The Visualization Pipeline is Broken

Chris North[1], Alex Endert[1,2], Christopher Andrews[1], Glenn Fink[2]

[1]The Center for Human Computer Interaction, Department of Computer Science, Virginia Tech
[2]Pacific Northwest National Laboratories

Abstract— Because of the emphasis on sensemaking and synthesis processes in visual analytics, interaction has evolved to take on a newer, deeper meaning in information visualization. Typically algorithmic data visualization and user-generated story synthesis tasks are separated in existing tools, but what if they could be better integrated? Malleable visualization is needed to enable analysts to interactively reorganize data visualizations into stories that capture deep semantics. However, the rigid mathematical nature of the visualization pipeline currently restricts visual representations and interactions, reducing their effectiveness – the visualization pipeline is broken. A new model is needed that better supports sensemaking and human thought. Can a visualization also provide methods for editing data, capturing thoughts, and setting the semantic framework chosen by analysts to complete their task successfully? How should user interactions in visual space propagate backwards up the pipeline through visual mappings and data space? Can we create such mappings that are useful, yet flexible? Or, is there a deeper problem stemming from theories of embodied interaction about the disconnect between human visual interaction and computational data processing? This broader view of visualization could give designers the ability to build new visualizations outside the traditional realm.

Index Terms—Visualization Pipeline, visual analytics, synthesis, embodied interaction

1 Introduction

Information visualization emphasizes the importance of showing the data to the user. As Ware describes in [1], "visualizations provide an ability to comprehend huge amounts of data". Tools have been developed that excel at this ability, processing large amounts of data and transforming it into an easy to "comprehend" visualization. These tools follow the model of the visualization pipeline, as shown in Figure 1.

In parallel, a collection of tools have been developed that focus on the semantics of sensemaking. They provide the user with an ability to represent their thoughts, insights, and hypotheses, whether by drawing or freely modelling.

In this paper, we analyze how visual analytics emphasizes these two types of tools: data visualization and sensemaking. The information, representations, and interactions associated with these tools comprise two domains: data and semantics. Ideally, these would be integrated, but the limitations of the current model of the visualization pipeline do not support this well.

The pipeline currently supports a limited class of user interactions in the data domain such as adjusting the parameters of the computational data transformation, visual mapping, and view transformations. But the pipeline inadequately supports the sensemaking interactions that belong primarily to the semantics domain - the area where users creatively tell stories. The analytic process involves both domains, but without adequate support for the semantics domain, users must mentally bridge the cognitive gap between the domains.

We conjecture that visualizations should support both domains resulting in visual models that directly reflect the way people think in both domains simultaneously. In this paper, we consider the intersection between the domains and illuminate new challenges for designers. Is it possible to create a model which binds these two domains?

2 Separate Domains: Data and Semantics

For the scope of this paper, we define the semantic domain as the sensemaking process and any mental processes occurring in conjunction. The data domain consists of the computations, transformations, algorithms, and raw data editing. Ideally, an analyst should be able to perform all the data and semantic transformations within a single visual framework. However, we found that visualizations typically only provide sufficient control over either the data or the semantics but not both. We witnessed analysts resorting to separate tools, workspaces, and representations to overcome this lack. Then they were forced to reintegrate their results manually without benefit of direct support by the visualization. Good visualizations of semantics do not support sensible manipulation of the data, nor do good data visualizations support meaningful representations of the semantics. Our research has revealed some examples of how certain tools force the user to choose between the data domain or the semantics domain (see Table 1).

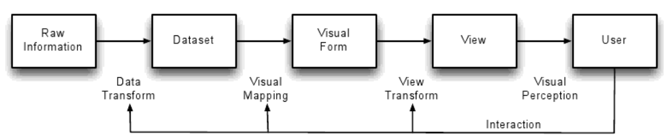

Figure 1 The Visualization Pipeline

2.1 The Data Domain

The visualization data pipeline is a model showing the series of events, calculations, and transformations that take place between the data and the user in order to achieve a visualization. It allows for certain types of interactions, such as changing data and visual transformations, as well as changing the visual mappings. This model can be seen in a set of visualizations thriving on these capabilities.

For example, the visual analytics tool IN-SPIRE [2] includes the ability to organize a large set of documents into a "Galaxy View" (a large spatial layout of the documents, represented by dots, algorithmically grouped into clusters). This visualization (Figure 2) gives the user a quick overview of a large collection of data. The galaxy view is created from the results of numerous keyword extraction and correlation algorithms. The series of steps propagate through the visualization pipeline producing the view into the data. Although the user is given a set of parameters to adjust, control over the presentation is limited and indirect.

