

# Supporting Effective Common Ground Construction in Asynchronous Collaborative Visual Analytics

Yang Chen\*

University of North Carolina at Charlotte

Jamal Alsakran†

Kent State University

Scott Barlowe‡

University of North Carolina at Charlotte

Jing Yang§

University of North Carolina at Charlotte

Ye Zhao¶

Kent State University

## ABSTRACT

Asynchronous Collaborative Visual Analytics (ACVA) leverages group sensemaking by releasing the constraints on when, where, and who works collaboratively. A significant task to be addressed before ACVA can reach its full potential is effective common ground construction, namely the process in which users evaluate insights from individual work to develop a shared understanding of insights and collectively pool them. This is challenging due to the lack of instant communication and scale of collaboration in ACVA. We propose a novel visual analytics approach that automatically gathers, organizes, and summarizes insights to form common ground with reduced human effort. The rich set of visualization and interaction techniques provided in our approach allows users to effectively and flexibly control the common ground construction and review, explore, and compare insights in detail. A working prototype of the approach has been implemented. We have conducted a case study and a user study to demonstrate its effectiveness.

**Keywords:** Visual analytics, asynchronous collaboration, insight, multidimensional visualization.

**Index Terms:** H.5.3 [Group and Organization Interfaces]: Collaborative computing—Web-based interaction;

## 1 INTRODUCTION

With the growth of public web-based visualization communities (e.g., Many Eyes [26] and sense.us [16]), people can collaboratively analyze data outside the constraints of time and space. Asynchronous Collaborative Visual Analytics (ACVA) allows shared access to resources, such as expertise and datasets by releasing the constraints on when, where, and who works together. ACVA turns collaborative visual analytics into a social process where everyone can participate and thus makes it possible to analyze datasets of much larger scales [16].

In this paper, we focus on common ground construction in ACVA. In social and psychological research, common ground is defined as the shared understanding enabling communication between conversational participants [11]. In collaborative visual analytics, common ground construction refers to the visual analytic process, in which users evaluate work that may have been created individually, to develop a shared understanding of data with a collection of insights and hypotheses [19]. Effective common ground construction may minimize the need to verbally confirm actions among collaborators, reduce the cost of collaborative effort [11].

\*e-mail: ychen61@uncc.edu

†e-mail: jalsakra@cs.kent.edu

‡e-mail: sabarlow@uncc.edu

§e-mail: jyang13@uncc.edu

¶e-mail: zhao@cs.kent.edu

In asynchronous settings, non-verbal cues for common ground is especially important since verbal communications between the collaborators are usually difficult or even impossible.

ACVA users face significant challenges in common ground construction. The scale of ACVA is usually larger than Collocated Collaborative Visual Analytics (CCVA). For example, Many Eyes received over 1463 registered users, 2100 datasets, and 450 users' comments in its first two months of life [26]. Browsing and organizing such a large amount of information from such diverse users and datasets are challenging tasks. To make it worse, there lacks instant communication among ACVA users. It is difficult for them to collaboratively identify significant insights and capture relationships among insights through face to face discussion and direct manipulation as in CCVA [19]. Thus, there is an urgent need for effective and efficient visual analytics tools for ACVA common ground construction, especially for the following tasks:

**Task 1 Generating Overview:** Common ground construction usually starts from forming an overview of the insights that have been recorded and gathered [19]. The overview presents the overall structure, key aspects, and evolution of the insights to help users gauge the context and determine future direction [5]. Existing ACVA systems provide limited capability with manual insight browsing, inspection, and organization, which hinders users' effort on quickly forming a mental map of existing insights. New approaches for overview generation that satisfy the following requirements must be developed:

- **Collecting information effectively and efficiently:** Rich semantic information about insights is needed for automated insight organization, retrieval, and association according to varying user interests. Collecting such information should not impose extra burden to users, i.e., their ongoing visual exploration process should not be disturbed or diverted.
- **Employing automatic insight analysis:** Manual insight association and grouping are not realistic for a fast growing pool of ACVA insights. Automatic insight analysis techniques, such as automated insight correlation, clustering, and summarization, are direly needed to be developed and integrated into ACVA systems for fast and operative overview generation.
- **Supporting dynamic overview construction:** ACVA users usually have diverse information needs. Dynamic overview construction should be supported in ACVA systems so that users can explore the insight space according to their specific needs. Moreover, the system should allow the users to dynamically manipulate visualization results according to their changing interests and developing understanding.
- **Providing a rich set of views and interactions:** Multiple coordinated views should be provided to allow users to examine insights from different aspects. For example, **Temporal visualization** helps users track and employ temporal evolution of insights, so that they can keep awareness of timing and preserve historical contents of insights [15, 19]. Furthermore,

**proximity-based projections**, where closely related insights are visually presented as clusters, facilitate users' ability to browse a large number of insights at a glance. In addition, interactions should be provided so that users can effectively navigate in the insight space, retrieve insights of interest, and manage overview construction.

**Task 2 Supporting Sensemaking:** After users identify interesting insights from the overview, they need to closely explore these insights for hypothesis generation and evaluation. We argue that the following tasks are important in the sensemaking process of ACVA users:

- **Comparison:** Scalable comparison among insights should be supported by ACVA systems. The comparison can be among insights generated by different users, from different datasets, during different time periods, or insights contributing to conflicting or relevant hypotheses. Overlapping information identified from comparison helps people associate insights by various collaborators for acquiring additional evidence for developing hypotheses [19]. It also helps people retrieve contextual information, examine the historical evolution of the reasoning process, and evaluate conflicting hypotheses.
- **Revisiting and refining:** Sensemaking is an iterative process. ACVA systems should allow users to revisit the sources of existing insights and refine them without disturbing the ongoing analytic process. This function is important in promoting new insights and hypotheses.
- **Result outreach:** A crucial function of common ground is to ensure that analysis results can be preserved and shared among collaborators [19]. Hence, ACVA systems should provide solutions for this task.

In this paper, we propose a novel ACVA toolkit to address these critical tasks. The toolkit is currently focused on common ground construction for analytic insights<sup>1</sup> discovered from multidimensional data<sup>2</sup>, but the approach is general enough to be extended to other types of insights. Our toolkit automatically retrieves important semantic information about insights, such as what they are (e.g., clusters, outliers, ranks, etc.), relevant data information (e.g., dimension names, data item names, etc.), and meta information (e.g., authors, timestamps, etc.), from semi-automatically generated, formalized insight annotations [6]. A rich set of views and interactions built upon automatic insight analysis are then provided. They allow users to browse semantics and identify clusters from a large collection of insights, track their temporal evolutions, retrieve insights of interest, compare groups of insights, and preserve and share results for effective common ground construction. We demonstrate the effectiveness of the toolkit through a case study and a user study.

## 2 RELATED WORK

Researchers have studied how collaborators construct common ground in a variety of collaborative work, such as emergency task management [13] and tactical operations planning [12]. In visual analytics fields, Robinson [19] examined how analysts collaboratively synthesize visual information by conducting an experiment using analytical artifacts printed on paper cards. He observed a set of activities commonly taken by analysts for achieving common ground, such as describing information development, reviewing prior work, and identifying overlapping information for relating

<sup>1</sup>Clusters, outliers, ranks, and dimension correlations are examples of analytic insights. Please refer to [18] for the definition of analytic insights and [7] for their categorization.

