

Scalable Visual Queries for Data Exploration on Large, High-Resolution 3D Displays

Khairi Reda*, Andrew Johnson*[‡], Victor Mateevitsi*, Catherine Offord[†] and Jason Leigh*

* Dept. of Computer Science, University of Illinois at Chicago, Chicago, Illinois, USA

[†] Dept. of Ecology and Evolutionary Biology, Princeton University, Princeton, New Jersey, USA

{mreda2, vmatee2, spiff}@uic.edu, [‡]aej@evl.uic.edu, cofford@princeton.edu

Abstract—As the scale and complexity of data continue to grow at unprecedented rates, scientists are increasingly relying on Large, High-Resolution Displays to visualize and analyze scientific datasets. Recent studies have demonstrated the effectiveness of these displays in supporting cognitively demanding data analysis and sensemaking tasks. While there has been an abundance of research on rendering algorithms for large, high-resolution displays, far less effort has gone into designing interactive visual analytic interfaces to effectively leverage these displays in visual exploration and sensemaking scenarios involving large collections of data. In this paper, we present an interactive visual analytics application for the exploration of large trajectory datasets. Our application utilizes large, high-resolution 3D display environments to simultaneously visualize and juxtapose a large number of trajectories. It also integrates a scalable visual query technique, which can be used to quickly formulate and verify hypotheses, encouraging scientists to contemplate multiple competing theories before drawing conclusions. We evaluate our design within the context of a behavioral ecology case study. We also share our observations from a pilot user study to provide insights on how scientists might utilize large display environments in visual exploration and sensemaking scenarios.

Keywords—Large, High-Resolution displays; visual exploration; sensemaking; trajectory analysis

I. INTRODUCTION

As the scale and complexity of scientific data continue to grow at unprecedented rates, visual exploration and analysis are taking central roles in the effort to make sense of today’s increasingly large and complex scientific datasets. Because of the vital role visualization plays in promoting scientific breakthroughs, we have to ensure that our techniques continue to develop to meet the challenges of big data. However, since visualizations are typically employed in larger analytical workflows, it is not sufficient to merely be able to send larger data volumes across the pipeline. An equally important endeavor is developing interactive analytical interfaces that seamlessly integrate with visualizations to enable scientists to see the data, to analyze it, and to think deeply about the phenomena depicted in these visualizations.

Scientific inquiry incorporates an iterative process of extracting evidence from raw data, representing that evidence in an appropriate form, and creating hypotheses and theories to present that evidence in a convincing narrative. This process is usually referred to as sensemaking [1]. Data visualization can be very helpful throughout the various

stages of sensemaking owing to the adeptness of the visual-perceptual system at detecting recurring patterns, prompting scientists to come up with hypotheses to explain those patterns. Visualizations can also be used to present and communicate evidence backing up hypotheses that were created earlier. Seasoned scientists strive during their investigation to maintain a broad perspective encompassing multiple hypotheses, as the quality of the analysis is largely dependent upon that. This is particularly important with today’s scientific datasets where the phenomena depicted are complex enough to warrant multiple complimentary or even competing theories. Human cognition however is naturally biased towards narratives that conform to the analyst’s intuition [2]. Our perception usually favors confirmatory evidence with a tendency to subdue contradictory information. These biases often cause distortions or even illogical interpretations [3]. With technological intervention however, the effect of cognitive bias can potentially be reduced [1]. By combining visualizations with interactive interfaces that make it easy to create and explore different hypotheses and narratives, we can assist scientists in maintaining a comprehensive analysis where a wide variety of views are considered before drawing conclusions.

One technology that can potentially provide a remedy to this problem is Large, High-Resolution Displays. These displays are being rapidly adopted in many scientific disciplines because they are an effective way to provide both context and detail when visualizing large-scale data [4]. Figure 1 illustrates example uses of this technology. The benefits of these displays go beyond their capacity to simply display larger amounts of data. There is strong evidence that larger display surfaces improve user performance under cognitively demanding tasks [5]–[7], while taking advantage of our embodied cognitive abilities [8], [9]. Recent studies demonstrated that such displays can help in complex analytical tasks where an analyst tries to make sense of a large body of data [10]. The capacity to visualize large amounts of information and their potential in augmenting our cognitive and perceptual faculties put large, high-resolution displays in a unique position to provide robust platforms for ultra-scale data visualization and analysis. Nevertheless, designing interactive visual analytic interfaces that effectively utilize these environments in large-scale data exploration scenarios is still largely an open question.




Figure 1. Large, high-resolution displays being used to explore nanoscale molecular datasets. The left side illustrates a small-multiple layout showing a small molecule under a variety of conditions. The right side illustrates a visualization of a large-scale molecular dynamics simulation comprising 5 million atoms.

In this paper, we provide a perspective on this problem space. We present an interactive visual analytics application for exploring and making sense of complex trajectory datasets. Our design is motivated by Pirolli and Card’s sensemaking model [1] with the goal of promoting an analysis that considers multiple competing hypotheses. The application takes advantage of large, high-resolution displays to simultaneously visualize a large number of trajectories in a small-multiple layout [11]. It also includes a Coordinated Brushing and Highlighting tool which can be used to formulate and test hypotheses with quick, scalable visual queries. To evaluate the design, we conducted a pilot user study with a domain expert- a behavioral ecology researcher investigating navigation strategies of insects by studying their movement. The study demonstrated the effectiveness of the application and provided insights on how scientists could utilize large-scale display environments in visual exploration and sensemaking scenarios. In summary, the contributions of this paper are:

- A general technique for the exploration of large collections of related data on large, high-resolution displays. The technique comprises a small-multiple layout and a coordinated brushing tool, which can be used to visually formulate and explore hypotheses about the data in a scalable manner.
- A pilot user study that evaluates the proposed technique within the context of a behavioral ecology application.