Observations of analysts show the utility of these overviews. A user would follow the normal routine of opening the documents in the document viewer, reading through the ones that are of interest, taking notes if needed, then returning to the galaxy view for further searching. The current model supports all of those interactions well, mimicking the loops in the visualization pipeline. The overview gives a shallow level of context for where the documents lie in the

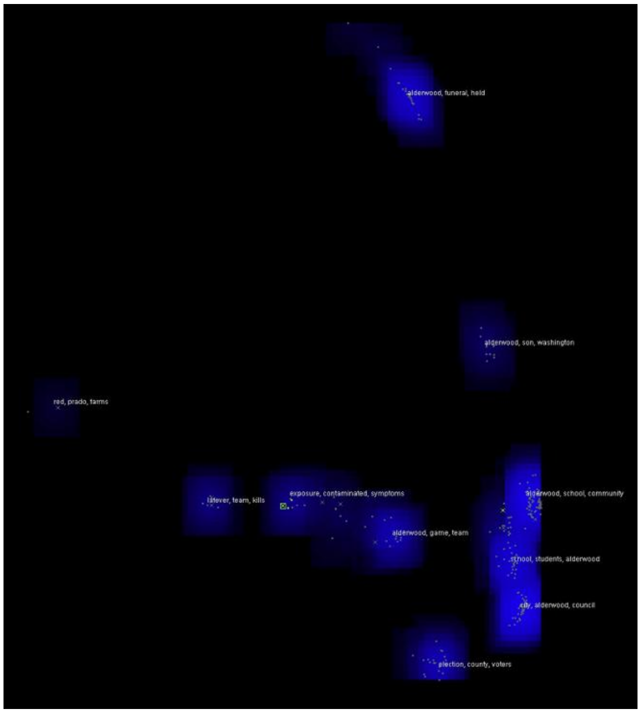

Figure 2 IN-SPIRE Galaxy View

overall document space. The user can interact directly with the mappings, transformations, or other aspect of the visualization. IN-SPIRE typifies visualization tools that specialize in semantic representation and manipulation to the exclusion of data manipulation.

Another commonly used tool in visual analytics is Spotfire [3]. Spotfire allows users to quickly visualize a set of data, then specify the type of representation they wish to view – for example, a scatter plot with visual encodings for size, shape, and color. Users can specify the axes to represent any value from the imported data, and scroll, zoom, and select detailed information on demand, as prescribed by the current visualization pipeline. Spotfire typifies the type of visualization that specializes in providing a wide variety of data domain manipulations while the semantics domain is almost entirely ignored.

2.2 Semantics and Sensemaking

After observing a visual representation of information, users begin to generate their own semantic connections within the data. They often do so by creating a story, a common way to understand and comprehend information [4]. Thus, in order to tell their story, they rely on the process of sensemaking. Sensemaking is described by Pirolli and Card [5]. An adapted version of their illustration (Figure 3) represents the analyst's mental process, which often cannot be computed or represented by a strict visualization algorithm or statistical structure.

The sensemaking loop demonstrates and supports the depth of interaction needed in visual analytics. At different stages an analysis task requires different kinds of interactions. Our study indicates that both bottom-up and top-down sensemaking processes are useful and desirable to users. An analyst must reason in both directions to create a story linking the raw data to a final report, possibly alternating between them numerous times.

Our study of analysts using visualizations showed that the subjects typically reached a point in their investigation where they require a means for recording their thoughts and hypotheses separate from the visualization. We observed analysts taking notes on a separate tool or even on their own notepads. Does this imply that these tools allowing a user to express their thoughts are based on a separate model.

Microsoft PowerPoint is the software package of choice when asking a user to present their thoughts. That being the case, we saw it used throughout the exploration of the data as a note taking tool, with the intent that these notes will be formalized and turned into a final report. Even as most of these tools attempt to provide the user with an environment suitable for exploring and investigating the data, users still turn to a presentation tools such as this when it comes to consolidating their thoughts into a concise bundle which others can understand. Observed users praise its ability to let them "do what [they] want" - the flexibility given for representing their thoughts is nearly endless.

Another tool user rely on heavily for externalizing their thoughts is Analyst's Notebook [6]. Although the feature set of this tool is far richer, it is often used as a mere drawing tool, with the ability to create entities, create links between them, and create note on the semantics behind the link. Thus, instead of importing the data through a series of algorithms provided for the sake of showing timeline, social network, or other visualization, users were often observed simply creating their own views. This simplified use of Analyst's Notebook proves valuable to the user, as they are able to store their possible hypotheses in a digital medium. Also, the freedom of placement and linking is brought to the foreground in a tool such as this, as there are hardly and restrictions for what can be an entity or a note on the link binding entities.

Sandbox, the sensemaking half of Occulus' nSpace suite [7] is another tool specifically designed for sensemaking. Sandbox provides the analyst with a flexible, open workspace where he or she can freely move data around, organizing, annotating, hypothesizing, etc. The space provides some specialized tools for activities like the analysis of competing hypotheses, but it primarily a free-form space that requires the analyst to provide meaning. In other words, Sandbox was designed to address the need for a purely semantic space for sensemaking.