<sup>2</sup>A *multidimensional dataset* is a spreadsheet where each column is called a *dimension* and each row is called a *data item*.

insights. Based on the experimental evidence, he emphasized the critical role of common ground in collaborative synthesis and suggested using multiple strategies and visual metaphors (e.g., timeline and category group) for organizing findings. Mahyar et al [17] studied how users take notes and organize notes in collocated collaborative analysis. Their results indicated that users often use multiple approaches (e.g., ordering by chronological history) to organize notes, which help them better communicate and discuss with each other. Experimental evidence of such studies, regardless of the specific tasks, resulted in similar lists of tasks and design implications critical for establishing common ground. These tasks and implications, as summarized in Section 1, guide us in better supporting common ground in asynchronous collaborative settings.

Researchers have developed tools for supporting collaborative visual analytics in both synchronous and asynchronous settings. In synchronous collaborative visual analytics systems, common ground is usually established through real time shared views and instant communication mechanisms. For example, VizCept [10] allows users to keep track of each other's findings and relations in a shared concept map. Users can refer to such a shared view to ground their actions. Reality Instant Messaging [9] integrates an online social tool into visualization systems, helping users coordinate their activities, interests, and emotions in decision making processes.

A few efforts have been made to support common ground construction in ACVA in recent years. Most of them fall into one of the following three categories:

**Using exploration histories:** Shrinivasan and Wijk propose an approach that records exploration histories of users by capturing their visualization states [24][25]. The exploration histories are then used to automatically connect related findings and help users understand their collaborators' analysis strategies [23]. Exploration histories are also recorded in the interaction (e.g. mouse clicks [24]) and action (e.g. zooming and panning [14]) level. However, it can be difficult to organize and summarize insights according to their high level semantic meaning, if the exploration histories consist of exploration steps with little semantic meaning. In addition, the large amount of exploration steps toward each insight may hinder an ACVA system on constructing common ground for a large number of insights.

**Using comments and annotations:** Most web-based collaborative visualization systems, such as sense.us [16] and Many Eyes [26], allow users to link free comments and graphic annotations to specific visualization views. The comments and annotations are usually manually generated by the users and contain high level semantic information of the insights associated with the views. The comments are often posted in a discussion forum where they can be retrieved by other users through browsing or keyword searches. CommentSpace [28] and Sandbox [29] go further by allowing users to tag insights and link them for supporting or conflicting hypotheses. However, manually generated comments can be incomplete and different users may describe the same information using different terms, possibly harming the effectiveness of insight retrieval. Users often have to manually examine lengthy discussion threads for reviewing insights. Our approach provides much richer views and interactions than existing collaborative visualization forums. In addition, it makes use of the rich semantic information automatically collected by a state-of-the-art insight annotation technique, which greatly improves the effectiveness and efficiency of common ground construction.

**Using structured insight organization:** Aruvi [24] and Analyst Notebook [1] allow users to manually organize findings and their relationships in a network structure. Tree Trellis and Table Trellis [8] support aggregation and comparison of linked free-text claims. Evidence matrices [4] aggregate and make inferences according to analytic evidence. Rows contain multiple hypotheses and columns

contain collected evidence. The above approaches require analysts to manually organize and associate insights and provide only static views for examining their structure. They are difficult to scale to the fast growing ACVA insights for users with diverse information needs. Our approach is much more scalable by allowing different users to efficiently construct their own overview.

### 3 TOOLKIT OVERVIEW

We have implemented our common ground construction approaches in ManyInsights, a web-based system for asynchronous collaborative analysis of multidimensional data [6]. Our previous work on semi-automatic insight annotation, named Click2Annotate [6], is integrated in ManyInsights and used as a basis of the proposed approaches in this paper. Besides commonly supported web-based visualization tasks such as uploading data, creating visualizations (e.g., ScatterPlot and Parallel Coordinates), and exploring the visualizations for insights, ManyInsights allows users to effectively conduct the common ground construction tasks identified in Section 1.

The common ground construction pipeline in ManyInsights consists of the following steps:

1. The system collects semantic information of the insights (we call it **insight contents**) when users semi-automatically annotate insights using Click2Annotate [6] (see Section 4.1). The contents are stored in an insight database, upon which insights can be searched and pairwise insight correlations can be calculated (see Section 4.2).
2. The users start their visual exploration by retrieving insights by contents and browsing them in an automatically generated overview (see Section 5). The overview reveals insight clusters, their temporal evolution, and their major semantics.
3. The users iteratively refine the overview to reflect their aims and concentrations. They can dynamically adjust the weights of different contents in the insight correlation calculation (see Section 4.2) to associate insights differently, or refine the search with instant visual feedbacks (see Section 5.2).
4. The users select and mark insights of interest from the overview. The selected insights are compared and examined in detail for sense making (see Section 5.5).
5. The users preserve and share their key findings and hypotheses for future use by other collaborators (see Section 5.7).

## 4 INSIGHT ANNOTATION AND CORRELATION CALCULATION

### 4.1 Insight Annotation

The following information is recorded for each insight annotated in ManyInsights: (1) data contents, such as dataset names, types of the insights (e.g., clusters, outliers, rank, correlation, and etc.), relevant dimensions and data items, and essential characteristics of the insights (such as the mean of clusters); (2) user-generated semantic information, such as hypotheses associated with the insights and tags given by users; (3) meta information, such as the name of the author who annotated an insight and a timestamp recording when the insight was annotated.

A semi-automatic insight annotation approach named Click2Annotate [6] is used to collect the above information in ManyInsights. Using pre-defined or customized annotation templates for typical analytic insight types, users can annotate most insights with a few mouse clicks. Please refer to [6] for details of the annotation process. Note that our common ground construction approaches can also be used on insights annotated by other approaches, given that all or part of the above information is provided in the annotations.

### 4.2 Insight Correlation Calculation

The correlation between two insights is calculated as a weighted sum of the following similarity measures. The measures are normalized to the range 0 to 1 and their weights can be interactively adjusted by users:

- **Closeness in data space:** We use data similarity ( $Sim_{data}$ ) to capture the closeness of two insights in the data space. It is calculated using Exact Transformation Measure (ETM) (please refer to [30] for details). ETM is based on transform cost and can handle subsets with different data populations efficiently, which makes it suitable in our application to calculate the closeness between insights of different types (e.g., an outlier and a large cluster). If two insights are not in the same dataset, their data similarity is set to 0.
- **Shared dimensions and data items:** We use dimension similarity ( $Sim_{dim}$ ) and data item similarity ( $Sim_{item}$ ) to capture the relationships between two insights involving the same dimensions and data items, respectively. Note that the two insights can be about different datasets that share dimensions and data items. By considering each dimension as a weighted keyword and an insight as a document, we use an improved cosine similarity measure [3] to calculate  $Sim_{dim}$ :

$$Sim_{dim}(I_i, I_j) = \frac{\sum_{k=1}^K (\log_2 \frac{N}{n_k} W_k) (\log_2 \frac{N}{n_k} W_k)}{\sqrt{\sum_{k=1}^K (\log_2 \frac{N}{n_k} W_k)^2 \cdot \sum_{k=1}^K (\log_2 \frac{N}{n_k} W_k)^2}} \quad (1)$$

where  $I_i$  and  $I_j$  are two insights,  $W_k$  is the importance of a shared dimension  $k$ ,  $K$  is the total number of shared dimensions,  $N$  is the total number of insights, and  $n_k$  is the number of insights sharing the dimension  $k$  inside  $N$ . Data item similarity ( $Sim_{item}$ ) is computed in a similar way.

