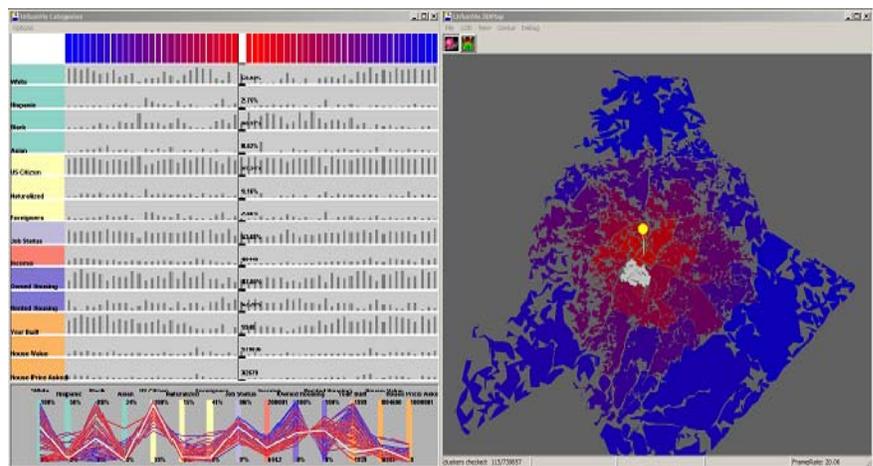


Figure 4. *UrbanVis*. Relational data view (upper left) Keyword-based Strings and Beads view (lower left) Geographic view with pointer that drives the visualization (right).

various selections, filter, and drill-down (through the transaction hierarchy of transactions) interactions appear in multiple visualizations. Further, very lightweight cursor passover interaction is enabled in several places (for example, passing over specific keywords). Finally, direct manipulation is used wherever possible to maintain user focus. As one of the papers co-authors remarked, “The interaction *is* the analysis” [25].

WireVis also has search by example, which has been singled out in our cognitive model because it is very general and it keeps the user in the context of her reasoning process without interrupting it to construct the appropriate search query, which can quite difficult to accomplish. In this case, the user select the keyword or transaction pattern she is thinking about to gather similar or dissimilar patterns. In fact, we believe that search by example should be part of any visual analytics interface involving analysis or reasoning tasks for large amounts of information. Search by example has been considered generally useful in other types of visualization, such as image analysis

[26], broadcast news event analysis [27], and terrorism event investigation [28]. In fact, we believe that search by example should be part of any visual analytics interface involving analysis or reasoning tasks for large amounts of information.

Jigsaw is a visual analytics system to support investigative analysis [23]. It works with large collections of reports or other text documents and with the entities extracted from them. Its two main goals are to permit investigators to handle efficiently and move quickly through large document collections and to support hypothesis formation and evidence gathering. *Jigsaw* won the university portion of the VAST 2007 Contest, which centered on a simulated investigation similar to those carried out by intelligence analysts.

As with *WireVis*, *Jigsaw* makes strong use of multiple windows with carefully tailored representations for complex investigative problems. The user is thought to be in an “information cockpit” with multiple monitors, in front of and above the user [23]. *Jigsaw* seeks to maximize pixel use to take advantage of both the user’s high acuity central focus and wide peripheral field. This is also a valid design point for *WireVis* (which requires at least two desktop monitors or a high resolution cinema display) or any other multi-window system.

However, although *Jigsaw* has some linking and brushing to integrate the windows, it does not have the balanced interaction *WireVis* employs. Based on the HCM guidelines, we would expect that users of *Jigsaw* would be less in the flow and require more cognitive effort than in *WireVis* during window management and connection. This is certainly a point worthy of further analysis and evaluation.

As a point of contrast, *Jigsaw* uses a bottom-up approach, employing an incremental, query-based method to bring in subsets of data for exploration and possible connection, as compared with *WireVis*’s top-down visualization of the whole dataset and its context. Undoubtedly both approaches are valid and could be available in a general tool for complex problem-solving, and will be the subject of future study.