We arrive at a conclusion that visualizations do not adequately allow analysts to manage their mental models of semantic connections nor do they enable them to inject their conjectures back into the visualizations in a natural way – thus breaking the current visualization pipeline. Since the visualization pipeline makes no provision for directly representing the relation between the semantics and the data, it forces users to switch between separate

Table 1 Characteristics of the Data Domain and the Semantic Domain.

Data Domain (Data Visualization)	Semantic Domain (Sensemaking)
• Visualization Pipeline model	• Drawing model
• Computationally-generated picture	• Analyst-generated picture
• External data	• Domain knowledge, hypotheses, insight
• Syntactic mappings	• Semantic spaces
• Representation = high-bandwidth view of data	• Representation = sensemaking space
• Interaction = adjust mapping parameters, multiple representations over time	• Interaction = sensemaking process, organize thoughts
• e.g. Galaxy View	• e.g. Sandbox, Analyst's Notebook

representations of the domains, increasing their mental workload. Compounding this problem is the need to refer back to the visualization, or even the raw data, to verify their recorded hypotheses. This requires users to simultaneously maintain multiple disjoint representations and the mental connections among them. By enabling rich interaction with the visualization and by allowing these interactions to propagate all the way back to the raw data, we aim to ease the frustration caused by the switching among different tools and representations. The user's need for a deeper interaction becomes clear – there must be a way to link the semantic workspace with the data visualization.

3 Fusing the Data and Semantic Domains

In the course of an analysis, analysts routinely work in both the data domain and the semantic domain. To the analyst process should be seamless – data-driven visualizations are fundamentally input to the semantically driven sensemaking process. However, this is not well supported by the available tool chains.

An important stage of the analytic process is when the analyst has to piece together the collection of thoughts and data that support their hypotheses – fusing the raw information with their sensemaking. All of their work on a scenario culminates into a form of written report which is passed on to the appropriate personnel. This report, which has many different forms, has to encompass many different types of information. These include, but are not limited to, a trail of evidence and how the links between them were established, background information the analyst knew before starting the task, sources of outside information pulled in through many different mediums, and visualizations highlighting their findings to provide a quick overview (see Table 2).

Creating this mosaic of semantics and data is challenging and difficult to the user, as their thoughts and the corresponding data are separated into two domains. First, an analyst is given a mathematically generated overview (e.g. galaxy view, scatter plot) that can be explored using operations like selection and filtering. In parallel, the analyst's thoughts regarding the information they are given are kept either in note form or to themselves. Finally, as the analysis draws to a close, the analyst is taxed with creating a conglomerate of data and thoughts into a coherent story. Many analysts vented their frustration with this step of the process, as their thoughts, notes, hypotheses, and evidence are strewn about these different tools – highlighting the rift between sensemaking and the visualization pipeline. Can we streamline this process?

How can the visualization pipeline adapt to the growing desire of users to interact deeply with a visualization? How can we bridge this gap between interactions tailored for visualizations and those corresponding to sensemaking?

3.1 Interpreting the Interaction

Visual analytics relies heavily on user interactions. However, these interactions go beyond the "control language" used within the visualization to make changes to the view. Interaction is also a way for analysts to offload their thought processes.

To illustrate this, we can consider Robinson's study of analysts performing synthesis with physical artifacts [8]. Analysts were provided with evidence separated out onto individual note cards and a space to lay them out. Robinson observed that the analysts used the space to create spatial metaphors that expressed various relationships between the pieces of evidence. In other words, they were manipulating objects within their workspace to help them to make sense of them. While this was done physically, it is not difficult to envision these same actions being taken within an application.

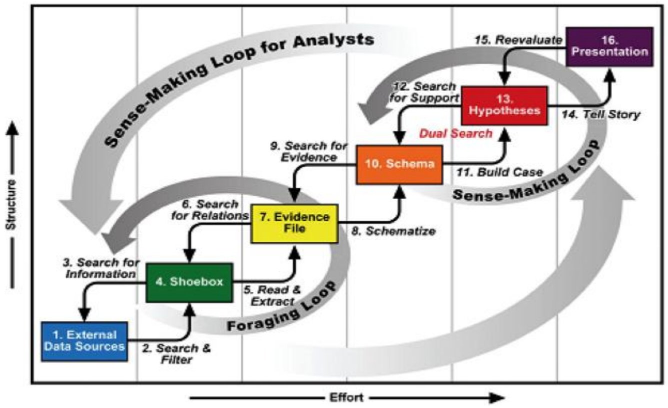

Figure 3 The Sensemaking Loop

A key problem lies in the interpretation of these actions. Controlling actions are easy to interpret – they are direct commands. It is more difficult to interpret an interaction made as part of the thought process. The context of these interactions extends beyond the confines of the data and includes all of the analyst's current tasks and experiences. The semantics and intent behind each action is likely to be highly individualistic and difficult to divine.

The problem of interpreting these semantic interactions can be largely ignored in systems dedicated to supporting this kind of work. The tool needs merely to provide the tools to permit the action and the interpretation is left to the user. However, if the goal is to fuse the data and semantic domains, the visualized representations have links back to raw data and actions taken by the user will have to be considered, lest the visualization lose its link back to the data.

