

Storygraph: Extracting patterns from spatio-temporal data

Ayush Shrestha
Department of Computer
Science
Georgia State University
Atlanta, Georgia
ashrestha2@cs.gsu.edu

Ben Miller
Department of Communication
Georgia State University
Atlanta, Georgia
miller@gsu.edu

Ying Zhu
Department of Computer
Science
Georgia State University
Atlanta, Georgia
yzhu@cs.gsu.edu

Yi Zhao
Department of Mathematics
Georgia State University
Atlanta, Georgia
yzhao6@gsu.edu

ABSTRACT

Analysis of spatio-temporal data often involves correlating different events in time and location to uncover relationships between them. It is also desirable to identify different patterns in the data. Visualizing time and space in the same chart is not trivial. Common methods includes plotting the latitude, longitude and time as three dimensions of a 3D chart. Drawbacks of these 3D charts include not being able to scale well due to cluttering, occlusion and difficulty to track time in case of clustered events. In this paper we present a novel 2D visualization technique called *Storygraph* which provides an integrated view of time and location to address these issues. We also present storylines based on *Storygraph* which show movement of the actors over time. Lastly, we present case studies to show the applications of *Storygraph*.

Categories and Subject Descriptors

I.3.3 [Computer Graphics]: Picture/Image Generation—*Line and curve generation*

Keywords

Spatio-temporal visualization, Information visualization

1. INTRODUCTION

With the advent of newer and cheaper location tracking devices and services, huge amount of spatio-temporal data is generated everyday. To visually analyze these kinds of data, presenting the time information and the location information separately is quite trivial. However, presenting spatio-temporal events together in an integrated view is not. For instance, it is not easy to present locations in a time series

chart, and it is not easy to present temporal information on a map. To address this issue, we present an interactive 2D visualization technique called *Storygraph* in this paper.

Storygraph consists of two parallel vertical axes similar to the vertical axes in the parallel coordinates, along with a horizontal axis. The vertical axes represent the location pair e.g. latitude and longitude while horizontal axis represent time. Each location (x, y) or $(latitude, longitude)$ is represented as a line segment referred to as *location line* in rest of the paper. Location lines are created by joining the corresponding values on the axes. Events occurring at that location at different times are plotted as markers along the location line. We assume that all events contain location (latitude and longitude) and time information in the data set.

In *Storygraph*, the location lines of two locations close to each other geographically are also close in the *Storygraph*. Thus it help users understand events in their spatial, temporal and in spatio-temporal context. This results in a unique and powerful ability of *Storygraph* to display recurring events or periodicity in data.

In many cases, data sets often contain attributes other than time and location. For such data sets, a common analysis technique is to highlight certain dimensions. For example, in a collection of human rights documents, besides the time and location of the violation, the name of perpetrators, victims, type, description is often present. In such data sets users might want to highlight a certain type of violence to better analyze it. To accomplish this, the *Storygraph* allows users to change the color, size and shape of the markers to highlight these dimensions.

Collection of events generally contain actors. Actors in this paper refer to groups of people, individuals or organizations present in the event. In a human rights violation event of "kidnapping of an individual by a guerrilla army and Red Cross helping the victims", Red Cross, guerrilla army and the individual could be three actors. Another common analysis is tracking the movement of characters. Using the same example as above, suppose the user wants to track the movement of the guerrilla army, and the army spends longer time in one place and lesser in others. Using traditional cartographic maps, users lose the sense of time. On the other hand, it is difficult to show location using timelines. To ad-

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

IDEA '13, August 11th, 2013, Chicago, IL, USA.

Copyright 2013 ACM 978-1-4503-2329-1 ...\$15.00.

dress this, we also present *storylines* based on Storygraph which shows the movement of actors in space and time.

We demonstrate the effectiveness of Storygraph using two case studies based on Afghanistan War Diary data and Vietnam War-era Laos bombing data. Storygraph reveals interesting patterns including gaps and trends in the events, thus providing insights into the data set.

The rest of the paper is organized as follows. In Section 2 we describe the related visualization techniques and recent work on the subject. In Section 3 we describe the nature of data that we used. Section 4 describes our main visualization technique: Storygraph. Following that we describe the implementation details in Section 5. Lastly we present the case studies in Section 6.

2. RELATED WORK

We are interested in visualizing spatio-temporal data sets and storylines based on it. Hence our analysis of related work focuses on the previous visualization methods for temporal and spatial data.

Much work has been done on time series data visualization some of which try to visualize storylines [2] [3] [10] [37] [20] [14] [13] [39] [23]. None, however, integrate time, location and actors. For example, visualization methods proposed in [10] [13] [23] can visualize time, actor, and sometimes context, but they do not include location.

Many methods have been proposed for spatio-temporal data visualization and analysis. Here we classify these methods based on whether it's 2D or 3D and whether it has integrated or separate spatial and temporal views. Thus we review the visualization methods in four categories: 2D separate views, 2D integrated view, 3D integrated view, and 3D separated views. Following that we review the existing work on storylines.

2.1 2D separated views

Many visualization methods present spatial and temporal data in different views. One key question is how to synchronize the presentation of temporal and spatial data. A number of methods have been proposed. For example, [8] [7] [31] used small multiple views to link spatial data with temporal data. Each small map view shows a subset of the temporal data. Plug et al. [31] use one map per weekday. Jern et al. [21] use color coding to link temporal data with spatial data. Fredrikson et al. [11] use a Gantt chart to link temporal data with spatial data. In this case, each small map view corresponds to a bar in the Gantt chart. Authors in [11] [28] use interactive synchronized views. When the user clicks on a data point in the temporal view, the corresponding point in the spatial view is automatically highlighted, and vice versa.