Each dimension or data item is assigned an importance in the above calculation. We allow users to interactively set the importance of individual dimensions and data items according to their exploration focus. For example, by assigning a high importance to a dimension of interest, insights containing this dimension are considered closely related.

- **Shared fact type:** People may want to examine and compare insights of the same type, such as all insights about ranks, at the same time. Type similarity ( $Sim_{type}$ ) allows users to conduct this task.  $Sim_{type}(I_i, I_j)$  is 1 if insight  $I_i$  and  $I_j$  have the same type and 0 otherwise.
- **Shared hypotheses:** In collaborative visual analytics, users often use insights as evidence to support or refute their hypotheses [19]. Insights that are associated with the same hypothesis may have semantic relationships. Hypothesis similarity  $Sim_{hypo}(I_i, I_j)$  is 1 if insight  $I_i$  and  $I_j$  are associated with the same hypothesis and 0 otherwise.
- **Shared tags:** Tags with descriptive text can be attached to an insight to express user interest, record their evaluations, and convey the semantic properties of that insight [28]. They are manually generated by users and shared among insights in ManyInsights. Sharing tags indicates semantic relationships between insights. Tag similarity ( $Sim_{tag}$ ) is used to capture such relationships. In particular, each tag is considered a keyword and  $Sim_{tag}$  is also calculated using Equation 1 with user-specified importance.
- **Author:** Users often want to examine insights from the same author. We define author similarity  $Sim_{author}(I_i, I_j)$  as 1 if insights  $I_i$  and  $I_j$  are created by the same user and 0 otherwise.

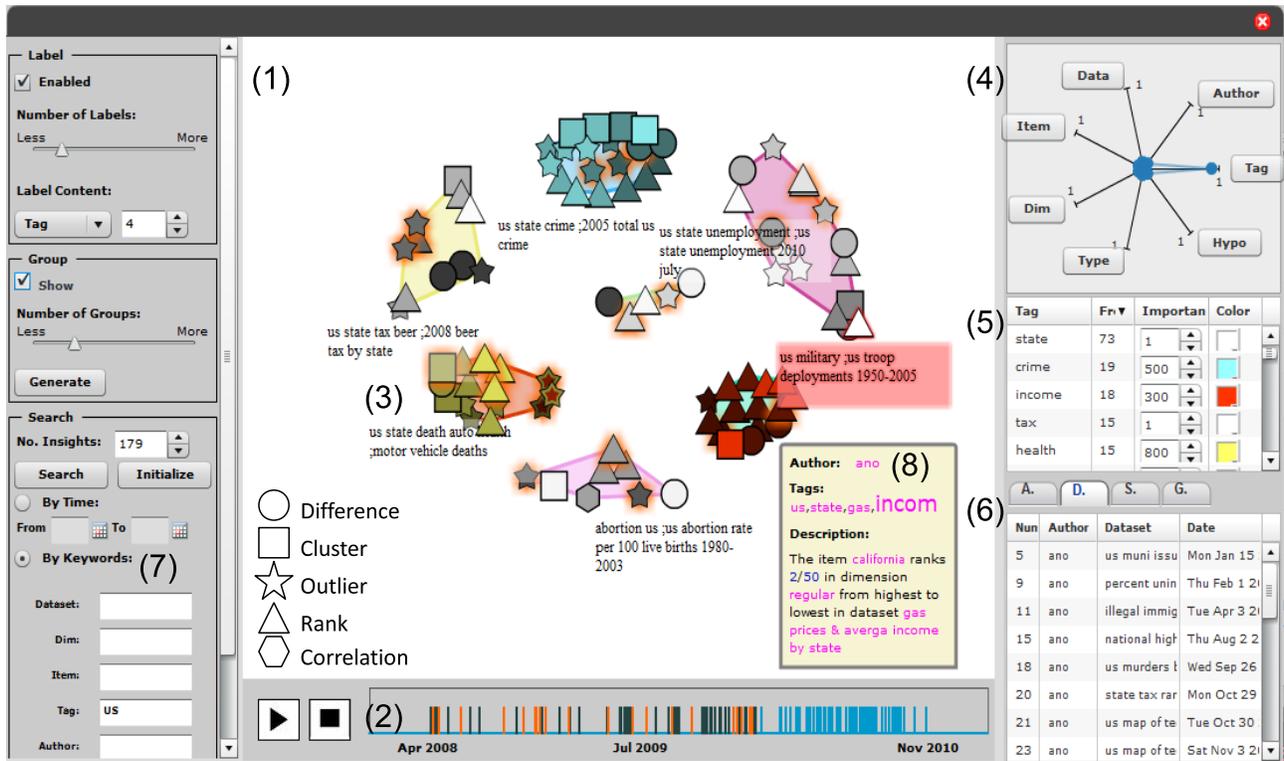


Figure 1: The overview interface. The left part includes label and group controls and search interfaces. The center part is the dynamic clustering display and the timeline (bottom). The right part includes weight controller, keyword tables, and insight tables. In the dynamic clustering display, the insights are represented by shaped particles, colored according to the keywords they contain (“health” - yellow, “income” - red, and “crime” - blue). Insights containing the keyword “Texas” are selected and highlighted by orange halos. The weight of tag similarity is set to 1 and others are set to 0.

We calculate insight correlation ( $Cor_{insight}$ ) between any pair of insights  $I_i$  and  $I_j$  using the similarity measures described above:

$$Cor_{insight}(I_i, I_j) = w_{data}Sim_{data}(I_i, I_j) + w_{dim}Sim_{dim}(I_i, I_j) + w_{item}Sim_{item}(I_i, I_j) + w_{type}Sim_{type}(I_i, I_j) + \dots + w_{author}Sim_{author}(I_i, I_j)$$

where  $w_{data}$ ,  $w_{dim}$ ,  $\dots$ ,  $w_{author}$  are user-controllable weights. The sum of all weights equals 1. By adjusting the weights, users can organize and associate insights according to a variety of interests. For example, if the users are interested in authors, they can set  $w_{author}$  to 1 and other weights to 0. Then the insights will be grouped by their authors in the visualization. Section 5.2 presents how users interactively adjust the weights and receive instant feedbacks in detail.

Among the similarity measures, the most computationally heavy ones are the data closeness, dimension, data item, and tag similarity calculations. They are either  $O(N^2)$  or  $O(N^3)$  approaches. However, once the calculation is performed, an individual measure is stored and only re-calculated when users adjust the importance (e.g. changing the importance of a dimension in dimension similarity). Hence, the modification of weights in insight correlation calculation can be performed efficiently.

## 5 VISUALIZATION

Multiple coordinated views are provided in our system to support the common ground construction tasks identified in Section 1. To generating overviews (T1), the dynamic clustering display, the con-

tent cloud, the timeline, and the insight table are provided. The dynamic clustering display (see Figure 1(1)) reveals the correlations among the insights by placing related insights close to each other. It can also reveal the temporal evolution of the insights through controllable animations. The content cloud visually summarizes the most significant semantic contents of an insight group (see Figure 2(a)(b)). The timeline allows users to examine the insights along a time axis (see Figure 1(2)). The insight table allows users to access the insights in a familiar table metaphor (see Figure 1(6)).