3.1.1 Example: Changing a Cluster

In our observations of users interacting with Galaxy View, we noticed a desire to select data points and drag them to new locations in the space. While this seems like a relatively straightforward action to the user, the interpretation of the action is not.

The decision to move a document within the space typically implies one of two goals: (1) Placing particular pieces of data near each other to illustrate a point when the two are juxtaposed to better present a story, or (2) attempting to correct the visualization's pre-defined clustering algorithm by placing a document in the "correct" cluster, with hope that the visualization will correctly group these documents in the future.

Provided the action was supported, how should it be handled by the underlying system? Each point is placed based on a mathematical mapping from the data to the space. Changing the position of a point breaks that mapping. However, the change does not break the visualization – presumably, it is now more informative from the perspective of the analyst. This is because the actual mapping

Table 2 Example Characteristics of Fusing the Data Doman and the Semantic Domain.

Fusing Data and Semantic Domains

- Unifies data visualization & sensemaking
- Picture iteratively generated cooperatively by analyst and system; computationally-generated visualization learns and responds to analyst's reorganization
- Domain knowledge, hypotheses, and insight, all grounded in external data
- Transform syntactic mappings into semantic spaces
- Representation = Data-rich sensemaking space; hypotheses containing source data
- Interaction = Inject domain knowledge into data, and evaluate alternative hypotheses, by directly manipulating data in the visual representation

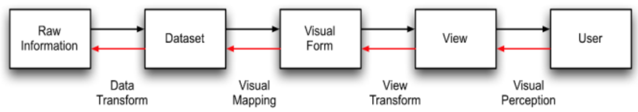

Figure 4 The Visualization Pipeline allowing bidirectional interactions – one possible response to users' demands

algorithm is not an important part of understanding the visualization. The clustering algorithms are using metrics such as word frequency and co-occurrence to create similarity measurements that are then mapped into 2D space. However, this is merely an approximation of how the space is perceived – as a space in which proximity indicates some form of similarity or relationship. In this context, it makes perfect sense to talk about moving dots around so that the space makes more sense.

The question is how the system should react to the change. The simplest reaction would be to do nothing – to behave as if the initial configuration is just provided as a starting point. However, this means that the first interaction breaks the connection back to the data. This leaves us with the problem of interpreting the intention of the user and feeding it back into the system. If the intention was to improve the clustering, then presumably this should be input back into the clustering algorithm and the whole space, which may potentially re-evaluate the entire space. A further challenge is that a clustering that made logical sense to the analyst may not be reflected in aspect of the data to which the algorithm has access. If the intent was to tell a story, then re-evaluating the clustering algorithm may not be the appropriate action. Instead, it might be appropriate to identify other documents that show similar links, again presuming that the link can be identified.

The goal of interpreting the user's interactions is to better inject the reasoning behind the interaction back into the visualization somehow. Instead of the analyst directly manipulating the algorithms, they would like to simply inject their thoughts into the visualization through a natural interaction, and have the corresponding events propagate down the pipeline, creating the updated mappings along the way. The current pipeline does not allow for such a level, or depth, of interaction. Although changing of mappings, transformations, and filters are currently often supported, observations have shown that the method for accomplishing these changes are clumsy, and do not mesh with the semantic process.

3.2 Tools Beginning to Fuse Semantics and Data

Some tools have attempted to provide a bridge between the semantics and data domains. These tools approach this challenge each from a different angle, directing the outcome of the tool's functionality each in a slightly different heading. At the heart of these is an attempt at providing a workspace for a user to express their thoughts freely; have it be through an organizational layout, a series of interactions, or direct data manipulation.

GeoTime Stories [9] is another prototype seeking to capture the analytic process. It breaks up into three main components: a story window, a visualization and annotation system, and a pattern matching system. These components are tailored to allow a user to tell a story, which provides a user with a way of capturing, organizing, and sharing the complex information which they have gathered [9].

When it comes to finally preparing a report, Microsoft PowerPoint is the software package of choice. That being the case, we saw it used throughout the exploration of the data as a notetaking tool, with the intent that these notes will be formalized and turned into the final report. Even as most of these tools attempt to provide the user with an environment suitable for exploring and investigating the data, users still turn to a presentation tools such as this when it comes to consolidating their thoughts into a concise bundle which others can understand. Observed users praise its ability to let them "do what [they] want" - the flexibility given for representing their thoughts is nearly endless.

The Fused Analytic Desktop Environment (FADE) is a domain-agnostic analysis software suite developed at Pacific Northwest National Laboratory (PNNL) to provide analysts a wide variety of information-driven fields with an integrated set of information analysis technologies**.** A single interface enables the tools to access a common repository of user-organized data modeled after a traditional PC desktop and to process or display the data to support different analytical purposes.