These methods have their limitations. The small multiple map views used by [8] [7] [31] cannot accommodate small temporal data scale. To avoid generating too many map views, temporal data is cut into large chunks. Color coding to link temporal data with spatial data is not intuitive, as do not have a standardized correlation, while time line is generally sequential. Interactive, synchronized spatio-temporal views are also limited, as they fail to present the continuous correlation between spatial and temporal data.

2.2 2D integrated views

In this category, the spatial and temporal information are

visualized in one view. For example, Andrienko et al. [5] propose a method that superimposes time graph on multiple maps. However the time graph may occlude the data points on the map view. In addition, events that happened at the same location but at different times may occlude each other. If many incidents happen on one location over a long period of time, many data points will be piled up on one spot, making it difficult to read and analyze.

2.3 3D integrated views

To address the problems of 2D integrated views, some researchers proposed 3D integrated visualization of spatio-temporal data. For example, in the space-time cube method [12], a synchronized 3D time line is plotted above a 2D map. Tominski et al. [35] place 3D icons on a 2D map to visualize events that happen on the same location over a period of time. In addition, time path is plotted above the 2D map. The benefits of 3D integrated view is that the time graph do not occlude the 2D map. If there are many data points at one location, there is a third dimension to accommodate the icons. However, the 3D integrated views also have its drawbacks. First, it's difficult to align time data with location in 3D. Techniques used by Kapler et al. [22] can be used to a certain extent to align the time data but as the number of events grows, the scalability reduces. These are the inherent problems of 3D data visualization.

2.4 3D separated views

Andrienko, et al. [1] use three separate but linked views: a 3D space-time cube, a 2D map view, and a time line view. A similar approach is used in Landesberger, et al. [24].

Compared with previous methods, our proposed method has the following benefits:

- It provides a big picture of the entire set of events in time as well as location with details displayed on demand.
- Spatial and temporal data are fully integrated in one view, and temporal and spatial data do not occlude each other.
- It is a 2D integrated visualization, thus avoiding the problems of 3D visualization.

2.5 Storylines

Starting with Minard's map of Napoleon's march to Moscow [36], there have been many efforts to create storylines [12] [4] [33] [9] [34]. Authors in [12] [4] [9] use 3D maps and plot the movement over it. Space-time paths are used to show the time spent by each of these characters in the location. To avoid over plotting when the number of actors increase, clustering is often used by the authors. Rinzivillo et al. [33] use 2D map and plot the trajectories of actors along streets. While it is ideal for actors moving with uniform velocity and provides highly accurate geographic information on the movement, it fails to show how much time did the actor spent at one point in case of non-uniform velocity of actors. Tanahashi et al. [34] create storylines of different characters in movies to show how they interact at what point in time. The resulting 2D diagram gives a good sense of time but only provide approximate location information based on the actors. Implementing storylines with Storygraph results in a visualization with high precision in both location and time making it more suited towards visualizing log data.

3. DATA CLASSIFICATION

Different types of visualizations work best on different types of data. In this paper we classify the data into three types: *Structured*, *Semi-structured* and *Unstructured*. Structured data sets contain precise geolocation and a time stamp for each event within the set. These data are uniform i.e., all the time stamps have equal precision (unlike some time stamps containing only year and another containing year, month and day) resulting in minimal or no ambiguity e.g. geo-sensor data, military operational logs etc. Most visualizations above use data from this category. Few common limitations include poor scalability though many authors have proposed workarounds by clustering. For example, Stardiates[25] works most appropriately on structured data[26] but fails to handle cases of high frequency data, Gravi++[18] works best for time dependent data series but has limitations in the frequency of the data as in case of Stardiates, method presented by Geng et al. [13] works well for large high-dimensional data but doesn't incorporate the temporal aspect well. Semi-structured data set include description of events with at least one of spatial or temporal component being in precise form. e.g. News report, Captain's Logs etc. In these data sets, the uniformity of the precision might not be guaranteed across the set, thus resulting in uncertainty. An example of this would be, data set containing three events with locations "Atlanta", "Atlanta, Georgia" and "10th Street, Atlanta Georgia". Similarly examples of non-uniformity in time include historic archives where some document have precise dates while others have only year, month-year present. Unstructured data include text reports where the extracted data may not have precise location and time, thus containing high amount of uncertainty. Examples include interview with the Fire fighters who were present in the scene during 9/11. In this paper, we only deal with structured data and leave the semi-structured and unstructured for future work.

4. METHOD

The Storygraph, is a 2D diagram consisting of two parallel vertical axes $V_\alpha \subset \mathfrak{R}$ and $V_\beta \subset \mathfrak{R}$ and an orthogonal horizontal axis $H \subset \mathfrak{R}$. Theoretically, all three of the axes, as in Cartesian graphs, are unbounded at both ends. In practice we take the maximum and minimum values in the data set for all corresponding axes and draw the graph accordingly. The values in the axes are ordered in ascending order: from left to right in horizontal axis and bottom to top in vertical axes. In this paper, vertical axes V_α and V_β represent the x and y coordinates of a point on a plane such as latitude and longitude. The horizontal axis, H , represents time. Thus a point plotted on Storygraph, which shall be referred to as *event* in the rest of the paper will have at least three dimensions: Parallel coordinates and a timestamp.