Our system also allows users to compare groups of insights, examine them in detail, and preserve and share the findings they derived (T2). The region graph (see Figure 3) presents the relationships among a group of insights in detail. It also allows the users to compare two groups of insights for shared or distinct information (see Figure 4). The users can further examine an insight in an annotation window (see Figure 2(c)) or revisit the data in a multi-dimensional display (see Figure 5(a)). They can also preserve and share their common ground construction results by creating new insights. In the following sections, we present the views and interactions in detail.

### 5.1 Content Cloud

Figure 2(a) shows a content cloud of 179 insights. Each tag is a frequently shared keyword in the insight annotations. Keywords from the same type of contents are grouped together with the same color. For example, the blue keywords are all from dimension names. The size of a keyword indicates its frequency of occurrence in all insights or the importance of the keyword assigned by users. In a cloud of all insights in an overview (see Figure 2(a)), the tags are ordered according to descending frequencies. In a cloud of a group



Figure 2: (a) A content cloud shows the most significant contents of 179 insights. (b) A content cloud shows the most significant contents of a group of insights selected from the 179 insights. Content colors: tag - pink, dimension - blue, data item - yellow, and type - red. (c) An annotation window, showing the visualization and contents of an insight in detail.

of insights selected from the overview (see Figure 2(b)), the tags are ordered by their TF-IDF weights [20] to emphasize salient features of the group. The TF-IDF weights are calculated using an improved TF-IDF weighting algorithm described in [27]. In Figure 2(b), keyword “smoke” is ranked high since it is significant in this group, despite that its global frequency is low.

**Interactions:** Content clouds provide users a convenient way to start exploring a set of unknown insights. In particular, by clicking a keyword in a cloud, users can select all insights with this keyword in their contents. They can also set the colors or importance of keywords to control the dynamic clustering display from the content cloud.

## 5.2 Dynamic Clustering Display

Figure 1(1) shows 90 insights in the dynamic clustering display. The insights are represented as particles with a variety of shapes indicating their types (see the shape legend in Figure 1(1)). The luminance of the particles indicates the age of the insights (the darker, the older). The insights are clustered according to their correlations. Please refer to [3] for details of the underlying force-based dynamic system of this view.

Labels are automatically generated for clusters of insights (see Figure 1(3) for an example). We allow users to define a combination of insight contents as a label to represent the semantics of a cluster. To avoid lengthy labels, we only use the top- $N$  most frequent keywords.

**Interactions:**

**Dynamic clustering:** Users can dynamically cluster the insights in this view to reflect their current exploration interest. For example, by setting the tag weight to 1 through the star glyph (see Figure 1(4)), insights are grouped by their tags. Users can also interactively adjust the importance of keywords from the keyword table <sup>3</sup> (see Figure 1(5)) to cluster the insights by keywords of interest,

<sup>3</sup>Keywords stored in this table are insight contents, such as dimension

as shown in Figure 1(1).

**Animation:** Users can play animations to examine temporal evolution of the insights. During the animation, insights are continuously injected into the display in chronological order. The layout gradually evolves to reveal how clusters are formed and evolving over time. Users can use play and stop buttons to pause and resume the animation. They can also jump to a particular moment using the timeline (see Section 5.3).

**Tracking keywords:** To track insights with keywords of interest, users can assign colors to them from the keyword table (see Figure 1(5)). An insight can have multiple colors if they contain multiple keywords of interest (see Figure 5(b-3)).

**Grouping and tracking insights:** The system can automatically divide the insights into groups according to a user-defined dissimilarity threshold. In Figure 1(1), the automatically generated groups are highlighted in different background colors. The groups are stored in a group table for further operations, such as comparison, viewing content clouds, and etc.. During an animation, users can highlight groups using colored halos for tracking their evolution (see Figure 5(b-3)).

## 5.3 Timeline

According to the experimental evidence reported in [17], users can easily understand the development of insights by organizing them in chronological history. Following this, the timeline represents insights as bars along a time axis (see Figure 1(2)), whose distribution reflects the temporal distribution of insights in the overview.

**Interaction:** The timeline is coordinated with the dynamic clustering display. Clicking a bar will navigate users to that particular moment in an animation. Bars in blue are insights yet to be displayed in the dynamic clustering display.

## 5.4 Insight Table

Users can examine a group of insights, all displayed insights, or the entire insight collection from the insight table (see Figure 1(6)). They can sort the insights by different contents and manually construct insights groups.

## 5.5 Region Graph

A region graph, inspired by the substrate graph [21], can have one or two columns, each for an insight group examined. In Figure 3, users examine the details and the relations among insights in the same group. In Figure 4, users compare and associate two groups of insights. Nodes displayed in the region graphs represent insights, which have the same visual representations (e.g., shapes and colors) with the insight particles in the dynamic clustering display.

**Layouts:** Layouts are based on user-defined content substrates, which are non-overlapping regions in which node placement is based on insight contents. For example, in Figure 3, each region, assigned to a distinct color, represents a dataset appearing in the insights. The name of the dataset is specified by the label on the bottom of the region. A region is evenly divided into rows, each of which presents a dimension of the dataset appearing in the insights. The label of a dimension is placed on the left. Users can also map other contents to rows (e.g., tags and authors). The horizontal placement of nodes is tied to their age, the oldest on the left and the newest on the right. The vertical position of a node is determined by the region and row contents of the insight. If an insight contains multiple dimensions inside a dataset, nodes will be placed in each of the associated rows, with the topmost one drawn in solid and others drawn with a blurring effect.

Users can learn the basic semantics of an insight by its spatial location and shape without reading its annotation. They can also easily identify how insights are temporally related. Another advantage

names, data item names, tags, authors, and etc..