The rest of the paper is divided as follows. In section 2 we survey the literature and place our work in context. In section 3, we give background on Pirolli and Card’s sensemaking model, which we use to motivate the design of our application. Section 4 describes the design of the trajectory exploration application within the context of a behavioral ecology use case. In section 5, we report on the results of a pilot user study with a domain expert. We discuss the results further in section 6, grounding our observations in the sensemaking model. Section 7 concludes the paper and gives future research directions.

II. RELATED WORK

Our work can generally be classified under the umbrella of visual analytics (VA). Technologically, our work takes inspiration from a large body of research on the benefits of using large, high-resolution displays in visualizations.

Visual analytics is the “science of analytical reasoning facilitated by visual, interactive interfaces” [12]. It is an emerging field that represents an outgrowth of visualization, drawing from many traditions including human-computer interaction, machine learning, and psychology. The main focus of VA is tackling problems of scale and complexity which requires coupling of machines and humans, particularly when the problem involves enormous amounts of data. In this arrangement, machines provide computational analysis and convert data into visual representations, while the human supplies judgment and interpretation. It is therefore essential to understand how people make these judgments and interpretations to come up with theories when faced with a large amount of data [13]. Perhaps the most recognized model for this process in the visualization community is Pirolli and Card’s sensemaking model [1]. Their notional model (described in section 3) provides the theoretical background for our work.

Many VA techniques have been developed to help in analysis scenarios involving big amounts of data. One particularly relevant example is the work of Schreck et al. on the visual analysis of large financial data [14]. Their visualization employs self-organizing maps (SOM) to cluster 2D time-dependent financial indicators as trajectories. Hurter et al. describe a number of interactive techniques that allow the exploration of trajectories using quick brushing, juxtaposition, and pick and drop operations [15]. Our application employs a small-multiple layout similar to Schreck et al’s, but we implement different interactive features geared towards large, high-resolution 3D displays. In addition to addressing the visualization aspect, we look at the thought process involved in making sense of large-scale trajectory data. Our design is also motivated by a desire to encourage a researcher to consider multiple hypotheses throughout the sensemaking process.

Some research has been done to develop analytical user interfaces with the explicit goal of facilitating hypotheses creation. Pike et al. describe an environment for collaborative sensemaking which ties data artifacts with visual depictions of the analysts’ thought processes and chains of reasoning [16]. Kehrer et al. describe a tool for multivariate data analysis which can be used to visually identify promising hypotheses, while narrowing down the parameter search space for further computational analysis [17]. The Jigsaw system supports intelligence analysis tasks involving a large body of documents [18]. It employs multiple coordinated views, helping analysts in combining scattered pieces of evidence into concrete theories. Maciejewski et al. describe

a dual-view visualization combining a geospatial and a time series view for the analysis of spatio-temporal datasets [19]. Their tool facilitate hypotheses creation by highlighting hotspots with abnormal data, allowing an analyst to further investigate them in the spacial and temporal domains. Our work combines ideas from the above research, employing coordinated multiple views coupled with interactive tools for data exploration. However, we extend the concept of coordinated views to large, high-resolution displays in order, scaling up the number of views to increase the amount of data that can be visualized.

There is a large body of research pointing to the advantages of using large, high-resolution displays in knowledge-based work environments. The benefit of these displays however goes beyond their capacity to simply display larger amounts of data. Their large display surface, which “approaches the sphere of perception and influence of the human”, enables users to exploit embodied cognitive abilities, such as spatial awareness and spatial memory [9]. With respect to visualization, large displays have been shown to improve many basic tasks. Ball et al. report improved performance in target search and pattern finding, two common visualization tasks [20]. In another study they found that users preferred physical navigation (such as walking in front of the display, and turning ones head) over virtual navigation (such as panning and zooming) [8]. Increased physical navigation was also correlated with improved performance in visual analysis tasks in the same study.

Few studies looked at how large, high-resolution displays could be leveraged to support complex sensemaking tasks with large amounts of data. One of the few studies conducted was Andrews et al’s, in which they showed that the large display served as fast external memory where a large amount of data can be externalized and consulted quickly during the analysis [10]. They also observed the empty surface of the display was leveraged to create cognitively efficient information organization schemes, where location of items encoded semantics. Andrews et al’s study provides evidence on the usefulness of large, high-resolution display environments in sensemaking scenarios involving a large amount of data. However, their study employed mostly textual data with impoverished visual content. The interaction paradigm was also limited to the traditional desktop metaphor, which allowed separate but disconnected windows to be resized and moved in space. The study did not make use of interactive features typical in modern visualization such as coordinated highlighting and multiple views.

It remains unclear how to design VA tools that leverage the affordances of high-resolution displays to support the visual exploration and analysis of scientific datasets. In fact, there have been recent calls to address this gap [9]. Our work provides perspective on this problem space. We also present a scalable visual query technique, which can be used to quickly formulate and verify hypotheses, encouraging

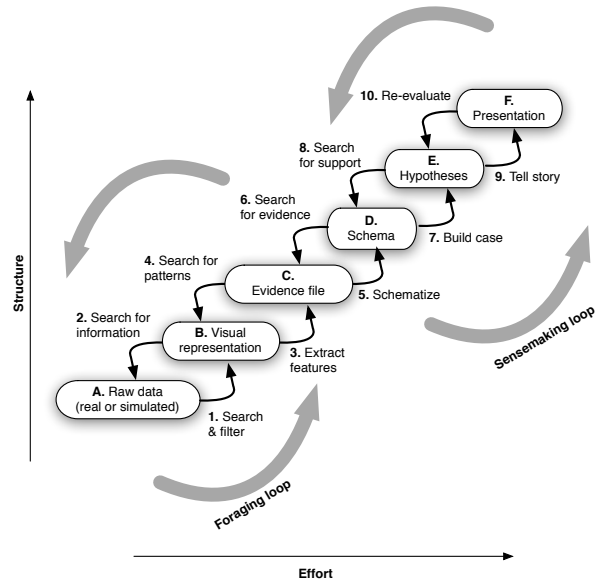


Figure 2. Pirolli and Card’s sensemaking model [1], with minor adaptations.

scientists to contemplate multiple competing hypotheses before drawing conclusions.