For any event occurring at (α, β) in time t and $t + 1$ as shown in Figure 1, our algorithm first draws a *location line* by connecting the points on the two axes, $\alpha \in V_\alpha$ and $\beta \in V_\beta$. The algorithm then returns the points on this line at time t and $t + 1$ respectively.

The function $f(\alpha, \beta, t) \rightarrow (x, y)$ which maps an event to the 2D Storygraph plane can be formally written as follows:

$$y = \frac{(\beta - \alpha)(x - T_{min})}{T_{max} - T_{min}} + \alpha \quad (1)$$




Figure 1: Left: Two events taking place at the location (α, β) at time t and $t + 1$ in the map, Right: Same two events represented in a Storygraph with parallel coordinates and timestamp.

$$x = t \quad (2)$$

where T_{min} and T_{max} are the maximum and minimum timestamps within the data set.

Figure 1 illustrates how a location on a regular 2D map, coded with a Cartesian coordinate, is presented in the Storygraph. Equation 1 and 2 is used to convert a location on a regular map to a Storygraph plane, and vice versa. As seen from the equations, such conversion is very efficient and can be done in real time allowing users more interactivity.

4.1 Cluttering and Ambiguity

While plotting a large number of events at different locations, often the location lines result in over plotting and cluttering. To address this issue, we set the alpha value of the location line as a function of some attribute selected by the user similar to the method described by Atero et al. [6].

Given some attribute a with maximum and minimum values a_{max} and a_{min} and a tuning parameter p , the color c of the location line is given by

$$c = \left(\frac{a - a_{min}}{a_{max} - a_{min}} * 255^{(1/p)} \right)^p \quad (3)$$

The tuning parameter p is used to control the difference between the higher values and the lower values. Figure 2 shows the change in the intensity of the location lines varying p . Higher the value of p , more evident is the difference between maximum and the minimum values of the attribute. Equation 3 is also used to color markers in the figure setting $p = 0.18$ in all four cases.

Our implementation of Storygraph also has a button to show/hide location lines. Hiding location lines is particularly useful when the data set is dense since the events align with each other over the hidden location lines giving a sense of the location. Hiding the location lines on the sparse data set may cause ambiguity since a point on the Storygraph may belong to multiple location lines. Another form of ambiguity occurs when two events at different locations get mapped to a single point on the Storygraph so that even though there are two location lines, its hard to distinguish which event belongs to which location. Our software alleviates this problem by providing zooming and filtering functions. Users can zoom into the Storygraph by clicking and dragging the mouse cursor over the Storygraph. The graph is redrawn expanding the selected area. Similarly, users can filter the range of latitude, longitude and time by specifying




Figure 2: The location lines drawn on the same set of storylines with different tuning parameters. From left to right, the first image has location lines painted with the same color. For the second, third and the fourth image, $p = 0.5$, $p = 1$ and $p = 2$ are used respectively. Higher value of p is useful for highlighting the attributes close to maximum. The same technique is used to paint the event markers as well in all the images with $p = 0.18$.

the numbers manually from a dialog box. Zooming and filtering minimizes the likelihood that two point overlap even after the Storygraph is redrawn.

4.2 Storyline generation

In this paper, we describe storyline as a series of events associated with an actor. Hence, in order to draw storylines, the data also needs to have actor information as one of its attributes. Storyline is constructed by connecting all the events sequentially in which the actor was involved resulting in a polyline. Multiple storylines can be drawn one for each actor thus allowing users to visualize the interactions between them. Storylines are especially helpful in visualizing movement data as in [9], [29], [30] because users can clearly see how different characters converge at certain events and then move on to different directions.

Since the event data is discrete, the movement of the actor between two time points cannot be determined. This uncertainty is presented in the Storylines by using furlough lines between two points as opposed to drawing solid lines.

Each storyline has a color value associated with it to distinguish one unit from the another. The color scheme in the implementation was adapted from Color Brewer [17]. This led to a problem when the number of actors increased beyond twelve - the upper cap of the brewer. To overcome this limitation, we hashed the unit names and converted those hashes into colors assuring each distinct unit to have a distinct color. In this process some colors differed from each other negligibly making two storylines almost indistinguishable by human eye. In this paper we only focus on storylines of less than twelve units and leave the color generation for higher numbers as future work.

Storylines based on Storygraph can be compared with the well established space-time path concept in Time-geography [16] because they share similar characteristics. Both techniques attempt to present temporal and spatial data in an integrated way. However, because Time-geography is a 3D visualization, it suffers from occlusion, cluttering, and difficult depth perception as more data points are plotted. On the other hand, Storygraph is a 2D visualization and therefore does not have any occlusion or depth problem. As can be seen from Figure 3, Storygraph can present co-location

in time or space and co-existence with lesser ambiguity than Time-geography.

4.3 Surroundings

When analyzing an event or a storyline, it is important to examine each event in its context or *surroundings*. In this paper, the surroundings refer to the events that have taken place in the temporal and spatial vicinity of the event. In Storygraph, surroundings can be defined in terms of space, time, actors or any of their combinations. For example, the surrounding of an event may be events that occurred in different neighboring locations. It could also be other events that an actor experienced in the past before that event. Analysis of surroundings in Storygraph can be compared to animated maps or videos of maps showing changes. An advantage of the Storygraph over the animated maps is the users do not need to track what happened where mentally frame by frame. Our implementation allows users to filter the view to get surroundings of a particular event. Users can adjust the coverage of the surroundings in terms of space, time as well as actors.