Finally, *Jigsaw* is usable and simple. The interface permits direction interaction with representations of entities and reports, changing focus and detail. As with *WireVis* and other tools described here, simplicity and intuitiveness are not just goals based on perception principles but also based on the need for cognitive support. The HCM provides a point of view for investigating these goals in that light.

5.2 UrbanVis

UrbanVis is a system created to combine views of detailed geographical information with relational views of associated abstract information[29]. The geographical view is based on urban legibility theories in architecture, and the overall system permits the user to interactively explore an urban environment to understand the detailed characteristics of a city.

As with *WireVis* and *Jigsaw* and for the same general reasons, *UrbanVis* is highly dependent upon multiple views, with a 3D multiresolution map driving the updates of a straight category view and a parallel coordinates view, giving the user a rich overview of many categories and relations at once (Figure 4). These views were carefully chosen after consultation with architects, urban planners, and GIS experts. *UrbanVis* provides a general approach to reasoning with the combination of geographic and abstract data. In addition to the data shown in Figure 4, we are applying *UrbanVis* to bridge management data over city and state regions and are planning to use it for situation awareness in emergency response. This and the other examples

in Section 5 show the generality of a multi-window approach that is designed with principles of human cognition in mind.

Users of *UrbanVis* interact directly with the information, moving a geographic pointer and highlighting areas of interest or conversely choosing categories or individual coordinates in the parallel coordinates view to highlight specific geographic areas. This makes it easier to discover hidden geographical patterns for combinations of demographic or other abstract data. In the same sense as with *WireVis*, *UrbanVis* provides balanced interaction. This combined with direct manipulation, makes interaction the central focus. In addition, *UrbanVis* also provides a top-down, exploratory view with drill down controlled by simple movement of the ball up and down. The interface has only one menu, as it strives to keep the user cognitively focused during problem solving.

Finally, since *UrbanVis* utilizes the central role of interaction, the visualization makes E, as defined in Section 4, the seminal focus. In this interface, designed for a broad cross-section of users, K_0 can be small or great, and an attempt is made through use of color and spatial organization to facilitate P.

5.3 GVis

Although the human is uniquely qualified for higher order

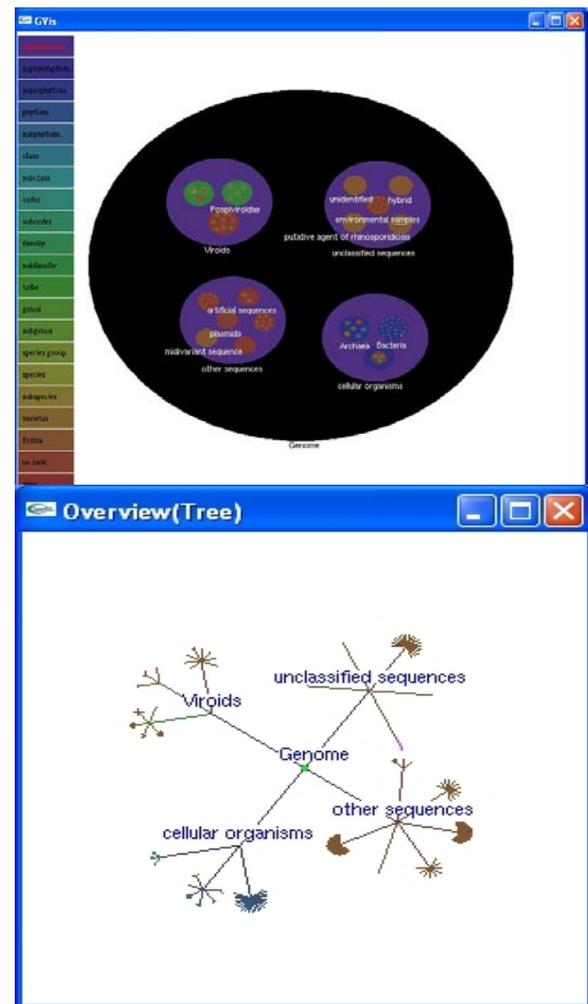


Figure 5: GVis Main navigation window(upper) and category tree view(lower).

reasoning, our human cognition model permits the computer to support this process in numerous ways. *GVis* in this section and *SRS* in the next provide some of this support. As several of the visualization approaches utilize similar methods, in these final sections we will highlight areas that are different from *WireVis*.