III. SENSEMAKING: A MODEL FOR SCIENTIFIC INQUIRY

To design effective visual analytic systems, we have to understand the process scientists employ to make sense of data, sensemaking. Based on this understanding, we can design visualization systems that conform to this process, taking advantage of our strengths, while providing technological interventions to overcome our perceptual and cognitive limitations when possible.

We first describe Pirolli and Card’s model of the sensemaking process, which has been widely adopted by the visualization community. We then address the question of how to design visual analytic platforms that improve the efficiency and quality of sensemaking when that process involves large-scale data that is potentially amenable to many interpretations.

A. Definition

Sensemaking is the process organizing scattered and incomplete pieces of data, extracting evidence form them, and combining that evidence into a cohesive narrative. Sensemaking is an inherently cognitive process that requires active human engagement to provide judgment and interpretation throughout. It arises in situations as diverse as collecting information from classified ads in order to apply to a suitable job to conducting a literature survey as part of writing a research paper. Researchers and scientists also engage in sensemaking in order to inspect real or simulated data,

analyze evidence, create hypotheses, and present them to peers.

Sensemaking is an integral part of scientific inquiry. As visualization tools start to take an increasingly important role in this process, it is important for visualization designers to understand this process [12]. Ignoring this body of knowledge could result in visualizations that hinder the performance of scientists instead of helping them [13].

B. Sensemaking: a model

One of the most widely accepted models of sensemaking was originally proposed by Russell et al. [21] and later refined by Pirolli and Card [1]. Russell et al. conducted observations of individuals and groups engaged in many sensemaking scenarios including designing a technical curriculum for printer maintenance, and analyzing intelligence to uncover a terrorist plot. Remarkably, Russell and Pirolli found that the general structure of sensemaking was consistent across a variety of tasks and disciplines.

The sensemaking process can be expressed as a bi-directional flow chart in which the raw data is transformed through a series of stages. These stages reflect processes that attempt to organize data into new representations so that the desired task can be performed more efficiently. The general structure of sensemaking is depicted in Figure 2. Here we adapt Pirolli et al's model so that it is more representative of the sensemaking process when working with big scientific datasets. Particularly, we emphasize data visualization in earlier stages of the model as it often represents the first opportunity to look at the data. The process broadly can be broken down into an information foraging loop, where the necessary evidence is extracted and placed into an intermediate repository, and a sensemaking loop, where the evidence is marshaled, and tied together into a narrative.

1) *Information foraging loop*: The process begins with the raw data, which could be real or simulated. The first step is to filter the raw data and select relevant subsets, such as a particular time window, or one or more relevant variables. The filtered data is then visualized, converting it into one or more visual representations such as volume visualizations, graphs, or scatter plots. The initial visual representations of large-scale data are often too complex to make sense of, so the researcher may wish to extract important features from the visualizations such as a cluster of data points in order to provide evidence for a theory. In interactive visualizations, the researcher can employ interactive tools such as zooming onto a particular portion of the data and highlighting relevant features. Those features along with any supporting information such as annotations comprise the 'evidence file' which contains low-level inferences about the data.

2) *Sensemaking loop*: As evidence accumulates, it is often desirable to illustrate it in a more efficient representation by juxtaposing snapshots from different visualizations, or by

annotating and drawing links between related items, for instance. The new representation is referred to as a 'Schema'. Schematization is essentially a process of marshaling scattered pieces of evidence into a cognitively efficient representation that makes certain aspects more salient, allowing the analyst to build a case. Interactive visual analytic tools can be extremely helpful in schematization by providing ways to cross-highlight related elements in different visualizations for example, elevating raw visualizations to schemas. The researcher then invokes his domain expertise, and attempts to piece the different fragments of evidence to build a case for a theory.

It is important to stress that the sensemaking model should not be interpreted as a linear, waterfall like process. In fact, the process is highly fluid and iterative with the analyst often going back to previous stages (as the back pointing arrows in Figure 2 suggest). As Pirolli et al. put it "Information processing can be driven by bottom-up processes (from data to theory) or top-down (from theory to data), often invoked in an opportunistic mix".

C. Technological support for sensemaking

The question is, how do we design visual analytic interfaces that augment our perceptual and cognitive abilities in order to improve the efficiency and quality of sensemaking? Pirolli and Card suggest leverage points where technology could be employed to augment different stages of the sensemaking process. For example, in the information foraging loop, we can increase the amount of data that can be assessed on a low-fidelity basis while highlighting relevant pieces of data using pre-attentive visual coding. This enables the researcher to investigate a larger portion of the data early on in the analysis. There are many well understood visualization techniques to take advantage of the human's visual perception system to encode information efficiently [22]. However, the role of high order cognition in sensemaking, and how to support that role with visual analytic tools is far less understood [23].

One of the key challenges during sensemaking is overcoming cognitive bias. Human perception is naturally biased towards confirmatory evidence that conforms to intuition, while tending to discard data that contradicts previously held beliefs. This can potentially cause the analysis to deviate from formal rationality [2], [24]. Seasoned researchers strive to maintain a wide perspective when evaluating evidence, keeping multiple competing hypotheses alive throughout their analysis [3]. Large, high-resolution displays can potentially provide a technological intervention to encourage this behavior. With properly designed tools, the ability to visualize a large volume of data can be leveraged to also meta-visualize a large number of questions and hypotheses, enabling scientists to quickly explore different theories and narratives throughout the analysis. In the following section, we describe how our design achieves that.