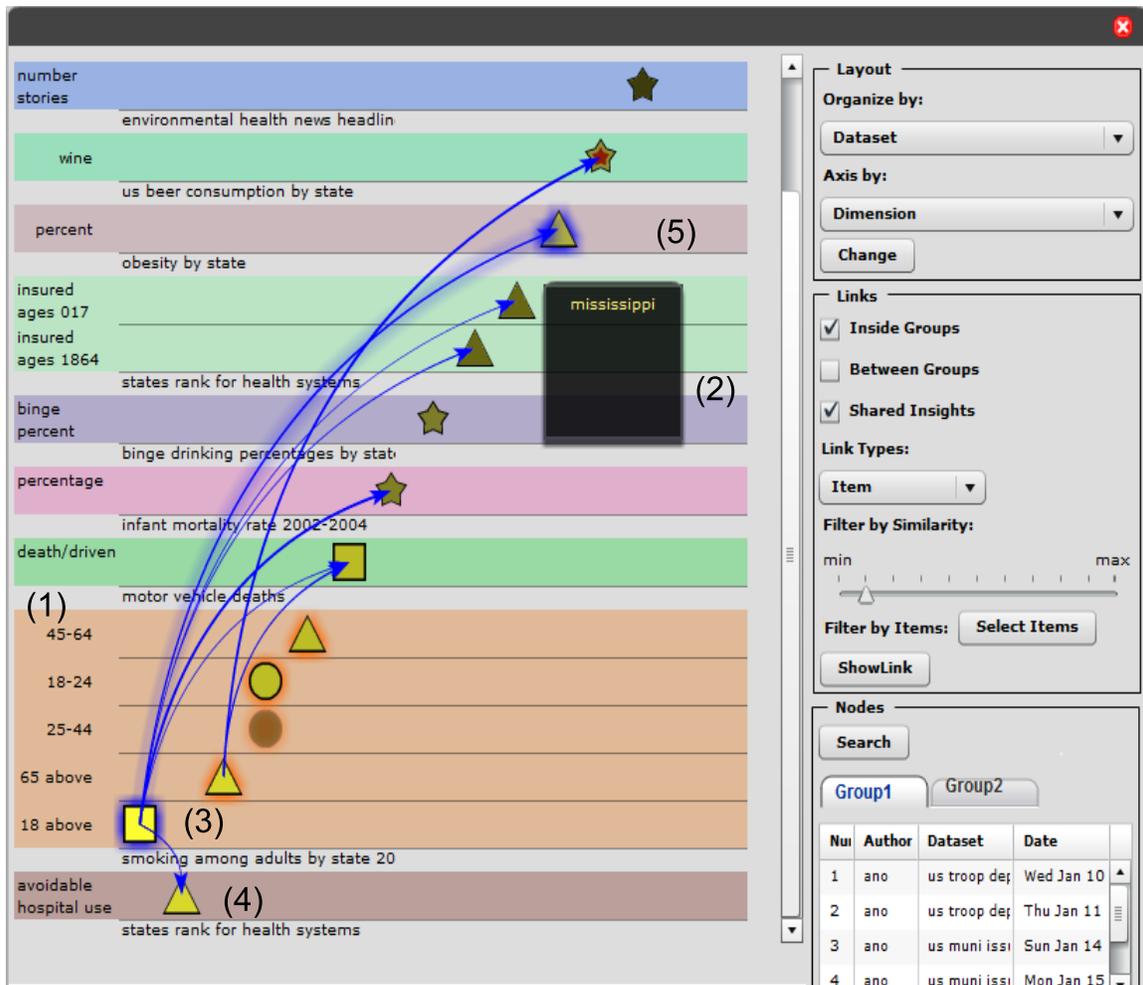


Figure 3: A region graph. The left part shows insights and their relationships. The right part includes layout controls, link visibility controls, and insight tables. Nodes with “health” are yellow and nodes with “income” are red. Data item links are blue.

is that the proportionally-sized regions indicate the relative cardinality of each region. For example, in Figure 3, users can quickly identify that the dataset “smoking among adults by state” (see Figure 3(1)) has more dimensions involved in the insights than other datasets.

**Links:** The region graph represents insight relationships using directed links between nodes. Pairwise insights could have multiple relationships displayed by different colors, such as shared tags and shared data items. The thickness of a link indicates the corresponding similarity measure. Users can hover their mouse over a link to examine the relationship in detail. For example, in Figure 3(2), two nodes are connected with a data item link since they share the same data item “Mississippi”.

**Alignment for two groups:** To compare two insight groups, the region graph horizontally places them in two columns (see Figure 4). To help users identify shared regions and rows (they indicate shared contents), we consider two goals when laying out the graphs: (1) any pairwise shared regions/rows should be placed closely to each other; and (2) all shared regions/rows should be grouped and placed in prominent positions (e.g., the topmost position). To achieve these goals, the following iterative, greedy algorithm is used (we assume that both columns use dataset-dimension layouts in the description):

**Step 1:** We denote a pairwise shared regions that represent the same dataset between two columns as  $PR_{common}$ . We denote the

difference of the height between  $PR_{common}$  as  $Diff_{pr}$ . Identify all  $PR_{common}$  between two columns and put them into a sorted queue (denoted as  $Q_{pr}$ ) where the one with the smallest  $Diff_{pr}$  is first. For each column, the rest of the regions (denoted as  $R_{uncommon}$ ) are placed into a sorted queue (denoted as  $Q_r$ ) where the one with smallest height is first.

**Step 2:** Take the  $PR_{common}$  at the front of  $Q_{pr}$ . Place each region of  $PR_{common}$  at the topmost position of the corresponding column. For each column, compute the total height of regions that have been placed. Then compute the difference of total height between two columns (denoted as  $Diff_{tr}$ ). For the column with the smaller total height, take the  $R_{uncommon}$  at the front of its  $Q_r$  and place it at the topmost position and then update  $Diff_{tr}$ . Repeat this step until  $Diff_{tr}$  reaches the minimum value.

**Step 3:** Repeat Step 2 until the  $Q_{pr}$  is empty.

**Step 4:** For each column, take the  $R_{uncommon}$  at the front of its  $Q_r$  and place it at the topmost position of the column. Repeat this step until the  $Q_r$  is empty.

**Step 5:** For each  $PR_{common}$ , sort their rows by identifying overlapping rows and place them on the topmost position of the region.

Figure 4 shows a result of the graph alignment. Links crossing two graphs represent insight relationships between the groups. If an insight is contained in both groups, the corresponding nodes in each column are connected by red, undirected dot links (see Figure

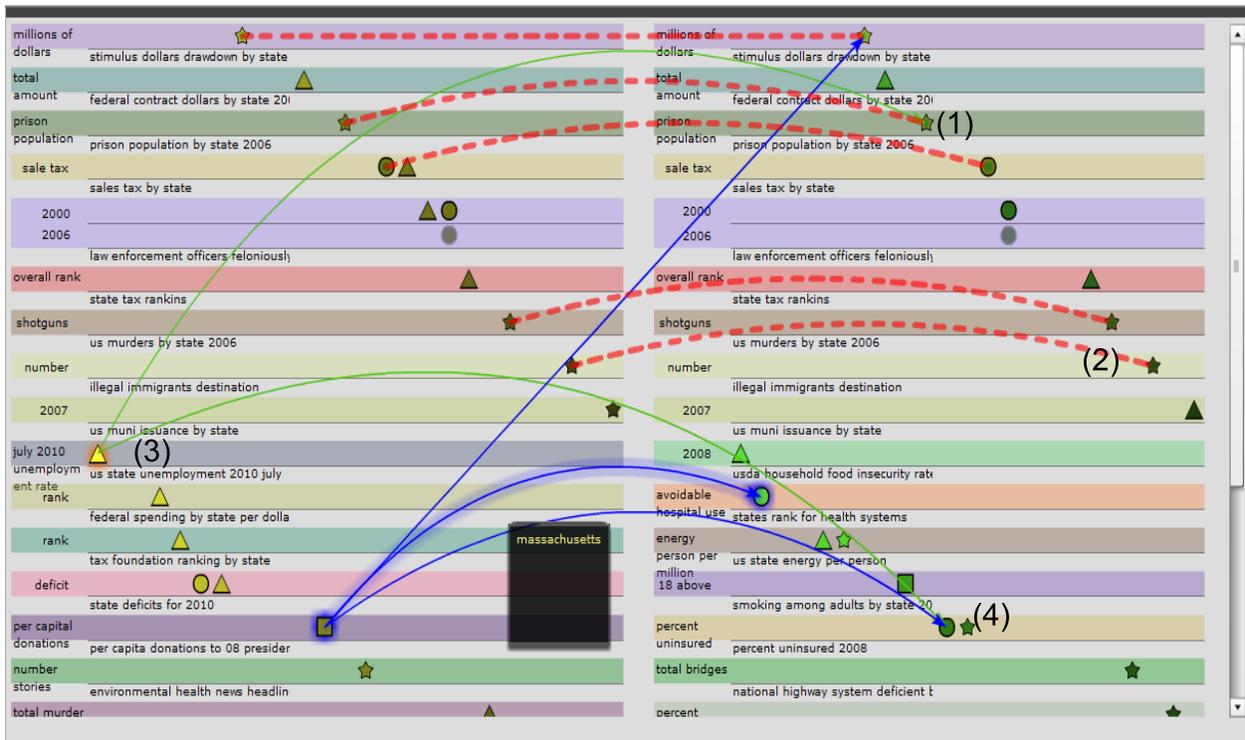


Figure 4: Compare and associate two insight groups using a region graph. Each column represents an insight group. Shared regions are indicated by the same colors and labels. Shared insights are connected by red, undirected dot links. Nodes with “Texas” are green and nodes with “California” are yellow. Data item links are blue and tag links are green.