5. IMPLEMENTATION

We implemented the Storygraph using Microsoft WPF and C#. We used MySQL database in the back end to store the data set. For visualizing the data we used a layer based approach as case of most GIS applications - the events, location lines and storylines plotted in different layers. Thus the users can get additional information about the data set from different sources and overlay it on the existing Storygraph to further facilitate the analysis of the events. We implemented this feature using the Wikileaks Afghanistan war logs [15] and Cablegate data [38] as shown in Figure 5. The Afghanistan war logs served as the first layer on top of which the Cablegate material was overlaid. Since the Cablegate material only had the embassy information for all the cables, we plotted them as lines instead of points on the Storygraph. Each vertical line in the Cablegate overlay corresponds to a cable sent on that day.

6. CASE STUDIES

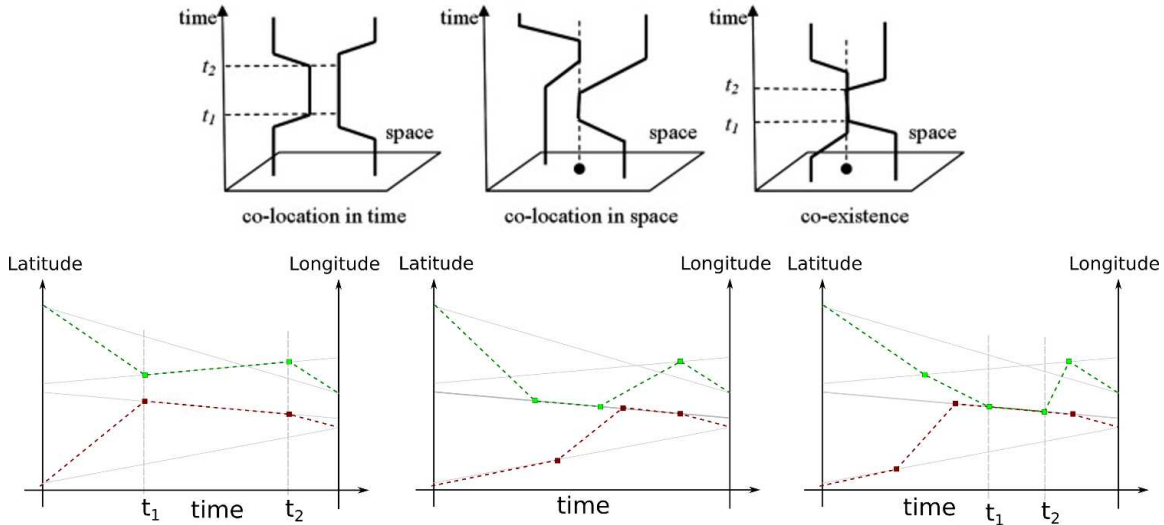


Figure 3: Comparison of the time-space paths using Time-Geography, adapted from [19] with storylines drawn using Storygraph. Starting from the left, the top figures directly compares to the bottom figure - first showing co-location in time, second showing co-location in space and third showing co-existence.

We applied our techniques on two databases: Wikileaks Afghanistan War Diary (War Diary) and data related to Laos from NARA’s Southeast Asia Data Base of Records About Air Sorties Flown in Southeast Asia (Laos UXO). Our visualizations, in some cases, revealed meaningful patterns in the data for which we were able to generate hypotheses.

6.1 Afghanistan War Log (2004-2010)

The War Diary comprises U.S. military significant activity reports from Afghanistan during the period 2004-2010 [15]. Consisting of approximately 60K highly structured records, the data provides a rigorously categorized and unprecedented look at the daily conduct of war. Figure 4 shows the Storygraph generated from all the events categorized as "Enemy Action" or "Explosive Hazard" during the war. We observed few interesting patterns marked by *A*, 1, 2 and 3.

The vertical bands marked by *A* in Figure 4 and *B* in Figure 5 are formed due to the events clustering between Jul-Oct 2005 and near Sept 2009. They indicate widespread and coordinated "Enemy Action" and "Explosive hazards" events at a very short period of time. By correlating these dates to broader events in Afghanistan, we discovered that these clusters were proximal to elections; our hypothesis is that the incidence of violence, and therefore the number of activity reports, increases during elections [38]. Although this is not new information, it demonstrates the ability of Storygraph to help identify significant event patterns.

Numbers 1 – 3 in the Figure 4 shows a periodicity of voids, or a lack of reports in those areas. These "holes" seem to appear around the end of the year. From location lines, we found that the geographic location by these holes correspond to a section of the Kabul-Kandahar highway, a key portion of national road system of Afghanistan and a main target of attacks. The Storygraph visualization shows that there have been regular quiet periods for that section of the highway around the end of 2005, 2006, 2007, and 2008. Since the data set does not extend beyond December 2009, we are unable to confirm if the pattern repeated in 2009. To our

knowledge, this pattern has not been identified or discussed in any published report. It is beyond the scope of this paper to evaluate the significance of this pattern and explain the cause of it. However, it demonstrates the effectiveness of the Storygraph to help discover patterns and form hypotheses.