IV. CASE STUDY: LARGE-SCALE VISUAL ANALYSIS OF INSECT TRAJECTORIES

In this section we describe the design of a visual analytics application for the exploration of complex trajectory datasets. Our application comprises an interactive visualization designed for a large, high-resolution 3D display environment. The application provides a set of intuitive yet expressive interactive features, which allow the analyst to perform rapid, scalable visual queries on a large portion of the dataset. By casting hypotheses in terms of straightforward visual queries, the researcher can easily formulate and contrast different hypotheses, and verify whether the data supports those hypotheses. First, we motivate the application by describing a use case in the domain of behavioral ecology where researchers need to explore and analyze a large number of trajectories depicting insect behavior. We then describe the visualization along with its interactive features.

A. Motivation

To understand the navigational strategy and decision-making processes animals employ, ecologists track and analyze their movement patterns. However, many organisms such as insects exhibit a stochastic, locally scoped behavior that is difficult to characterize on a case-by-case basis. Therefore, entomologists resort to collecting a large sample of trajectories under varying conditions to tease out the general insect behavior. Due to the large number of plausible explanations and hypotheses concerning an observed behavior, entomologists need a scalable and efficient way of exploring these different theories and narratives. The sheer number of trajectories collected during experimentation makes them extremely difficult to visualize on traditional desktop screens. While few trajectories could be visualized simultaneously on a desktop screen, the researchers would need to switch between different sets of trajectories to cover the data. This makes it hard to perform comparison across a large set of trajectories, a task that is crucial in this domain. Furthermore, it is difficult to formulate and test hypotheses to explain general insect behavior by looking at few instances of trajectories at a time.

This use case therefore provides a good reason to use large, high-resolution displays so that a larger portion of the dataset can be visualized at the same instant. However, here we also ask two additional questions: does the visualization encourage the researcher to create, verify, and contrast multiple hypotheses in an attempt to explain different strategies insect employ in navigation? How do we channel the affordances of the large display to promote a multi-perspective analysis that explores different theories? We designed our visualization with these two questions in sight. Generally speaking, we included a set of interactive features that make it possible for the researcher to perform fast, scalable visual queries on a large number of trajectories simultaneously. The rationale behind this is that given the ability to formulate



Figure 3. Our trajectory analysis application running on a 19 Megapixels, thin-bezel, tiled 3D display wall. Trajectories are juxtaposed side-by-side and grouped into “bins” depending on their associated meta data. A stereoscopic 3D visual encoding is used to convey spatial and temporal features in the trajectories. The frame at the right bottom of the figure illustrates *Coordinated Brushing*. A blue paintbrush (circled) is being used to highlight insect movement in the center of the experimental arena.

queries with little effort and time investment, a researcher will likely use this feature to create and follow up on a larger number of theories and narratives throughout the analysis.

B. Data

Our dataset comprises approximately 500 trajectories, which represent movement of ants under different experimental conditions. The trajectories were obtained by tracking ants’ movement in the field at approximately 3mm spatial resolution. Each trajectory represents the movement of a single ant (of the Kenyan Seed Harvester Ant specie - *Messor cephalotes*), which has been captured, taken away from its colony, and placed on the center of an experimental arena in order to study its navigational strategy. Trajectories range in duration from 10 seconds to 3 minutes. The goal of the experiments was to analyze the navigational strategy employed by the ants. Trajectories were categorized based on the state of ant when captured. Some of the variables include: position relative to the main foraging trail, journey direction (heading away from or returning to the colony), and whether the ant was carrying a seed.

While the raw size of this dataset is relatively modest, nonetheless, it poses significant visualization and analytical challenges that would be difficult to address with traditional desktop screens. Furthermore, due to stochasticity inherent in this dataset, it is susceptible to a large number of plausible theories and explanations, which provide a good test case to evaluate the scalability of the visual query technique.

C. Design

Figure 3 shows a picture of the visualization environment. The visualization was rendered on a wall-sized, stereoscopic, tiled LCD display in a 6 x 3 arrangement with a total size of 7 x 3 meters (approximately 23 x 10 feet). The visualization

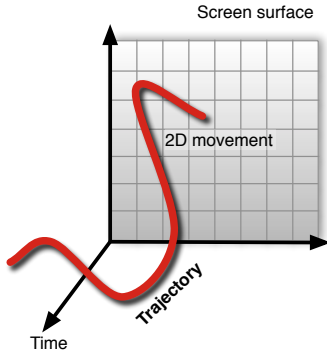


Figure 4. Visual encoding of an ant trajectory, with stereoscopic depth cues to convey time.

application utilized 2/3 of the total display surface at a resolution of 8,192 x 1,536 (approximately 12.5 million pixels). The user interacted with the display using a mouse and a keyboard placed on a desk that was positioned about 3 meters (10 feet) in front of the display. We describe the visualization and follow with an illustration of its interactive features.

1) *Visualization*: Each trajectory was rendered in stereoscopic 3D using the space-time cube metaphor [25]. The XY plane (the display surface) encodes 2D movement, while the Z+ axis (away from display) encodes time. Figure 4 illustrates the visual encoding for a single trajectory. To avoid perspective distortion, an orthographic projection was employed. The visualization renders two separate views for each eye, enabling a viewer with a pair of polarized 3D glasses to perceive the trajectory in stereoscopic 3D. The trajectories appear as a cylinder starting at the display surface, extending out to ‘float’ in front of the display. The stereoscopic view makes spatio-temporal patterns evident and enables the analyst to disambiguate situation where a trajectory contains overlapping segments. We visualized the trajectories in a small-multiple layout, making it possible to look at a large subset of them simultaneously and facilitating comparison (Figure 3).