FADE supports *ad hoc* data sources where information is not specified *a priori,* but rather the analyst identifies the appropriate data to be used for each analysis task. FADE's Universal Parsing Agent (UPA) [10] automates the data domain extraction and manipulation to support activities in the data domain. FADE's Concept-Based Clustering (CBC) technology enables each user to create personalized folder organizations into which the system will automatically categorize documents. After placing several example documents in a user-created folder hierarchy CBC learns from the placement of those documents and automatically places new documents with similar themes into corresponding folders. Using a statistical algorithm that associates word usage with individual categories, CBC can rapidly categorize new documents without the user reading them, providing a first-cut document triage.FADE then provides a set of semantic tools (including IN-SPIRE, Frame-of-Reference Visualization, and the Analyst-driven Knowledge Enhancement and Analysis tools) to allow the user to see an overview story of the whole data set or parts of it. FADE does not specialize in enabling the user to generate his own stories or hypotheses, but does come close to providing a full suite of analytical tools to support both the data and semantics domains. Observing these set of tools in use, it becomes clear that the merging of these two domains is one possibility, although it may still not be enough. There are further possibilities when modelling this new type of visualization.

4 Examining the Pipeline

As we attempt to fuse the semantic and the data domains together, it becomes clear that the visualization pipeline cannot support interactions at the necessary level. The pipeline is data driven and designed to support user control. There is no support for pumping information back into the pipeline from the user's end. This is further complicated by a model in which changes are expressed by direct manipulations of the view, rather than by explicitly addressing the various stages of the pipeline to control it. In this section we will discuss a number of possible models to support this change in usage scenarios.

4.1 The Bi-Directional Pipeline

The first approach is to augment the visualization pipeline to support back propagation (Figure 4). Expanding the visualization pipeline to include means for propagating back to the data changes the emphasis placed on an interaction. This could imply a variety of changes. For example, a change to the visualization might result in a transformation in the algorithm performing the data transformation, such as the clustering algorithm discussed previously. A change, like lengthening a bar in a bar chart, could feed back and alter the actual data so that it reflected that altered visualization.

The VITE system is an example of this later form of feeding back into the pipeline [11]. VITE uses a two-way visual mapping controlled by a central mapping engine. Users define how various visual forms such as color, shape, and position are mapped to attributes of their data, much like other visualization tools like Spotfire. However, the user can directly change the visual encodings of individual pieces of data within the visualization, causing the raw data values to change. Interestingly, a number of users discovered that some visual encodings could be left unmapped, thus freeing it to be used in ways not directly linked to the data. In other words, they were assigning a semantic meaning to the encoding that could not be encompassed by the strictly data based system.

For less direct visualizations, feeding back into the pipeline becomes progressively more complex. The model to support this would have to include the ability to take the inverse of the functions currently driving the pipeline. These mapping (depicted in black in Figure 4) functions, however, do not necessarily have a working inverse function each time. Without being able to generate one of these functions, how can we propagate the information back to the raw data correctly, while creating the correct mappings and functions along the way?

Taking into account all the possible variations of a user's reason behind grouping a set of elements, this challenge is daunting. The solution may very well not exist in the traditional mathematical realm. When analyzing a user created cluster, for example, there may not exist any statistically significant correlation between the grouping and the documents within the group.

4.2 Modeling the Connection between Data and Semantics

Creating a bidirectional visualization pipeline that enables a reverse mapping of user interactions backwards through the view, visual, and data mappings, is a good first step. However, there are additional fundamental problems that break this bidirectional pipeline.

The disconnect between the semantic and data domain is apparent. The mapping that links these two domains is not clear, yet important. When analysts view and interact with visualizations they are inherently deriving and contributing semantic meaning. The system on the other hand is operating on data, performing computational transformations, oblivious to subtle semantics from the user.

The domain knowledge an analyst brings with them is often critical in their task. At times, they create constructs that coincide with this background knowledge. These constructs can include grouping data corresponding to information that carries a particular meaning to them, or perhaps reminds them of a scenario they have seen before. How can the visualization pipeline be augmented to accommodate this type of information?

The concept of intention is important. In terms of visual mappings, the mapping is syntactic, but what were the semantic intentions of its designer? What meaning was the designer attempting to convey with the mapping? E.g. the Galaxy visualization uses a series of syntactic data operations to generate a starfield of documents, but its designers intent was to capture the essence of 'themes' in the document collection, sometimes successfully so and sometimes not. Likewise, in terms of interactions, the analyst attempts to perform semantic operations, but ultimately these are translated into syntactic data edits or parameter changes on mappings. E.g. the analyst's intention is to position a lynch-pin piece of evidence into a critical whole in the story line, based on her domain expertise and knowledge about semantic relationships, not merely to specify x, y coordinates of data object.

A similar problem is the ambiguity of the relationship between the visual world and the data world. The user is operating primarily in the visual world, while the system is operating primarily in the data world (although clearly assisting in the mappings between). Together these problems creates a fundamental disconnect in the visualization pipeline (see Figure 5).