Figure 6 shows the storylines of seven Afghanistan based combat units from October 1-20, 2009. It can be seen that units *tfprotector* and *tfcyclone* were at the same location during October 5 - 7 marked by *D* in Figure 6 and also during 12 - 14. Similarly *E1* and *E2* shows the presence of two units alternated in that location. It is also clear that *tfwhiteeagle* travelled a wider range of locations than the rest, *tflethal* travelled a lesser range and *tfprotector* didn't travel at all. Hence, besides showing the mobility of the units and how the unit interacted with other units, the greatest strength of the storylines based on Storygraph is it also shows the time period during which the unit was most mobile or least mobile.

6.2 Unexploded Ordnance Laos (1964-1973)

An ongoing problem in war torn regions is the persistence of unexploded ordnance. Established by the Lao government with support by various NGOs, the Lao National Unexploded Ordnance Programme (UXO Lao) addresses this ongoing problem so as to reduce the number of new casualties and increase the amount of land available for agriculture. By their estimates, approximately 80m unexploded bombs remain in Laos as a result of U.S. Air Force (USAF) bombing missions during the Vietnam Conflict. These bombs affect 25% of all villages, all 17 provinces, and have resulted in approximately 25k injuries and fatalities in the post-war period, 1974-2008. This data set details bombings in Laos by the USAF during the period of the war, 1968-1975 and is also a military, descriptive log of the bombings. It consists of 60K reports documenting approximately 2m tons of ordnance. Figure 7 shows the graph obtained from this data set.

Figure 7 shows all the events in the UXO data set. The




Figure 4: Storygraph showing all the events categorized as "Enemy Action" or "Explosive Hazard" during Afghanistan war from 2004 to 2009 across different regions of the country. Pattern marked by A shows that a lot of events took place during the same time period throughout the country. Patterns marked by 1 – 3 show voids - indicating that the attacks were periodic in nature.

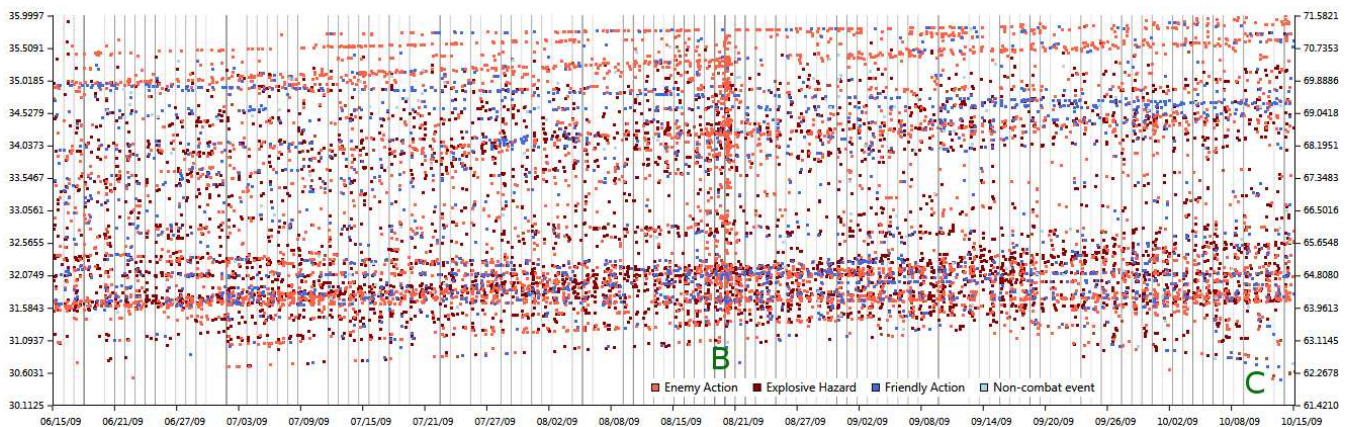


Figure 5: The Afghanistan war log from June 2009 to October 2009 with Cablegate cables within the same time frame plotted overlaid. The figure shows four types of events "Enemy Action", "Explosive Hazard", "Friendly Action" and "Non-combat event". Pattern B, similar to pattern A in Figure 4 show a number of events taking place in the entire region within the small time frame. Pattern C shows events taking place in near regular intervals (not induced due to the lack of precision in time).

markers are colored according to the number of bombs dropped. Patterns marked by A1 and A2 show that the bombings intensified during this period. A common error in interpreting Storygraph arises in cases like A2 where the users might be tempted to argue that more bombs were dropped in a certain location during that time. This kind of confusion resulting from crossing location lines can be alleviated in our implementation by zooming into the region. Figure 8 shows the zoomed view of region marked by A2 in Figure 7. The alpha value of the location line has been changed according to the number of bombings so that the higher number of bombings are highlighted. The figure shows two crossing bands implying that these locations were constantly bombed (as opposed to bombings concentrating at one location).

The Storygraph obtained from this data set shows three clear patterns.

1. Events clustering on location lines to form distinct lines can be seen from the start to June 1966. These

lines signify that the bombings were focused at certain locations at almost regular intervals.

2. The bombings reduced drastically after March 1972 when most US troops left Vietnam.
3. The patterns of bombing data are interspersed with periodic bands of white.

We have two hypothesis for those voids: either they could mean that bombings were paused during that range of time, it could also mean that the data correlated to those raids during that period was redacted from the set. Like much military data, classified operations such as those by special forces are frequently not contained within general operation activity reports. One way to test this second hypothesis would be to correlate bombing data to peace talks, truces, and ceasefires as documented in other sources. Identifying the focal locations of bombing campaigns helps groups like UXO Lao address the areas most in need of re-mediation.