2) *Interactive features*: A number of interactive features were implemented to facilitate analysis and comparison of the trajectories, and to enable scalable visual queries. These features were implemented to take advantage of the small-multiple layout, the large display surface, and the stereoscopic 3D capability of the display:

- *Small-multiple Layout*: The number of trajectories visualized can be varied. The user can switch between a number of configurations by pressing a number on the keypad: ‘1’, ‘2’, etc... Some of the pre-configured layout provided include a 15x4, 24x6, and 36x12. These configurations were chosen to avoid a trajectory overlapping with a bezel (the gap created by overlapping borders of adjacent LCD panels). Although the bezels in our display were thin (less than




Figure 5. Illustration of how Coordinated Brushing can be used to test a hypothesis by means of a visual query. In this example, the researcher checks whether “Ants that were captured *east* of the colony’s foraging trail will exit the experimental arena from the *west* side.” A red highlight in majority of trajectories indicates the hypothesis is supported by the data.

1cm in thickness), some users reported discomfort when stereoscopic 3D content overlaps a bezel. Furthermore, we wanted to make use of bezels as natural dividers, since this behavior was commonly observed when working with large, tiled displays [20].

- *Trajectory Grouping*: The user can define rectangular groups that encompass a contiguous subset of trajectories. A set of filters can be associated with each group to show trajectories that satisfy the filter only. Groups can be given different background colors to easily distinguish them. For example, Figure 3 illustrates five groups corresponding to ants captured *on* the main trail (blue background), *west* (red), *east* (yellow), *north* (gray), and *south* (green) of the main trail.

- *Coordinated Brushing*: Generally speaking, coordinated brushing allows the user to highlight a particular data item in one view, causing related items in all other views to be auto highlighted as well. In our application, the user can brush the background of a single trajectory using a paintbrush tool. This causes segments in all currently displayed trajectories to be highlighted when the insect moves over a brushed area. Figure 5 illustrates this feature.

When we designed this feature, we envisioned that it would be helpful when looking at trajectory similarities across the entire layout. For example, the user can brush a portion of one interesting trajectory, which would cause trajectories with a similar movement pattern to be highlighted. Much to our surprise, early trials showed that this simple feature played a far more important role: it enabled the researcher to perform fast visual queries on a large portion of the dataset. The results of queries are highlighted in a distinct color, making them pre-attentively perceivable. Often, the researcher was able to formulate visual queries that corresponded to hypotheses. We discuss this facet in detail and give examples in sections 5 and 6.

- *Temporal Filter*: A time-window can be specified, caus-

ing the visualization to display segments of trajectories corresponding to insect movement during the specified time window only. For example, the researcher can display the beginning, the middle, or the end period of the experiment. The time window of interest can be selected using a range slider.

•*Ergonomic considerations*: Prolonged viewing of stereoscopic images has been known to cause discomfort for some viewers, mainly due to excessive binocular parallax and accommodation-convergence conflict [26]. To reduce the chance of fatigue, we included a set of controls in the visualization to modify the 3D view to allow for comfortable, prolonged viewing. A slider allows the user to push trajectories so that they lie in front of the display surface, behind the display surface, or somewhere in between. Additionally, the time-scale can be (de)exaggerated using a second slider. Using these two sliders, the user can control the maximum amount of binocular parallax and keep it within a comfortable range while maintaining sufficient depth cues.

V. EVALUATION

To evaluate the usefulness of our visual analytics application, we conducted a pilot user study with a domain expert. Our goal was two folds. First, we wanted to get a sense of how a researcher could utilize the interactive tools that we built into the visualization to explore a complex dataset. Secondly, we wanted to understand how the researcher adapts her sensemaking workflow to take advantage of the large display. Particularly, we wanted to know if the application encouraged the research to contemplate and compare several hypotheses before coming to conclusions, and whether the interactive tools made this easy.

For this pilot user study we had a single participant. One of the authors, a behavioral ecology doctoral student who was actively involved in the ant navigation research project used the visualization to explore the ant trajectory dataset. The session was video and audio taped. To analyze the recording, we developed a coding scheme to tag the video, indicating instances when:

- The researcher made an observation about the data.
- The researcher created a hypothesis.
- The researcher utilized one of the interactive tools (such as the coordinated brush), along with the question or hypothesis she was trying to answer.

The tagged video recording and transcripts were analyzed to shed a light on the two research questions we outlined earlier, namely, how the researcher utilized the visualization, and what effect this had on the sensemaking process. We report on our qualitative observations below. In section 6, we draw connections between user behavior and the sensemaking model outlined in section 3 in order to illustrate how the visualization and the screen form factor shaped the user’s sensemaking process.

A. Data grouping for comparative analysis

One feature that was used extensively is Trajectory Grouping, which allowed the researcher to divide the large display surface into bins that held distinct groups of trajectories. Most of the time, the display was divided into two to five groups containing trajectories of ants captured under different experimental conditions. For example, Figure 3 illustrates one scheme where five groups contained trajectories of ants that were captured east, west, south, and north, and on the colony’s main foraging trail. The extensive use of data binning is consistent with other studies; users tend to utilize the large display surface to partition the data, placing related data items close to each others [10], [27].