Using information available in public biological databases such as NCBI, *GVis* pictures the relationships and publications known about genomic data (Figure 5) [30]. Due to the detail inherent in genomic data, the amount of information presently viewable during drill down in direct interaction becomes quickly overwhelming; the use of multiple views to visualize multiple levels of information hierarchies prevents the human from “losing their place” in the taxonomy. Also, similarly to *WireVis*, *GVis* is an expert visualization, requiring a rather sizeable K_0 to focus effective Exploration. This may mitigate the cognitive overload effects of information on P and K in Figure 2.

A popup menu allows the user to view and explore publications on the spheres in the main view. This is perhaps not an optimum solution, as use of the menu is not intimate and can obstruct the field of view, which could threaten reasoning flow insularity and reduce the value of direct interaction.

The visualization uses color and simple circles to highlight groups, insulating reasoning flow and focusing P . In addition, it employs the notion of a stretchable canvas, similar to *Pad++* [31], to handle detail at all scales. The latest version of *GVis* applies the precepts of *knowledge visualization*, relying on taxonomic and ontological representations of genomic knowledge. When combined with the stretchable canvas, important knowledge at any scale can be made visible to support the task at hand. Thus, for example, glyphs showing how many genes are annotated (and what types) are made visible at the genome or even the cluster level, even though the individual genes are sub-pixel at these levels. This provides important support for the human reasoning process.

5.4 Scalable Reasoning System (SRS)

In *SRS*, Pike *et al.* demonstrate nicely the capacity of visual analytics to aid in hypothesis generation and analysis [32]. For example, while searching by example in *SRS* is limited to text searches, queried information can become “reasoning artifacts” to be used as the basis of hypotheses, or as evidence for hypotheses represented in “thought chains.” The human is free to manipulate these artifacts directly in a sandbox-like information space, which encourages reasoning flow insularity and focuses P as described in Section 4. Additionally, by allowing the human to arrange artifacts, interaction becomes the principle objective. While *SRS* does not use multiple views to display information, allowing rearrangement of artifacts encourages the human to organize them in a way that is meaningful to the individual.

Additionally, *SRS* is more than a display. New knowledge can be created by creating relationships between reasoning artifacts. Thus, as the human generates hypotheses and their evidence, new knowledge that is created during the process is not lost, either to the current analysis, or to all other humans who use *SRS* after its creation. These new relationships are given editable confidence ratings, which aids in the weighing of evidence in hypothesis analysis.

CONCLUSIONS

If we in the visual analytics community are to attain the aspiration of more effective, more human-perceptive

visualizations, we must begin to understand how humans manipulate semantically-complex information. It is not enough to understand what is being seen and given attention. Nor is it appropriate to infer combinatorial, individually-variable reasoning heuristics from binary, “lower level” cognitive behaviors; just as an understanding of perceptual cognition based on evaluative research have been employed in creation of information-valuable displays, “higher cognition”: reasoning, problem-solving, and decision-making must be employed in the development of mixed-initiative analytical interfaces.

What we have proposed is not a derived working model, but is, as yet, a framework of human cognition. Our goal is to sketch out a system of “thinking about thinking” as a first step to interface interaction which is no longer between user and software, but between human and computer partners, collaborating in the discovery of information and the creation and extension of knowledge.

What we have proposed is a model still in its rudimentary stages, whose future will undoubtedly be marked by additions, corrections, and multiple series of empirical evaluation. In some ways, this work still has the emergent expectation and general outline of an excavation; we cannot pretend to have all of the answers, but we know we’re digging at the right spot.

And yet, it doesn’t take a bulldozer to uncover that humans, like any other thinking system, behave in foreseeable ways. Thanks to decades of psychological research and some research in human-computer interaction, we aren’t starting from scratch or relying on anecdotal evidence; we have, at the very least, a thorough understanding of how the model is shaped and where to go from here.