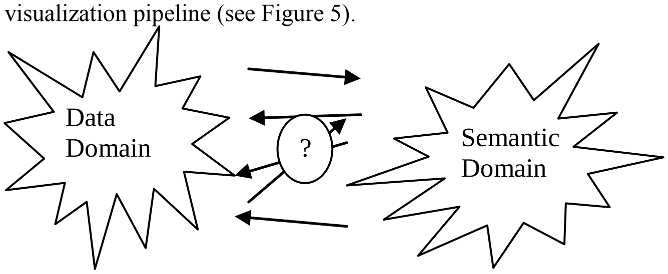

Figure 5 The Disconnect between Data and Semantic Domains.

This makes difficult problems in the mathematical modelling of a new pipeline to support the fusion of data and semantic worlds. Determining how the interactions in the semantic domain correlate to the calculations in the data domain can bridge the gap, but how?

The mappings are not necessarily purely deterministic, and carry with them an inherent level of uncertainty. Also, they may not be consistent between or within types of interactions. The mappings an analyst creates during sensemaking are not complete, but often only partial.

Thus, potential avenues for developing new visualization models might in involve probabilistic methods or partial evaluation of mapping functions. Another possibility is to refocus the computational model of the data domain to the visual domain, and operate on a spatial basis, perhaps exploiting spatial data structures and algorithms or vision based systems.

4.3 Towards New Models

A user's quest for more control over their visualization leads to a new opportunity in the realm of mathematical mappings. This inherent level of uncertainty of mapping an interaction which so heavily relies on subjective semantics can pose new challenges for researchers in other disciplines such statisticians.

4.3.1 Example: Bayesian Analysis

One possible approach is to turn to Bayesian Analysis. For example, analysts' interaction to reorganize a visualization can be thought of as a form of domain knowledge input by the user. Bayesian models could be used in visual mappings to essentially support an inverted mapping between domain knowledge and prior distribution specifications. Bayesian models could also support uncertainty to capture the ambiguity of analysts' semantic intentions.

The major challenge is mapping the Bayesian model concepts into elements of the visual representation, and mapping interactive operations back into the Bayesian model. Of course, this is not a panacea solution. These models still must make inferences about intention, but it gives the analyst more power in expressing their intensions. This could lead to greater burden on the user to specify these parameters, but the goal would be to map these as cleanly to the targeted domain visualization and interaction as possible.

In the case of the Galaxy example, document movement operations, representing semantic domain knowledge, could be mapped back into the statistical model for update in the form of prior distributions. The analyst could specify the strengths of their beliefs using various interactive controls. The enhanced model could then update the display based on posterior distributions, perhaps revealing a new clustering based on the users input. Probability values could be mapped into the visual representation to provide valuable feedback to the analyst in the form of a quantitative measure of uncertainty of the mappings.

This leads to another challenge: deciding how much control over the Bayesian model a user should have, and how direct it should be. This level of control can range from being *implicit* to *explicit*. A user can be given *implicit control* based on their normal interactions with the data, e.g. using default values or spatial metrics. Also, a user can have *explicit control*, e.g. sliders or other representations which allow for the changing of values. For example, the Bayesian strength metric in the Galaxy could be expressed as spatial proximity or a radius around a dot. A likely approach would enable the novice user to start with a set of default implicit controls, leaving the opportunity for the user to gain experience and smoothly transition from implicit to explicit modes of control.

Admittedly such models introduce design challenges. Bayesian concepts are notoriously difficult for novices to comprehend. However, mapping them into the targeted domain-specific visualization that is appropriate for the given analytic task could provide the necessary semantics to cast the Bayesian concepts and operations in a form that is understandable by analysts, and therefore support analysts' natural sensemaking process. For example, in the

Galaxy, the Bayesian strength metric could be expressed as spatial proximity or a radius around a dot.

A potential advantage of the Bayesian approach is that it support imprecision in the analysts' interaction. That is, in some cases analysts may be uncertain about hypotheses or simply might not care about exact placement. For example, feedback could be provided about possible alternatives placements that improve the fit.

This possibility of receiving feedback from the system allows for an opportunity for the user to gain confidence. Using the earlier example of dragging a data point in the IN-SPIRE galaxy, as the user drags the data and contemplates on a place to put it, the system should provide feedback on the potential intermediate effects on choosing a particular location. Hence, the user would still have the ability to position where the data where they please, but also receives feedback from the system about how well it is capturing their intention.

4.3.2 Data-Centric Models

A good example of past work that focuses on data models is Visage and CoMotion by Roth [12]. It gives the user the ability to directly manipulate data, such as drag data between different, user-determined views, cutting down on the often tedious task of importing and exporting data from one tool to another. Thus, the user is able to setup a workspace, consisting of a number of views, each aiding in representing a piece of their story. More importantly, this supports the analytic task of synthesis, in which the analyst must gather data from diverse visual sources and integrate into a combined view. The data in this combined view is still live, in that it continues to update with dynamic streaming data, and links back to original data sources and relationships, but now also indicates relationships between data that were previously disconnected. This is a good example of how analysts want to mold their working environment, and get their data "out of the window".