B. Scalable visual query using coordinated brushing

Coordinated brushing proved to be intuitive and extremely flexible, enabling the researcher to quickly formulate visual queries on the data. Once she had a question or hypothesis in mind, the researcher needed only to apply the paintbrush to a single trajectory, which produced a color-highlight in all other trajectories when the insect crossed the brushed area. The user appeared to rely on her pre-attentive visual processing to determine which trajectories satisfied the query. In some situations, the entire dataset could be visually queried in a matter of few seconds. For instance, to test whether ants captured *east* of the main foraging trail exit the experimental arena from the *west* side in an attempt to get back to the trail, the researcher brushed the left (west) part of the arena with red. Her hypothesis was quickly confirmed upon seeing a concentration of red highlight in the ‘east’ group. Figure 5 illustrates this.

Coordinated brushing was often used in conjunction with the temporal filter. This allowed for effective testing of complex spatio-temporal patterns. For example, to determine whether ants that have dropped the seed they were carrying spend more time in the center searching for the seed before deciding which direction to take, the user would brush the center of the experimental arena with green and set the temporal filter to display the beginning of the experiment. The hypothesis can be verified by looking for green segments that are roughly perpendicular to the display surface, which would indicate a stationary ant. This form of visual query can be made thanks to the stereoscopic view, which encodes time as length in the Z axis.

C. Benefits of stereoscopic 3D

The 3D stereoscopic view made spatio-temporal patterns evident and enabled the user to disambiguate situation in which a trajectory’s segments overlap in space, making it clear when an ant returns to an earlier spot on the experimental arena, for example. The user also indicated that, with the stereo 3D view, she was able to perceive the periodicity of ant behavior not only on a single-trajectory basis, but also on a larger scale. This suggests that stereoscopic depth

cues can be leveraged to encode visual features that can be processed pre-attentively on a large 3D display.

VI. DISCUSSION

Overall, the experience of the researcher was overwhelmingly positive. Originally, the researcher used Matlab as her analysis platform, visualizing trajectories one at a time. The new visualization proved to be far more helpful in the analysis, making it “easier to think about [the problem] visually than in Matlab”. The user study also illuminated a number of unique ways in which the large display was utilized to further the analysis task. Of particular interest is the use of coordinated brushing as a way to perform scalable visual queries on a large portion of the dataset. This capability translated into an ability to rapidly test hypotheses using one or more visual queries.

We can recognize two broad behaviors that the researcher engaged in during the study, namely, making comparisons between groups of trajectories, and creating hypotheses and weighing them against the evidence. We discuss these two behaviors, drawing connections back to Pirolli and Card’s sensemaking model.

A. Making comparisons and low-level inferences

A significant portion of the analysis workflow comprised comparisons in which groups of trajectories were visually compared and contrasted. The small-multiple layout facilitated this behavior, enabling the researcher to shift her attention between trajectory groups and compare them both at the individual as well as the group level. Although such comparative analysis could be in theory performed on a traditional desktop display, the size and resolution of the desktop display would severely limit the number of trajectories that can be compared simultaneously. This would force users to switch between different views, potentially increasing the cognitive workload and distracting them from the analytical task at hand. In some studies users also showed signs of frustration at the increased amount of window switching and virtual navigation they had to perform with small displays [27]. On a large, high-resolution display on the other hand, a larger portion of the data can be visualized at once, making it possible to compare a large number of data items at literally a flick of an eye. By virtue of having a larger portion of the data available for immediate visual consultation, the analysis would not only be sped up, but one could potentially make inferences and discoveries that would be extremely difficult to make without seeing all the relevant items at once.

Comparative analysis can be tied to Steps 3 (*extract features*) and 4 (*search for patterns*) of the foraging loop in the sensemaking model (Figure 2). During this stage, a analyst attempts to make low-level inferences concerning the visual similarity (or dissimilarity) of trajectories. For example, during our study, the researcher described some

trajectories as being ‘more windy’ corresponding to ants captured on the main foraging trail, while trajectories of ants captured off the trail were characterized as being ‘more direct’, reflecting a desire to head in a certain direction. The purpose of these low-level inferences is to elevate the visual representations (Box B in Figure 2) to an *evidence file* (Box C) so that they could be used as supporting evidence for hypotheses created later during the analysis. Although the user did not explicitly generate an evidence file as a separate artifact, the fact that all trajectories were persistently available on the screen for consultation likely eliminated the need for one. Therefore, the small-multiple layout in itself could be considered an evidence file in our case. The low-level inferences derived from comparison were often recalled later to provide supporting evidence for hypotheses. One limitation in our design is that there was no explicit way of recording or tagging those inferences. A future iteration of the design could add this feature.

B. Hypothesis creation and verification

Another tool that was used extensively is coordinated brushing. Analysis of the video recording reveals an interesting sequence with respect to how this tool was employed. We noticed that comparison of trajectory groups often prompted the researcher to formulate a question or a hypothesis. For instance, “Do ants that were captured *east* of the main foraging trail exit the experimental arena from the *west* side in an attempt to get back to the trail?” Once the hypothesis is formulated verbally, the researcher would proceed to weigh the data against that hypothesis. The first step is isolating the relevant data instances into a distinct group. In this case, filters were set to show trajectories of ants captured *east* of the trail only. Once the relevant data is isolated, the researcher would cast her question in terms of a visual query. The query can be conceived as “look for trajectories that terminate at the *west* (left) side of the experimental arena and determine if they constitute a majority.”

Nevertheless, in typical circumstances, this remains a complicated visual query which requires the researcher to individually inspect every single relevant trajectory to determine whether the ant had exited from the correct side. Using coordinated brushing however, the query can be further simplified. In this case, the researcher can brush the left side of the arena with red causing a red highlight in all other trajectories whenever the ant is over the *west* side (Figure 5). The researcher can also set the temporal filter to only show the last few seconds of the experiment. With this combination the original query is reduced to searching for red segments in the visualization, which can be easily performed by quickly glancing at the small-multiples layout, owing to the pre-attentive visual encoding of query results. With a traditional desktop screen, checking this is still a tedious, slow task given the large number of instances that need to be checked one by one. On a large, high-resolution

display on the other hand, visual queries could be scaled up to ‘cover’ a much larger portion of the dataset. In our setup, it was possible to simultaneously visualize 432 trajectories. Thus, the researcher could quickly apply her queries and instantly see the results on 85% of the data. Thanks to the pre-attentive encoding, the results of queries could be perceived in a matter of few seconds.