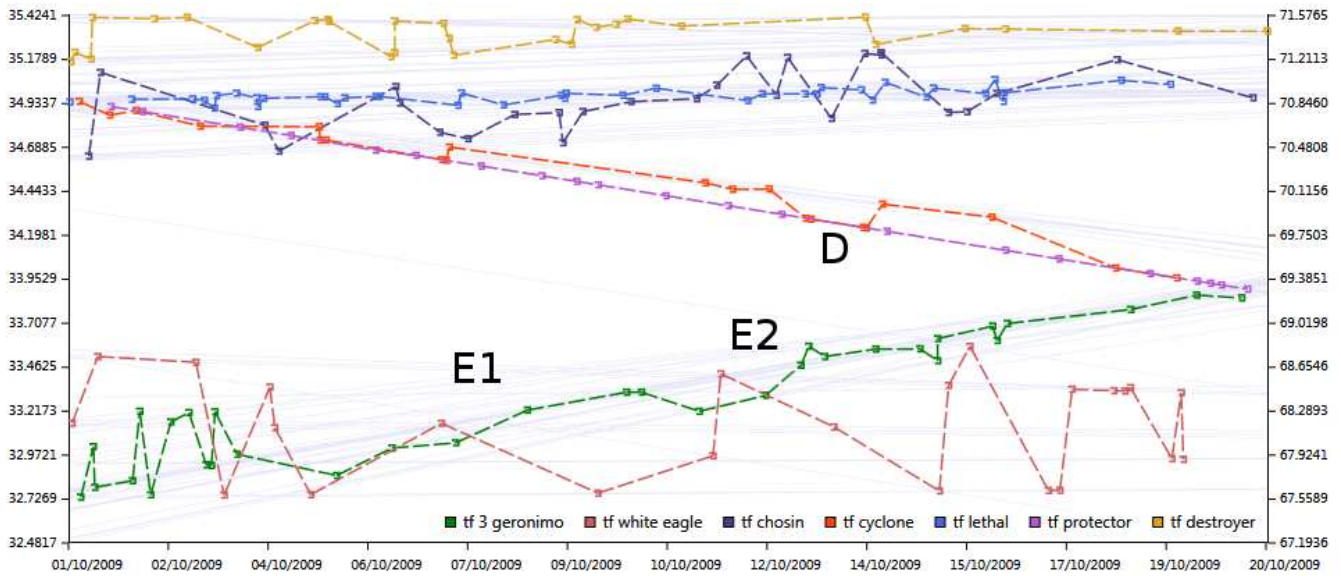


Figure 6: Storylines of seven Afghanistan based combat units from October 1 - 20, 2009. The pattern marked by *D* shows co-location in space and time - meaning that the two units were at the same location. Pattern marked by *E1* and *E2* show that co-location in time - meaning that one unit was present in the location before another.

7. CONCLUSION

In this paper, we have presented our novel visualization technique, Storygraph, which provides an integrated view containing locations and time. We also presented storylines based on Storygraph. Storygraph addresses the problems in previous 3D and integrated 2D spatio-temporal data visualization by minimizing glyph occlusion and cluttering. This improved design help users better correlate events on both the spatial and temporal dimensions. Storylines help track the movement of characters and the interactions between them. In the future we plan to extend the visualization to incorporate uncertainty in location and time. We also intend to implement data clustering algorithms to handle large scale data sets.

8. ACKNOWLEDGMENT

This material is based upon work supported by the National Science Foundation under Grant No. 1209172. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

9. REFERENCES

- [1] G. Adrienko, N. Adrienko, M. Mladenov, M. Mock, and C. Politz. Identifying place histories from activity traces with an eye to parameter impact. *IEEE Transactions on Visualization and Computer Graphics*, 18(5):675–688, May 2012.
- [2] W. Aigner, S. Miksch, W. Muller, H. Schumann, and C. Tominski. Visualizing time-oriented data - a systematic view. *Computers and Graphics*, 31(3):401 – 409, 2007.
- [3] W. Aigner, S. Miksch, W. Muller, H. Schumann, and C. Tominski. Visual methods for analyzing time-oriented data. *IEEE Transactions on Visualization and Computer Graphics*, 14(1):47 –60, jan.-feb. 2008.
- [4] G. Andrienko, N. Andrienko, and S. Wrobel. Visual analytics tools for analysis of movement data. *ACM SIGKDD Explorations Newsletter*, 9(2):38–46, 2007.
- [5] N. Andrienko, G. Andrienko, and P. Gatalsky. Visualization of spatio-temporal information in the internet. In *Proceedings of 11th International Workshop on Database and Expert Systems Applications*, pages 577 –585, 2000.
- [6] A. O. Artero, M. C. F. de Oliveira, and H. Levkowitz. Uncovering clusters in crowded parallel coordinates visualizations. In *Information Visualization, 2004. INFOVIS 2004. IEEE Symposium on*, pages 81–88. IEEE, 2004.
- [7] A. Asgary, A. Ghaffari, and J. Levy. Spatial and temporal analyses of structural fire incidents and their causes: A case of Toronto, Canada. *Fire Safety Journal*, 45(1):44 – 57, 2010.
- [8] C. Brunsdon, J. Corcoran, and G. Higgs. Visualising space and time in crime patterns: A comparison of methods. *Computers, Environment and Urban Systems*, 31(1):52 – 75, 2007.
- [9] U. Demsar and K. Virrantaus. Space-time density of trajectories: exploring spatio-temporal patterns in movement data. *International Journal of Geographical Information Science*, 24(10):1527–1542, 2010.
- [10] D. Fisher, A. Hoff, G. Robertson, and M. Hurst. Narratives: A visualization to track narrative events as they develop. In *Proceedings of IEEE Symposium on Visual Analytics Science and Technology*, pages 115 –122, oct. 2008.
- [11] A. Fredrikson, C. North, C. Plaisant, and B. Shneiderman. Temporal, geographical and