4(1)).

#### Interactions:

**Changing layout:** Users can change the contents mapped to the regions and rows through a control panel, and hence organize the insights in different ways. They can also manually adjust the order of vertical placement of the regions for a graph to place regions of interest in prominent positions.

**Filtering links:** Users have multiple options to control the visibility of links to reduce clutter (see Section 6 for examples).

**Visualizing data:** Users can select dimensions from the region graph to open a multidimensional display (see Figure 5(a-1)). In the display the related insights will be highlighted, with flags indicating their types [6]. In this way, users can explore the visualization for new insights or examine existing insights for refinement.

### 5.6 Other Interactions

**Search:** Users can search insights by a variety of insight contents (see Figure 1(7)).

**Manual selection:** Users can manually select insights from any views by clicking an insight.

**Annotation card:** Users can hover their mouse over an insight to examine its annotation card (see Figure 1(8)). It provides a preview of the insight by summarizing its essential information using descriptive language [6]. Keywords in the annotation card are highlighted. The keywords with high importance are enlarged.

### 5.7 Preserving and Sharing Results

Collocated collaborators preserve and share their common ground construction results in a shared work space. In ACVA, users with different interests may want to construct their common ground in different ways. In addition, a static common ground view is not suitable for the dynamic environment of ACVA. Therefore, rather

than providing a shared work space for persevering and sharing common ground, we allow users to record their results of common ground construction through a special type of insights, namely the hypothesis insights. A hypothesis insight contains a tag given by users. The tag is also assigned to insights related to it as their hypothesis contents. A screenshot of the common ground view is attached to the insights and users can also make free notes. To review the work of other users, a user can search for hypothesis insights. Furthermore, they can organize the insights by their hypotheses (see Section 4.2).

## 6 USE CASE

To demonstrate how the system can be used to explore a large amount of insights generated by a diversity of real users from real datasets, we imported 239 insights of 102 datasets from Many Eyes [2] to ManyInsights (ManyInsights is not public yet and there are no such real data available). Many Eyes [2] is a popular web-based collaborative visual analytics system, where users visually explore datasets contributed by themselves or others. They share their insights by posting comments linked to specific visualization displays. A large number of insights are reported daily by users from a variety of domains [26]. These insights come from a wide range of datasets.

All the insights we imported were generated from multidimensional visualization displays. They all contain one of two popular tags: “US” and “world”. Most of them were generated from users’ comments linked to visualization displays. We reviewed the comments and displays to extract their semantic contents, and manually generated a formalized annotation as described in Section 4.1 for each insight.

Consider Mary and Tom, two graduate students majoring in sociology in different universities, are both investigating the quality

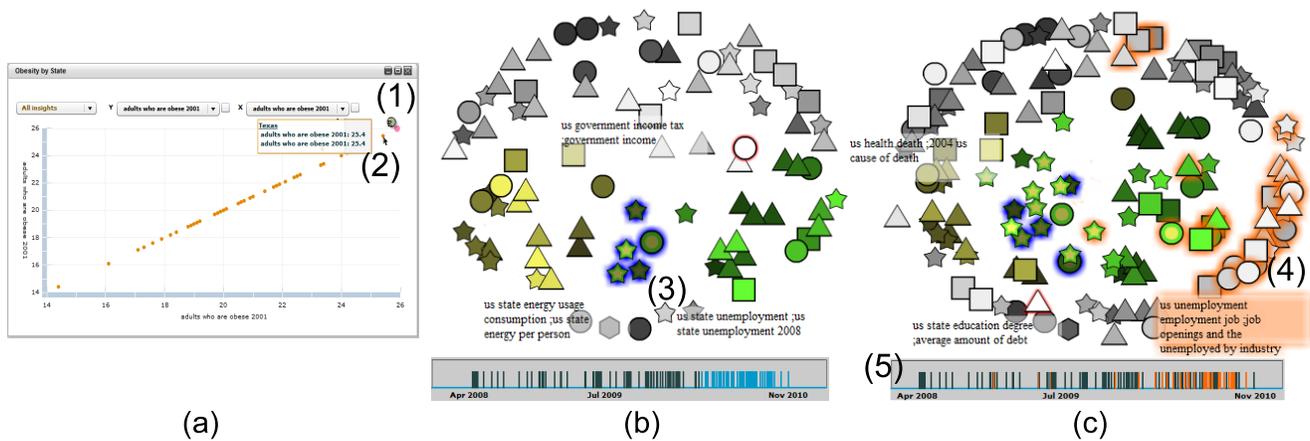


Figure 5: (a) Revisit the visualization of the dataset “obesity by state”. An insight is displayed in the visualization. (b) 97 insights in dynamic clustering display (Jan. 2010). (c) 169 insights in dynamic clustering display (Nov. 2010). Data item keyword colors: “Texas” - green and “California” - yellow. Content weights: data item - 0.6 and tag - 0.4. (3) is a cluster of insights with both keyword “Texas” and “California”.

of living by state in America. To acquire information, they search Many Eyes for comments of related data. A large number of comments are returned and it is time consuming to read them one by one. Therefore, they use our system to review and organize these insights.

**Identifying Clusters:** One day, Mary logs into the system and searches insights with the tag “US” (see Figure 1(7)). 179 insights are returned and displayed in the overview interface. From the content cloud (see Figure 2(a)), she immediately identifies three tags - “health”, “income”, and “crime”, which are important factors related to the quality of living. To cluster the insights by these factors, she increases the weight of the tag similarity to 0.9 and sets the importance of these three tags to a value much higher than the remaining tags. She also assigns colors to insights with the three tags, for example, yellow for insights with the tag “health”. She starts an animation in the dynamic clustering display, and soon notices several clusters (see Figure 1(1)). She is interested in the group of insights with the tag “health” (see Figure 1(3)). She pauses the animation and selects this group for further exploration.

**Examining a cluster in detail:** Mary opens a content cloud for this group, as shown in Figure 2(b). The data item “Texas” in this figure catches her attention since it is significant in this cloud. Mary then examines the group in a region graph (see Figure 3). She quickly locates a dataset, named “smoking among adults by state” (see Figure 3(3)), involving more dimensions and insights than others. She thinks that the smoking population is quite related to the quality of living, so she focuses her exploration on this dataset. Based on the shapes and the annotation cards of the nodes, she quickly learns the essential content of insights in this dataset. To further investigate how insights about smoking are related to other insights, she displays their links to others. An insight about a cluster (see Figure 3(3)) catches her attention since it has many links to other insights. She clicks the insight to explore it in the annotation window (see Figure 2(c)). By reading the annotation, she realizes that “Mississippi”, “Texas”, and other two states have similar high values in percentage of smoking people. Mary finds many interesting links from this insight to other insights. For example, “Texas” is highly ranked in percentage of smoking people (see Figure 3(3)) and is ranked low in health systems (see Figure 3(4)).