What makes visual queries take an even more important role is that, in a many cases, a query corresponds to a hypothesis. The ability to express high-level questions in terms of easily perceivable visual features translated to an ability to formulate and verify hypotheses in rapid succession. During the user study, several hypotheses could be formulated and tested within a span of few minutes. In fact, the researcher spent most of the time contemplating a variety of theories and scenarios and evaluating them with quick visual queries. While visual queries may not be enough to fully substantiate a particular theory, they nevertheless provide a high-fidelity, low-cost data assessment scheme, which can be used to explore a larger number of hypotheses and identify the promising ones for further analysis.

The use of coordinated brushing can be linked to schematization (Step 5 in Figure 2). The small-multiple visualization with relevant trajectories could be regarded as an *evidence file* comprising scattered pieces of information that needs further refined. Brushing and highlighting amounts to a refinement process that elevates the evidence file to a schema- a higher-order representation that provides concrete support for a particular theory.

Although coordinated brushing as a technique is not new, its application in large, high-resolution display environments gives rise to intriguing new possibilities. The ability to formulate flexible visual queries on a large display provides a powerful method for visually exploring big data collections from different narratives and perspectives. Thanks to the large-scale, pre-attentive visual encoding of query results, a researcher can follow up on multiple hypotheses and rapidly determine whether those hypotheses are supported by the data at hand, enabling him/her to quickly explore a complex hypothesis space.

C. Scalability

Although the case study presented in this paper dealt with a moderately sized dataset, we note that visual queries can potentially operate at different levels of scale, making them scalable to larger datasets. This potentially allows us to employ a similar visualization to study a trajectory dataset comprising 10,000, 100,000 or perhaps even a million traces. Instead of showing individual trajectories, we can cluster those trajectories based on feature similarly by employing self-organizing maps for example. The unit of exploration becomes a cluster of trajectories that exhibit similar spatio-temporal patterns. The small-multiple layout would be adapted to visualize and juxtapose cluster aver-

ages instead of showing individual trajectories. Coordinated brushing can still be employed to explore those clusters in a similar manner. While this does change the granularity of the analysis, one can analyze a larger dataset using such visualization. Moreover, a user can interactive ‘zoom in’ on a particular cluster of interest and query the cluster at the individual-trajectory level, enabling one to explore the dataset at multiple scales.

Alternatively, one can scale up the amount of data instances that can be visualized in a small-multiple layout by employing more compact visual encodings. For example, a representation that shows general trajectory shape while discarding high-frequency features could be employed. This reduces the amount of screen real-estate needed for a single instance, allowing a larger number of instances to be visualized simultaneously.

VII. CONCLUSIONS AND FUTURE WORK

Large, High-Resolution displays are being adopted in many scientific disciplines because they provide an effective way to see both context and detail when visualizing massive amounts of data. The benefits of these displays however go beyond their capacity to simply visualize larger data volumes. When coupled with specially designed visual analytic interfaces, these displays can provide robust analytical platforms to enable scientists to make sense of complex datasets, and explore the data from multiple perspectives and narratives.

In this paper, we presented an application for the visual exploration and analysis of complex trajectory datasets. Our application utilizes a small-multiple layout to visualize and juxtapose a large number of trajectories. It also embodies a scalable visual query technique that takes advantage of large, high-resolution 3D displays. A pilot user study demonstrated the effectiveness of the design in supporting high-level sensemaking processes, enabling the creation of rich visual representations of *evidence files* with little effort. This in turn translated into an ability to formulate and explore hypotheses in a scalable manner, encouraging users to consider a large number of theories throughout the analysis.

Although the scenario explored in this paper pertains to the analysis of movement trajectories in the context of behavioral ecology, we believe the techniques illustrated here are applicable in different science domains. In particular, we believe the concept of scalable visual queries could be generalized to other applications especially when dealing with large collections of related data instances, such as ensembles of simulation runs under different conditions. In the future, we will continue to work with different communities of scientists in order to explore those possibilities. We will also look at ways of integrating our application into larger scientific workflows to support evidence and insight provenance. Although preliminary evaluation demonstrated the usefulness of the techniques presented, long-term studies are

needed to fully understand the impact of display resolution and form factor on exploratory data analysis with interactive visualizations.

ACKNOWLEDGMENTS

Part of this work was performed during the joined Princeton and UIC Field Computational Ecology course in Spring 2012 (www.cs.uic.edu/bin/view/ComputationalEcology), with co-instructors Tanya Berger-Wolf, Daniel Rubenstein and Iain Couzin, who were instrumental in several parts of this research. We thank the staff at Mpala Research Center (Kenya), Ol' Pejeta Nature Conservancy (Kenya), and fellow graduate students at EEB-Princeton University, CS at University of Illinois at Chicago, and University of Nairobi. Funding for this work was provided by the Department of Ecology and Evolutionary Biology at Princeton University, and National Science Foundation grants OCI-1152895, IIS-0747369, OCI-0943559, and CNS-0959053.