We have been able to show extant examples of several of the HCM’s submodels, but there is still work to be done. There are, as yet, no spontaneous methods of searching by analogical structure. Computer-aided discovery will require both a better understanding of learning interfaces as well as a comprehensive understanding of human iterative thought chaining.

There is also the pesky problem of a unified theory of reasoning. Understanding the fragmented research that is available to us is a good foundation, but we must approach a more complete discernment of how all of the pieces work together in an information space to better be able to define and evaluate the best ways to insulate reasoning flow., mitigate cognitive load, facilitate appropriate task switching, and minimize attentional interference during the reasoning process.

Finally, there is the issue of better understanding the temporal coordination of human reasoning and computation and presentation of information. The temporal dynamics of control actions and cognitive processing were addressed early in the history of HCI with GOMS and keystroke analysis. However the dynamic coupling of scene gist perception, eye movements, attentional allocation, cognitive processing and microstrategic perception/action patterns [33] remains to be explored. Recent advances in the application of nonlinear dynamical modeling [34] may provide sufficient predictive validity for incorporation into models of sensemaking in visual analytics

These tasks are broad items on a bold agenda. But our evaluation has also uncovered multiple practical problems, and directed the search for how best to tackle them: what number of multiple views maximize the cognitive return on investment, the best way to suggest unconsidered information without interrupting – or annoying – the human at work, and methods of maintaining the interactive process so as to keep cognition in the flow, whatever the task. Finally, it is also clear that there must be strong support for permitting and managing human

annotations. But again, it all starts with daring to peek into reasoning's black box.

ACKNOWLEDGEMENTS

This work was performed with support from the National Visualization and Analytics Center (NVAC_{TM}), a U.S. Department of Homeland Security Program, under the auspices of the SouthEast Regional Visualization and Analytics Center. Special thanks to Benjamin F. Green for helpful insights.