**Hypothesis generation and validation:** Following the link between (3) and (5) in Figure 3, Mary observes that “Mississippi” is also highly ranked in percentage of obesity (see Figure 3(5)). Mary thus makes a hypothesis that “Texas” is also high in obesity. To validate her hypothesis, she selects the dataset “obesity by state” and

the dimension “percentage” from the graph to create a visualization (see Figure 5(a)). From the visualization, she easily discovers that “Texas” (see Figure 5(a-2)) is ranked the second highest in obesity percentage just behind “Mississippi” (see Figure 5(a-1)). Therefore, her hypothesis is confirmed. Mary makes an annotation for this new insight.

Mary then highlights all insights with “Texas” in the dynamic clustering display (see Figure 1(1)). She observes that several other insight clusters also contain insights about “Texas”. By examining them she collects more insights showing low quality of living in “Texas”. She saves this result in a new hypothesis insight “living quality in Texas is low” and posts it in ManyInsights. The hypothesis is also attached to the relevant insights Mary found.

**Comparing insight groups:** Later on, Tom logs into the system and organizes the insights by their hypotheses in the dynamic clustering display. He finds the group of insights showing low quality of living in “Texas”, which was discovered previously by Mary. Tom has investigated the quality of living in “California”, so he is interested in continuing Mary’s investigation to compare insights about “Texas” and “California”. He does so by making an insight group for “Texas” and another for “California” and compares them in the region graph. He quickly locates several shared insights, datasets, and dimensions on the top positions of the region graph (see Figure 4). For example, he finds that “Texas” and “California” are extremely high in prison population (see Figure 4(1)) and illegal immigration population (see Figure 4(2)). Tom then selects two interesting insights in the “California” group (left column) and examines their links to the “Texas” group (right column). Here, he chooses to display the tag and data item links and filters out links that contain general terms such as “US” and “Texas”. While examining this display, he realizes that the insights about “unemployment population” (see Figure 4(3)) are related to insights about “uninsured population” (see Figure 4(4)) and “prison population” (see Figure 4(1)) because they share the same tags.

**Tracking insight evolution:** Tom examines the temporal trends by playing the animation in the dynamic clustering display. Figure 5(b) and Figure 5(c) show two screenshots during the animation, where the insights about both “Texas” and “California” Tom has explored are highlighted in blue halos. During the animation, Tom notices that the highlighted group shown in Figure 5(b-3) gets much bigger by adding several outliers about “Texas” and “California”. He also notices a newly formed insight group (indicated by light darkness, see Figure 5(c-4)). He quickly learns that this group is about “unemployment” and “job” issues according to the label and the con-

tent cloud. He highlights insights with these keywords. From the timeline view (see Figure 5(c-5)), he finds that most highlighted insights (they are in orange) were developed after Jul. 2009. Based on this pattern, Tom thinks that more and more people are concerned about how “unemployment” affects their quality of living after the economic crisis. Therefore, “unemployment” and “job” should be important topics for his further investigation of quality of living.

## 7 PRELIMINARY EVALUATION

We conducted a preliminary user study to evaluate our system with the following goals: (1) to understand how users construct common ground using our system, and (2) to learn how well the various provided features support their efforts in this process. Our study was essentially a longitudinal analysis [22] that was focused on small numbers of users (both experts and novice users) over long time periods. In particular, our study consisted of two sessions, namely an individual analysis session and an asynchronous collaborative session. They were designed to simulate real world collaborative analytic process.

### 7.1 Procedure

We first ran a 2-week individual analysis session with 5 graduate students, all of whom were computer science major and had participated in our previous user study with Click2Annotate [6]. They were asked to explore two datasets individually using either scatterplot or parallel coordinates, and annotate their insights using the Click2Annotate tool. The first dataset is the NFL football season data (75 dimensions and 32 data items). Participants were instructed to discover insights about key factors for a football team to win more games. The second dataset is the fast food nutrition data (9 dimensions and 274 data items). Participants were instructed to discover insights that determine the healthiest fast food restaurant. The tools and data were installed in portable laptops so participants could perform the tasks whenever convenient. They tagged insights with predefined keywords or those created by themselves. The generated insight annotations were automatically collected. After the first session, we collected 43 insights for the NFL football season data and 67 insights for the fast food nutrition data.

Next, the asynchronous collaborative session was conducted. Participants were instructed to use the system separately to review the insights created in the first session. Seven graduate students attended this session, including five existing students (experts) and two novice students not participating in the first session. After training, participants were instructed to review the entire insight collections for each dataset, following by free exploration in which they queried and reviewed insights of interest. There was no time limit but all the participants completed their work within 3 hours. We observed and recorded the screen of the whole process and conducted interviews after the session.

### 7.2 Findings

Our key findings were derived from the asynchronous collaborative session with observations and users’ feedback. First, we observed that author information were commonly used to organize insights in the initial period of common ground construction. In particular, three participants used authors to group insights at the beginning. Thereafter, they used region graphs looking for shared information between authors. Four participants used color encodings to distinct insights generated by different authors. One participant grouped insights by dimensions and colored them by authors, when reviewing fast food nutrition data. He then divided five authors into three groups according to their exploration focuses. This finding answers the call by Robinson [19] for supporting role assignment in collaborative visual analysis.

Second, during the free exploration process, we observed that the rich set of views provided by our system allowed experts and

novice users to use different exploration strategies. Most experts first searched insights generated by themselves and select them on the timeline. Then, they used the timeline to navigate to a particular moment, created a group for the selected insights, and highlighted the group. By tracking the evolution of this group and manipulating the visual structure, they continuously added correlated insights into the group. They further investigated the group using either annotation cards or region graph to browse and relate them. In this process, searching and sorting on insight or keyword tables are the most frequent actions taken by the experts. In contrast, we noticed that novice users mostly relied on content clouds to manipulate and organize insights. They would spend longer time on the clouds and find important items to guide them in further exploration. However, both novice users and experts used content clouds intensively in exploring individual groups.

The feedbacks from the participants indicated that the system helped them understand and manipulate each other’ insights, and was useful in complex collaborative tasks. When asked about specific features, participants were greatly impressed by the dynamic organization of insights, the abundant interactions (e.g., color coding, annotation card, and multiple selection tools), and the ability of comparing insight groups with region graphs. One participant with Many Eyes experiences compared our system to Many Eyes, “*I really like the way to visually present and group insights. Even more I can change the group as will. In Many Eyes, I have to endlessly search keywords and read hundreds of posts. It is really boring*”. Regarding specific tasks, one participant emphasized that grouping insights by data similarity was particularly powerful in understanding NSF football data, “*When I originally explored this data, It really messed me up since there is more than 70 dimensions! But after I grouped insights and reviewed them, I suddenly got some interesting correlations about dimensions. It really helps*”.

## 8 CONCLUSION AND FUTURE WORK

The visual analytics toolkit presented in this paper is among the first efforts on supporting effective common ground construction for ACVA. Our case studies and user studies suggested that such a toolkit, which allows users to quickly construct exploratory overviews for a large amount of evolving insights, flexibly study their relations and patterns, as well as effectively share and exchange insights, greatly reduces human efforts and enhances the visual sense making process in asynchronous collaborations. We argue that such features are essential for ACVA and should be supported by all ACVA systems. Our preliminary user study has also led to many valuable insights for ACVA researchers, such as the different exploration strategies used by domain experts and novices in ACVA common ground construction. In addition, our approach, namely the semi-automatic annotation combined with semi-automatic common ground construction approach, is general enough to be extended to other data types, such as geospatial data and graph data. The approach is independent from the visualization platforms where the insights are discovered and thus it can be used in a wide range of collaborative visual analytics applications.