REFERENCES

- [1] P. Pirolli and S. Card, "The sensemaking process and leverage points for analyst technology as identified through cognitive task analysis," in *Proceedings of International Conference on Intelligence Analysis*, vol. 5, 2005.
- [2] A. Tversky and D. Kahneman, "Judgment under uncertainty: Heuristics and biases," *science*, vol. 185, no. 4157, pp. 1124–1131, 1974.
- [3] R. Heuer, *Psychology of intelligence analysis*. US Government Printing Office, 1999.
- [4] J. Leigh, A. Johnson, L. Renambot, T. Peterka, B. Jeong, D. Sandin, J. Talandis, R. Jagodic, S. Nam, H. Hur *et al.*, "Scalable resolution display walls," *Proceedings of the IEEE*, no. 99, pp. 1–15, 2012.
- [5] M. Czerwinski, G. Robertson, B. Meyers, G. Smith, D. Robbins, and D. Tan, "Large display research overview," in *Proc. of CHI'06*. ACM, 2006, pp. 69–74.
- [6] M. Czerwinski, G. Smith, T. Regan, B. Meyers, G. Robertson, and G. Starkweather, "Toward characterizing the productivity benefits of very large displays," in *Proc. Interact 2003*, vol. 3, 2003, pp. 9–16.
- [7] D. Tan, J. Stefanucci, D. Proffitt, and R. Pausch, "The info-cockpit: Providing location and place to aid human memory," in *Proceedings of the 2001 workshop on Perceptive user interfaces*. ACM, 2001, pp. 1–4.
- [8] R. Ball, C. North, and D. Bowman, "Move to improve: promoting physical navigation to increase user performance with large displays," in *Proceedings of the SIGCHI conference on Human factors in computing systems*. ACM, 2007, pp. 191–200.
- [9] C. Andrews, A. Endert, B. Yost, and C. North, "Information visualization on large, high-resolution displays: Issues, challenges, and opportunities," *Information Visualization*, vol. 10, no. 4, pp. 341–355, 2011.
- [10] C. Andrews, A. Endert, and C. North, "Space to think: large high-resolution displays for sensemaking," in *Proc. of CHI'10*. ACM, 2010, pp. 55–64.
- [11] E. Tufte and P. Graves-Morris, *The visual display of quantitative information*. Graphics press Cheshire, CT, 1983, vol. 31.
- [12] J. Thomas and K. Cook, Eds., *Illuminating the path: The research and development agenda for visual analytics*. IEEE Computer Society, 2005.
- [13] D. Russell, M. Slaney, Y. Qu, and M. Houston, "Being literate with large document collections: Observational studies and cost structure tradeoffs," in *System Sciences, 2006. HICSS'06. Proceedings of the 39th Annual Hawaii International Conference on*, vol. 3. IEEE, 2006, pp. 55–55.
- [14] T. Schreck, T. Tekušová, J. Kohlhammer, and D. Fellner, "Trajectory-based visual analysis of large financial time series data," *ACM SIGKDD Explorations Newsletter*, vol. 9, no. 2, pp. 30–37, 2007.
- [15] C. Hurter, B. Tissoires, and S. Conversy, "Fromdady: Spreading aircraft trajectories across views to support iterative queries," *Visualization and Computer Graphics, IEEE Transactions on*, vol. 15, no. 6, pp. 1017–1024, 2009.
- [16] W. Pike, R. May, B. Baddeley, R. Riensche, J. Bruce, and K. Younkin, "Scalable visual reasoning: supporting collaboration through distributed analysis," in *Collaborative Technologies and Systems, 2007. CTS 2007. International Symposium on*. IEEE, 2007, pp. 24–32.
- [17] J. Kehler, F. Ladstadter, P. Muigg, H. Doleisch, A. Steiner, and H. Hauser, "Hypothesis generation in climate research with interactive visual data exploration," *Visualization and Computer Graphics, IEEE Transactions on*, vol. 14, no. 6, pp. 1579–1586, 2008.
- [18] J. Stasko, C. Görg, and Z. Liu, "Jigsaw: supporting investigative analysis through interactive visualization," *Information visualization*, vol. 7, no. 2, pp. 118–132, 2008.
- [19] R. Maciejewski, S. Rudolph, R. Hafen, A. Abusalah, M. Yakout, M. Ouzzani, W. Cleveland, S. Grannis, and D. Ebert, "A visual analytics approach to understanding spatiotemporal hotspots," *Visualization and Computer Graphics, IEEE Transactions on*, vol. 16, no. 2, pp. 205–220, 2010.
- [20] R. Ball and C. North, "Effects of tiled high-resolution display on basic visualization and navigation tasks," in *Proc. of CHI'05 extended abstracts*. ACM, 2005, pp. 1196–1199.
- [21] D. Russell, M. Stefik, P. Pirolli, and S. Card, "The cost structure of sensemaking," in *Proceedings of the INTERACT'93 and CHI'93 conference on Human factors in computing systems*. ACM, 1993, pp. 269–276.
- [22] C. Ware, *Visual thinking for design*. Morgan Kaufmann Pub, 2008.
- [23] T. Green, W. Ribarsky, and B. Fisher, "Building and applying a human cognition model for visual analytics," *Information visualization*, vol. 8, no. 1, pp. 1–13, 2009.
- [24] D. Kahneman and A. Tversky, "Subjective probability: A judgment of representativeness," *Cognitive psychology*, vol. 3, no. 3, pp. 430–454, 1972.
- [25] M. Kraak, "The space-time cube revisited from a geovisualization perspective," in *Proc. 21st International Cartographic Conference*, 2003, pp. 1988–1996.
- [26] M. Lambooj, W. IJsselsteijn, and I. Heynderickx, "Visual discomfort in stereoscopic displays: a review," *Stereoscopic Displays and Virtual Reality Systems XIV*, vol. 6490, no. 1, 2007.
- [27] R. Ball and C. North, "Analysis of user behavior on high-resolution tiled displays," *Human-Computer Interaction-INTERACT 2005*, pp. 350–363, 2005.