REFERENCES

- [1] Previous work
- [2] A. Newell, You can't play 20 questions with nature and win: Projective comments on the papers of this symposium. In *Visual Information Processing*. Chase, W.G. (Ed). New York: Academic Press. (1973).
- [3] J. Piaget, "Piaget's theory" *Cognitive development to adolescence*, K. Richardson, and S. Sheldon, (Eds.) Erlbaum: Hillsdale, NJ, 3 – 18 (1988).
- [4] K.L. Komatsu, "Recent views of conceptual structure," *Psychological Bulletin*, 112, 500-526 (1992).
- [5] G. Gigerenzer, and D. G. Goldstein, Reasoning the fast and frugal way: Models of bounded rationality. *Psychological Review*, 103 (4) (1996).
- [6] I. Biederman, "Recognition by components: A theory of human image understanding," *Psychological Review*, 94 (2), p. 115 -147, (1987).
- [7] R.A. Rensink, A dynamic representation of scenes. *Visual Cognition*, 7, (2000).
- [8] Z. Pylyshyn, *Seeing and Visualizing: It's not what you think*, MIT Press/Bradford Books, Cambridge, MA: (2003).
- [9] Z. Pylyshyn, *Things and Places*. IT Press/Bradford Books, Cambridge, MA, (2007).
- [10] S. Ullman, Visual routines. *Cognition*, 18, p. 97-159, (1984).
- [11] I. Rock. *The logic of perception*. MIT Press/Bradford Books, Cambridge, MA, (1985).
- [12] M.J. Tarr, Learning to see faces and objects. *Trends in Cognitive Sciences*, 7(1), (2003).
- [13] J. Kozielicki, "Elements of a psychological decision theory," *Studia Psychologica*, 13(1), p. 53-60 (1971).
- [14] P. Cherubini, and A. Mazocco, "From models to rules: Mechanization of reasoning as a way to cope with cognitive overloading in combinatorial problems," *Acta Psychologica*, 16(3), p. 223-243 (2004).
- [15] G.A. Miller, "The magic number seven, plus or minus two: Some limits on our capacity for processing information," *Psychological Review*, 63, p. 81-97.
- [16] P.C. Wason, "On the failure to eliminate hypotheses in a conceptual task," *Quarterly Journal of Experimental Psychology*, 12, p. 129-140 (1960).
- [17] J.St. B.T. Evans, J. Varston, and P. Pollard, "On the conflict between logic and belief in syllogistic reasoning," *Memory and Cognition*, 11, p. 295-306 (1983).
- [18] M.Q. W. Baldonado, A. Woodruff, and A. Kuchinsky, "Guidelines for using multiple views in information visualization." In *Proceedings of Working Conference on Advanced Visual Interfaces*, Palermo, Italy, (2000).
- [19] L.S.Vygotsky, *Mind and society: The development of higher psychological processes*. Harvard University Press, Cambridge, MA. (1978).
- [20] R. J Heuer, *The Psychology of Intelligence Analysis Center for the Study of Intelligence*. CIA. (1999).
- [21] J.J. van Wijk, "The value of visualization." *IEEE Visualization, Proceedings of Vis 05*, p. 79 – 86. (2005).
- [22] R. Chang, M. Ghoniem, R. Kosara, W. Ribarsky, J. Yang, E. Suma, C. Ziemkiewicz, D. Kern, A. Sudjianto, "WireVis: Visualization of categorical, time-varying data from financial transactions." In *Proceedings of the 2007 IEEE Symposium on Visual Analytics Science and Technology* Sacramento, CA, October 2007, p. 155-162. IEEE Computer Society, 1995.
- [23] J. Stasko, C. Gorg, Z. Liu and K. Singhal, "Jigsaw: Supporting Investigative Analysis through Interactive Visualization", *Proceedings of 2007 IEEE Symposium on Visual Analytics Science and Technology*, Sacramento, CA, October 2007, pp. 131-138
- [24] J. A. Wise, J. J. Thomas, K. Pennock, D. Lantrip, M. Pottier, A. Schur, and V. Crow. "Visualizing the non-visual: spatial analysis and interaction with information from text documents." In *INFOVIS '95: Proceedings of the 1995 IEEE Symposium on Information V visualization*, page 51. IEEE Computer Society, 1995.
- [25] Private communication, Remco Chang.
- [26] J. Yang, J. Fan, D. Hubball, Y. Gao, H. Luo, W. Ribarsky, and M. Ward., "Semantic Image Browser: Bridging Information Visualization with Automated Intelligent Image Analysis,". *Proceedings of IEEE VAST 2006*, p. 191-198. (2006).
- [27] M. Ghoniem, D. Luo, J. Yang, and W. Ribarsky., "NewsLab: Exploratory Broadcast News Video Analysis,". *IEEE VAST 2007*, p. 123-130, (2007).
- [28] X. Wang, R. Chang, R. Kosara, W. Ribarsky, K. Smarick, and E. Miller, "Investigative Visual Analysis of Global Terrorism," (Accepted for publication). *EG/IEEE EuroVis 2008*. (2008)
- [29] R. Chang, G. Wessel, R. Kosara, E. Sauda, and W. Ribarsky, "Legible Cities: Focus-Dependent Multi-Resolution Visualization of Urban Relationships". *IEEE Transactions on Visualization and Computer Graphics (TVCG) InfoVis 2007*, Sacramento, CA, October 2007
- [30] J. Hong, D. H. Jeong, C. D. Shaw, W. Ribarsky, M. Borodovsky, and C. Song, GVis: A Scalable Visualization Framework for Genomic Data, " In *Proceedings of EuroVis 2005* p.191-198, (2005).
- [31] B.B. Bederson, J.D. Hollan, "Pad++: A zooming graphical interface for exploring alternate interface physics". In: *UIST '94*, p. 17-26, (1994)
- [32] A. P. Pike, R. May, B. Baddeley, R. Riensche, J. Bruce, and K. Younkin, "Scalable visual reasoning: supporting collaboration through distributed analysis," *Proceedings of International Symposium on Collaborative Technologies and Systems (IEEE)*, (2007).
- [33] W. D. Gray, and D. A. Boehm-Davis, "Milliseconds Matter: An introduction to microstrategies and to their use in describing and predicting interactive behavior." *Journal of Experiment Psychology: Applied*, 6(4), p. 322-335. (2000)
- [34] M. Spivey, "The Continuity of Mind" Oxford Press, New York, NY (2007).