Future work will focus on three main directions. First, we plan to extend the toolkit to support collaborative work on insights generated from miscellaneous data sources and by different visualization tools. Such a generalized approach accommodating various datasets and scenarios will benefit a diverse range of communities across scientific and social domains. Second, we will improve our system with enhanced functionalities, including insight notification, insight recommendation, and hypothesis generation, to support more comprehensive and encouraging collaborative exploration. Finally, we will publish ManyInsights online for public tests. We will collect user feedbacks to evaluate its utility, usability, and scalability, and thus refine our toolkits. Eventually, we will promote it to a variety of realistic applications.

## ACKNOWLEDGEMENTS

This work was performed with partial support from the DHS Visual Analytics for Command, Control, and Interoperability (VACCINE) Center of Excellence, under the auspices of the SouthEast Regional Visual Analytics Center. The work is also partially supported by National Science Foundation under grant number IIS-0915528, IIS-0916131 and NSFDACS10P1309.

## REFERENCES

- [1] Analyst's notebook. <http://i2group.com/Analysts-Notebook>.
- [2] Many-eyes. <http://www.many-eyes.com>.
- [3] J. Alsakran, Y. Chen, Y. Zhao, J. Yang, and D. Luo. Streamit: Dynamic visualization and interactive exploration of text streams. *Proc. IEEE Symposium on Pacific Visualization*, 2011.
- [4] D. Billman, G. Convertino, J. Shrager, P. Pirolli, and J. Massar. Collaborative intelligence analysis with cache and its effects on information gathering and cognitive bias. *Proc. HCI Consortium Workshop*, 2006.
- [5] J. Carroll, M. Rosson, G. Convertino, and C. Ganoe. Awareness and teamwork in computer-supported collaborations. *Interacting with Computers*, 18(1):21–46, 2005.
- [6] Y. Chen, S. Barlowe, and J. Yang. Click2annotate: Automated insight externalization with rich semantics. *Proc. IEEE Symposium on Visual Analytics Science and Technology*, pages 155–162, 2010.
- [7] Y. Chen, J. Yang, and W. Ribarsky. Toward effective insight management in visual analytics systems. *Proc. IEEE Symposium on Pacific Visualization*, pages 49–56, 2009.
- [8] T. Chklovski, V. Ratnakar, and Y. Gil. User interfaces with semi-formal representations: a study of designing argumentation structures. *Proc. ACM Conference on Intelligent User Interfaces*, pages 130–136, 2005.
- [9] M. Chuah and S. Roth. Visualizing common ground. *Proc. 7th International Conference on Information Visualization*, pages 365–372, 2003.
- [10] H. Chung, S. Yang, N. Massjouni, C. Andrews, R. Kanna, and C. North. Vizcept: supporting synchronous collaboration for constructing visualizations in intelligence analysis. *Proc. IEEE Symposium on Visual Analytics Science and Technology*, pages 107–114, 2010.
- [11] H. Clark and S. Brennan. Grounding in communication. in perspectives on social shared cognition. *American Psychological Association*, pages 127–149, 1991.
- [12] G. Convertino, C. Ganoe, W. Schafer, B. Yost, and J. Carrol. A multiple view approach to support common ground in distributed and synchronous geo-collaboration. *Proc. Coordinated and Multiple Views in Exploratory Visualization*, pages 121–132, 2005.
- [13] G. Convertino, H. Mentis, M. Rosson, M. Rosson, J. Carrol, A. Slavkoci, and C. Ganoe. Articulating common ground in cooperative work: content and process. *Proc. ACM SIGCHI Conference on Human Factors in Computing Systems*, pages 1637–1646, 2008.
- [14] D. Gotz and M. Zhou. Characterizing users visual analytic activity for insight provenance. *Proc. IEEE Symposium on Visual Analytics Science and Technology*, pages 123–130, 2008.
- [15] J. Heer and M. Agrawala. Design considerations for collaborative visual analytics. *Proc. IEEE Symposium on Visual Analytics Science and Technology*, pages 171–178, 2007.
- [16] J. Heer, F. Viegas, and M. Wattenberg. Voyagers and voyeurs: Supporting asynchronous collaborative information visualization. *Proc. ACM SIGCHI Conference on Human Factors in Computing Systems*, pages 1029–1038, 2007.
- [17] N. Mahyar, A. Sarvghad, and M. Tory. A closer look at note taking in the co-located collaborative visual analytics process. *Proc. IEEE Symposium on Visual Analytics Science and Technology*, pages 171–178, 2010.
- [18] Z. Pousman and J. Stasko. Casual information visualization: Depictions of data in everyday life. *Proc. IEEE Symposium on Information Visualization*, pages 1145–1152, 2007.
- [19] A. Robinson. Collaborative synthesis of visual analytic results. *Proc. IEEE Symposium on Visual Analytics Science and Technology*, pages 67–74, 2008.
- [20] G. Salton and M. McGill. *Introduction to Modern Information Retrieval*. McGraw-Hill Press, 1986.
- [21] B. Shneiderman and A. Aris. Network visualization by semantic substrates. *IEEE Transactions on Visualization and Computer Graphics*, 12(5):733–740, 2006.
- [22] B. Shneiderman and C. Plaisant. Strategies for evaluating information visualization tools: multi-dimensional in-depth long-term case studies. *Proc. AVI Workshop on beyond time and errors: novel evaluation methods for information visualization*, pages 1–7, 2006.
- [23] Y. Shrinivasan, D. Gotz, and J. Lu. Connecting the dots in visual analysis. *Proc. IEEE Symposium on Visual Analytics Science and Technology*, pages 123–130, 2009.
- [24] Y. Shrinivasan and J. van Wijk. Supporting the analytical reasoning process in information visualization. *Proc. ACM SIGCHI Conference on Human Factors in Computing Systems*, pages 1237–1246, 2008.
- [25] Y. Shrinivasan and J. van Wijk. Supporting exploration awareness in information visualization. *IEEE Computer Graphics Applications*, 29(5):34–43, 2009.
- [26] F. Viegas, M. Wattenberg, F. van Ham, J. Kriss, and M. McKeon. Manyeyes: a site for visualization at internet scale. *IEEE Transactions on Visualization and Computer Graphics*, 13(6):1121–1128, 2007.
- [27] F. Wei, S. Liu, Y. Song, S. Pan, M. Zhou, W. Qian, L. Shi, L. Tan, and Q. Zhang. Tiara: a visual exploratory text analytic system. In *Proc. KDD*, pages 153–162, 2010.
- [28] W. Willett, J. Heer, J. Hellerstein, and M. Agrawala. Commentospace: Structured support for collaborative visual analysis. *Proc. ACM SIGCHI Conference on Human Factors in Computing Systems*, 2011, to appear.
- [29] W. Wright, D. Schroh, P. Proulx, A. Skaburskis, and B. Cort. The sandbox for analysis: Concepts and methods. *Proc. ACM SIGCHI Conference on Human Factors in Computing Systems*, pages 801–810, 2006.
- [30] D. Yang, Z. Xie, E. Rundensteiner, and M. Ward. Managing discoveries in the visual analytics process. *SIGKDD Explorations*, 9(2):22–29, 2007.