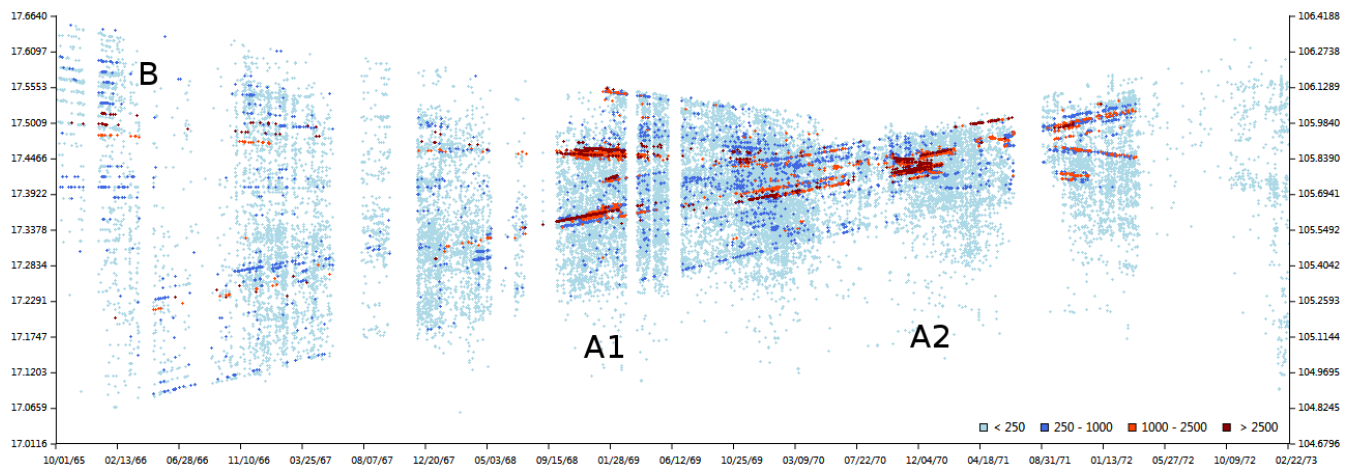


Figure 7: Storygraph generated from Laos data set. Each event corresponds to an instance of bombing. Patterns marked by A1 and A2 showing bombing intensified during the period. Pattern B shows that some places were bombed regularly. Also evident in the figure are the vertical bands of white either meaning missing data or the stopped bombings.

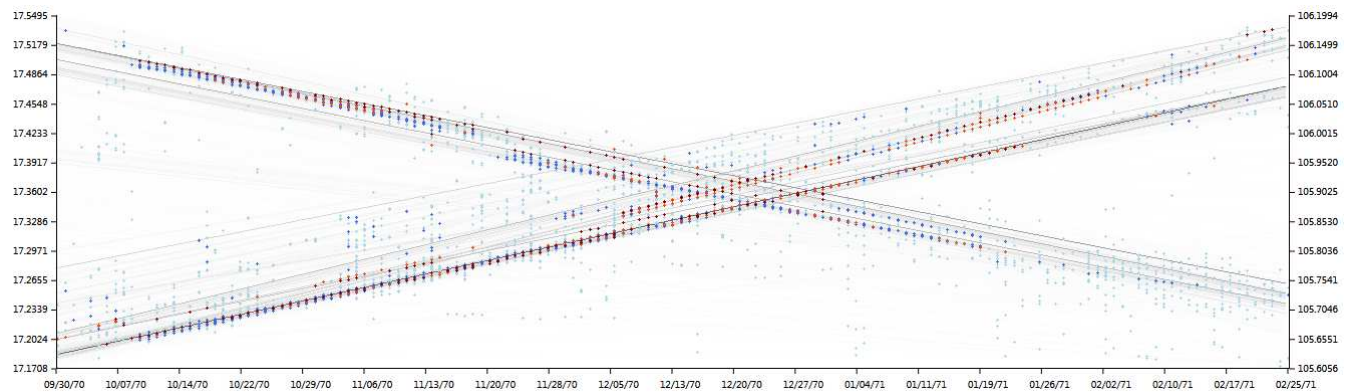


Figure 8: Area marked by A2 in Figure 7 zoomed together with location lines. The alpha value of the lines were set according to the number of bombs dropped so that more prominent areas and bands of location are highlighted.

categorical aggregations viewed through coordinated displays: a case study with highway incident data. In *Proceedings of the 1999 workshop on new paradigms in information visualization and manipulation in conjunction with the eighth ACM international conference on Information and knowledge management*, NPIVM '99, pages 26–34, New York, NY, USA, 1999. ACM.

[12] P. Gatalsky, N. Andrienko, and G. Andrienko. Interactive analysis of event data using space-time cube. In *Proceedings. Eighth International Conference on Information Visualisation*, pages 145 – 152, July 2004.