This requires a very flexible data model that can support integration of heterogeneous data on the fly, and support integrated visualization of them.

4.4 Encodings Supporting Semantics

Can we approach this challenge by creating visual encodings which inherently carry meaning? Currently, data is an intrinsic part of a visual encoding. Whether referring to color, size, shape, or other means of visually encoding information, the *information* encoded refers to data – not semantics. We can attempt to move the semantics into the visual space by providing the analyst with encodings that have the ability to display more than just data values.

One possible approach is to augment the encoding to account for the connection between the encoding and the user. As a user interacts with the encoding, can it slightly alter its appearance to account for the meaning behind the user's semantic intention? This goes beyond a simple change in color of a document that has been read and one that has not. Further development of such encodings lead to changing of the size of the glyph, altering its shape, or even creating complex glyphs comprised of many basic encodings – which when combined in a particular way, hold a semantic meaning to the user.

As can be seen in Figure 6, when a user interacts with a workspace, they change the layout of a standard visualization tremendously. What used to resemble a timeline (Figure 6a) turned into a workspace which does not have a particular mathematical layout, yet proves much more valuable and usable to the user. Treating the user-generated clusters as encodings, we can see how they have taken on a slight personal semantic meaning. When taking a closer look at Figure 6b, the slight distinctions from one cluster to another become apparent: some are neatly organized, some are stacked on top of each other, some are aligned in some fashion, and others take on a particular location on the workspace. Can further investigation of these clustering techniques give some insight of how semantic meaning can be incorporated into visual encodings?

In investigating these possible encodings, one should be aware that certain tradeoffs occur. As the level of personal meaning tied into the encodings increases, the more the visualization as a whole takes on meaning to the individual user. Without a model for translating these encodings to be meaningful for others, they become quite useless when taken out of the hands of the original creator. So how can we model these types of interactions?

5 Platforms for Deeper Interaction

These new deeper forms of interaction enabled by new interactive visualization models emphasize directly manipulating large quantities of information in new ways. These new models will place interaction at the heart of the visualization model. It is possible that interaction could take primacy over visual representations. In a sense, visual representations will become emergent spaces, brought to being through deep interaction and perhaps collaborative interaction. Thus, computational emphasis is placed less on the representation mappings and more on the way that analysts iteratively form the mappings in an emergent interactive sensemaking process. To accomplish this, analysts will need rich new spaces that afford such deep interaction. This will require likely require fundamentally new interaction techniques and technologies. Here we discuss two examples of such key enabling technologies.

5.1 Large High-Resolution Displays

Large high-resolution displays offer several key enabling capabilities for this type of deep interaction:

- They provide 'space to think'. Embodied Interaction and Distributed Cognition theory suggests that interaction is fundamental to cognition, and space plays a crucial role in cognition. For example, analysts use space to organize their thoughts in many different ways during the sensemaking process [8].
- Interaction is itself space consuming. When analysts reorganize information, whitespace is needed during the reorganization process for temporary arrangements. If many interactions are involved, interaction must be efficient. This suggests larger interaction targets that can be directly manipulated. For example, tiny pixel sized dots in an INSPIRE Galaxy may be too difficult to manipulate.
- Larger space provides greater visual opportunity for directly embedding or integrating data visualization and sensemaking tools [13], and therefore better enables the fusion of the data and semantic domains and the interaction between them.
- They afford new types of interaction in the form of *physical navigation*, that offers improved performance for navigating and understanding large quantities of information [14].

5.2 Multi-Touch Interaction

Multi-touch interaction technologies, such as the Microsoft Surface, offer two important enabling capabilities:

- It offers significantly greater efficiency of interaction when manipulating larger numbers of visual data objects [15], since it simultaneously exploits greater degrees of freedom (10 fingers and 2 hands) and more natural interaction. Analysts can exploit faster, less precise ways to produce layouts they desire.
- It also offers the opportunity for deep co-located collaborative interaction that exploit the benefits of emergent sensemaking [16].

5.3 A Futuristic Example

Figure 6 (a) A mathematically generated spatial layout as the task begins, showing titles of documents arranged in a clustered timeline and (b) the final spatial layout created by the user fusing both data and semantics.

Figure 6 demonstrates a potential example, based on scaling up the INSPIRE Galaxy into a new visualization with fused data and semantic domains. Consider a workspace designed to provide a visualization not simply as a rigid data visualization, but rather as a foundation upon which semantics may be added. As seen by the screenshots, the user has the ability to use the tool as both a visualization as well as a semantic workspace. Note taking, re-grouping, highlighting, and other sensemaking interactions occur directly within the visualization, to ultimately tell the story of the data. Documents are viewed as titles or as full text within the space, including highlights and annotations in a large space. Analysts can grab, push, shove groups of documents around in the space to organize thoughts, explore alternative hypotheses, and settle on a storyline.

The starting point of the workspace, Figure 6a, is a timeline of documents. From there, the analyst can reorganize to inject domain knowledge. Looking at the final result, Figure 6b, the basic outline of their organizational constructs used can be seen. For example, clusters can be seen represented in different ways. At times, simply tossing documents near each other was adequate for the task. Other times, these clusters had a more formal structure within them. Instances where these groups had documents placed in a line, sometimes preserving the temporal order of the documents in the timeline, and other times 'fudging' the order to represent which the documents were read first. Beyond just editing a computationally arranged layout, the system responds to the analysts organizations to organize other related data accordingly, show newly related themes, or point out potential supporting or missing data in the story.

6 Conclusion

Visualization is an important sensemaking tool. However, there is a disconnect between the data domain of pure data visualization and the semantic domain in which sensemaking takes place. In the semantic domain, interaction is about meaning rather than control. Analysts need the freedom to manipulate the environment, creating complex organizational structures that are transformed into schemas and hypotheses. Most tools treat these as completely separate domains and leave it to the users to bridge the gaps themselves. Fusing the data and semantic domains within a common environment offers new visualization and interaction opportunities. The addition of semantic content to data-driven visualization forces us to reconsider the architecture of the visualization pipeline, which is currently unable to support semantic input from the user. In light of this new model for visualization, the next step is building such visualizations. A new opportunity space has emerged, and the challenge is now to explore it. There is ample opportunity for research in this new and exciting area, and this work only touches on a few landmarks. The challenges of creating a visualization that supports this new pipeline are both plentiful and difficult. However, the possible opportunities for researchers are tremendous.

References

[1] Ware, C.: 'Information visualization : perception for design' (Morgan Kaufman, 2004. 2004).

[2] Pak Chung, W., Hetzler, B., Posse, C., Whiting, M., Havre, S., Cramer, N., Anuj, S., Singhal, M., Turner, A., and Thomas, J.: 'IN-SPIRE InfoVis 2004 Contest Entry', in 'Book IN-SPIRE InfoVis 2004 Contest Entry' (2004, edn.).

[3] Ahlberg, C.: 'Spotfire: an information exploration environment', SIGMOD Rec., 1996, 25, (4), pp. 25-29

[4] Fisher, W.R.: 'The Narrative Paradigm: In the Beginning', The Journal of Communication, 1985, 35, (4), pp. 74-89.

[5] Pirolli, P., and Card, S.: 'Sensemaking Processes of Intelligence Analysts and Possible Leverage Points as Identified Though Cognitive Task Analysis ', Proceedings of the 2005 International Conference on Intelligence Analysis, McLean, Virginia, 2005, pp. 6

[6] http://www.i2inc.com/products/analysts_notebook/, accessed 4/28/2009 2009

[7] Wright, W., Schroh, D., Proulx, P., Skaburskis, A., and Cort, B.: 'The Sandbox for analysis: concepts and methods'. Proc. CHI '06, New York, NY2006.

[8] Robinson, A.C.: 'Design for Synthesis in Geovisualization'. PhD thesis, Pennsylvannia State University, 2008.

[9] Eccles, R., Kapler, T., Harper, R., and Wright, W.: 'Stories in GeoTime', Information Visualization, 2008, 7, (1), pp. 3-17.

[10] Whiting, M.A., Cowley, W., Cramer, N., Gibson, A., Hohimer, R., Scott, R., and Tratz, S.: 'Enabling massive scale document

transformation for the semantic web: the universal parsing agent'. Proc. Proceedings of the 2005 ACM symposium on Document engineering, Bristol, United Kingdom2005.
[11] Hsieh, H., and Shipman, F.M.: 'Manipulating structured information in a visual workspace'. Proc. UIST '02, New York, NY, USA2002.
[12] Roth, S.F., Lucas, P., Senn, J.A., Gomberg, C.C., Burks, M.B., Stroffolino, P.J., Kolojechick, A.J., and Dunmire, C.: 'Visage: a user interface environment for exploring information', in: 'Book Visage: a user interface environment for exploring information' (1996, edn.), pp. 3-12, 116
[13] Stasko, J., Carsten, G., and Liu, Z.: 'Jigsaw: supporting investigative analysis through interactive visualization', Information Visualization, 2008, 7, (2), pp. 118-132
[14] Ball, R., North, C., and A. Bowman, D.: 'Move to improve: promoting physical navigation to increase user performance with large displays'. Proc. Proceedings of the SIGCHI conference on Human factors in computing systems, San Jose, California, USA2007.
[15] North, C., Tim, D., Isenberg, P., Fisher, D., Lee, B., Robertson, G., and Quinn, K.I.: 'Understanding Multi-touch Manipulation for Surface Computing', Proc. IFIP INTERACT 2009, 2009, to appear Sept. 2009, pp. 8
[16] Isenberg, P., Tang, A., and Carpendale, S.: 'An exploratory study of visual information analysis'. Proc. Proceeding of the twenty-sixth annual SIGCHI conference on Human factors in computing systems, Florence, Italy2008.