[13] Z. Geng, Z. Peng, R. Laramee, R. Walker, and J. Roberts. Angular histograms: Frequency-based visualizations for large, high dimensional data. *IEEE Transactions on Visualization and Computer Graphics*, 17(12):2572 –2580, Dec. 2011.

[14] T. Gschwandtner, W. Aigner, K. Kaiser, S. Miksch, and A. Seyfang. CareCruiser: Exploring and visualizing plans, events, and effects interactively. In *Proceedings of IEEE Pacific Visualization Symposium (PacificVis)*, pages 43 –50, March 2011.

[15] Guardian.co.uk. Afghanistan war logs 2003-2010, July 2010.

[16] T. Hägerstrand et al. Time-geography: focus on the corporeality of man, society, and environment. *The science and praxis of complexity*, pages 193–216, 1985.

[17] M. Harrower and C. Brewer. Colorbrewer.org: an online tool for selecting colour schemes for maps. *Cartographic Journal*, The, 40(1):27–37, 2003.

[18] K. Hinum, S. Miksch, W. Aigner, S. Ohmann, C. Popow, M. Pohl, and M. Rester. Gravi++: Interactive information visualization of highly structured temporal data. *Workshop IDAMAP 2005 Intelligent Data Analysis in Medicine and Pharmacology*, pages 67 – 72, 2005.

[19] T.-H. Hsieh, J.-J. J. Chen, L.-H. Chen, P.-T. Chiang, and H.-Y. Lee. Time-course gait analysis of

- hemiparkinsonian rats following 6-hydroxydopamine lesion. *Behavioural Brain Research*, 222(1):1 – 9, 2011.
- [20] W. Javed, B. McDonnel, and N. Elmqvist. Graphical perception of multiple time series. *IEEE Transactions on Visualization and Computer Graphics*, 16(6):927 – 934, nov.-dec. 2010.
- [21] M. Jern and J. Franzen. GeoAnalytics - exploring spatio-temporal and multivariate data. In *Proceedings of Tenth International Conference on Information Visualization*, pages 25 – 31, july 2006.
- [22] T. Kapler and W. Wright. Geotime information visualization. *Information Visualization*, 4(2):136–146, 2005.
- [23] M. Krstajic, E. Bertini, and D. Keim. CloudLines: Compact display of event episodes in multiple time-series. *IEEE Transactions on Visualization and Computer Graphics*, 17(12):2432 – 2439, dec. 2011.
- [24] T. V. Landesberger, S. Bremm, N. Andrienko, G. Adrienko, and M. Tekusova. Visual analytics for categoric spatio-temporal data. In *Proceedings of IEEE Conference on Visual Analytics Science and Technology (VAST 2012)*, pages 183–192, October 2012.
- [25] M. Lanzenberger, S. Miksch, and M. Pohl. The Stardines - visualizing highly structured data. In *Proceedings of Seventh International Conference on Information Visualization (IV 2003)*, pages 47 – 52, july 2003.
- [26] M. Lanzenberger, S. Miksch, and M. Pohl. Exploring highly structured data: a comparative study of stardines and parallel coordinates. In *Proceedings. Ninth International Conference on Information Visualisation*, pages 312 – 320, july 2005.
- [27] U. LAO. Lao national unexploded ordnance programme, 2011.
- [28] 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. *IEEE Transactions on Visualization and Computer Graphics*, 16(2):205 – 220, march-april 2010.
- [29] A. Murray, Y. Liu, S. Rey, and L. Anselin. Exploring movement object patterns. *The Annals of Regional Science*, pages 1–14, 2012. 10.1007/s00168-011-0459-z.
- [30] D. Orellana, A. K. Bregt, A. Ligtenberg, and M. Wachowicz. Exploring visitor movement patterns in natural recreational areas. *Tourism Management*, 33(3):672 – 682, 2012.
- [31] C. Plug, J. C. Xia, and C. Caulfield. Spatial and temporal visualisation techniques for crash analysis. *Accident Analysis and Prevention*, 43(6):1937 – 1946, 2011.
- [32] L. N. U. O. Programme. Annual report 2007, 2008.
- [33] S. Rinzivillo, D. Pedreschi, M. Nanni, F. Giannotti, N. Andrienko, and G. Andrienko. Visually driven analysis of movement data by progressive clustering. *Information Visualization*, 7(3-4):225–239, 2008.
- [34] Y. Tanahashi and K.-L. Ma. Design considerations for optimizing storyline visualizations. *Visualization and Computer Graphics, IEEE Transactions on*, 18(12):2679–2688, 2012.
- [35] C. Tominski, P. Schulze-Wollgast, and H. Schumann. 3D information visualization for time dependent data on maps. In *Proceedings. Ninth International Conference on Information Visualisation*, pages 175 – 181, july 2005.
- [36] E. R. Tufte. *Beautiful evidence*, volume 23. Graphics Press Cheshire, CT, 2006.
- [37] K. Vrotsou, J. Johansson, and M. Cooper. ActiviTree: Interactive visual exploration of sequences in event-based data using graph similarity. *IEEE Transactions on Visualization and Computer Graphics*, 15(6):945 – 952, nov.-dec. 2009.
- [38] Wikileaks. Cablegate’s cables, 2010.
- [39] J. Zhao, F. Chevalier, E. Pietriga, and R. Balakrishnan. Exploratory analysis of time-series with ChronoLenses. *IEEE Transactions on Visualization and Computer Graphics*, 17(12):2422 – 2431, dec. 2011.