

New Methods for Representing and Interacting with Qualitative Geographic Information

Contract #: W9132V-11-P-0010

Contract Period: May 23, 2011 – October 31, 2012

Principal Investigators:

Dr. Alan M. MacEachren, GeoVISTA Center, Penn State University

Dr. Anthony Robinson, GeoVISTA Center, Penn State University

Report on Component 3: SensePlace2 Evaluation

*Anthony C. Robinson, Scott Pezanowski, Alexander Savelyev,
and Alan M. MacEachren*

<arobinson, spezanowski, savelyev, maceachren>@psu.edu

GeoVISTA Center, Department of Geography, The Pennsylvania State University
Submitted, Oct. 31, 2012

REPORT DOCUMENTATION PAGE				Form Approved OMB No. 0704-0188	
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing this collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number. PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.					
1. REPORT DATE (DD-MM-YYYY) 05-11-2012		2. REPORT TYPE Final		3. DATES COVERED (From - To) May 23, 2011 – Oct. 31, 2012	
4. TITLE AND SUBTITLE Report on Component 3: Evaluating New Methods for Visualizing Text in Spatial Contexts				5a. CONTRACT NUMBER: W9132V-11-P-0010	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S) Alexander Savelyev, Scott Pezanowski, Anthony C. Robinson, and Alan M. MacEachren				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) PENNSYLVANIA STATE UNIVERSITY , THE 408 OLD MAIN UNIVERSITY PARK PA 16802-1505				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES) US Army Engineer Research and Development Center (ERDC) Topographic Engineering Center (TEC) 7701 Telegraph Road Alexandria, VA 22135-3864				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION / AVAILABILITY STATEMENT Approved for public release; Distribution is unlimited					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT Our aim for this component of research was to evaluate the prototype SensePlace2 environment to gauge its support for key tasks related to spatio-temporal analysis of qualitative data derived from social media sources (with the focus on Twitter). Our preliminary findings reveal that participants view SensePlace2 as having the capability to integrate and analyze geospatial dimensions of social media, but that the execution of the interface has many limitations related to ease of use and support for efficient analysis. Participants proposed many additional features they felt would improve utility and a range of interface adjustments they felt would improve usability. Qualitative feedback from our tasks shows that users were able to generate good answers to our task prompts in most instances. However, users frequently mention that their answers were difficult to generate and that they were uncertain about the quality of those answers. This further supports the overall finding that key mechanisms may exist to support solid analysis, but that the means for interacting with these mechanisms require significant further refinement. Results from this evaluation will be applied to specific interface improvements and the further development of refined analytical methods.					
15. SUBJECT TERMS geovisualization, visual analytics, social media, microblogs, cartography, qualitative geographic information, text analytics					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT	b. ABSTRACT	c. THIS PAGE			19b. TELEPHONE NUMBER (include area code)
SAR	SAR	SAR	SAR	35	

1 Introduction and Methodology

Our aim for this component of research was to evaluate the prototype SensePlace2 environment to gauge its support for key tasks related to spatio-temporal analysis of qualitative data derived from social media sources (with the focus on Twitter). Results from this evaluation will then in turn lead to specific interface improvements and the further development of refined analytical methods.

To satisfy these evaluation goals we developed a multi-part user study featuring task analysis and survey components to elicit qualitative and quantitative feedback on a range of related areas of concern. Eight participants (3 female, 5 male, all between the ages of 20-29) were recruited for our study from a graduate seminar course focusing on geographical analysis of social media. All participants are currently pursuing a graduate degree (6 in Geography, 1 in Criminal Justice, 1 in Information Science and Technology). We asked participants to rate their expertise in several broad areas, and they indicated their expertise was primarily in Geographic Information Systems, Information Science, and the Social Sciences (Figure 1).




Figure 1: Participants self-evaluated areas of expertise.

The study procedure includes three key parts. First, participants were given a tutorial document providing an overview of the key functions of SensePlace2, along with sample tasks to complete (Appendix A). Second, participants completed three representative tasks using SensePlace2 (detailed below). Finally, after completing these tasks, users completed a usability and utility survey to rate SensePlace2 against a wide range of metrics (Appendix B).

A brief introduction to SensePlace 2 was given during a class session (by MacEachren) in which the application was demonstrated briefly and students were invited to participate in the user study. The study procedure took place as a self-paced, distributed activity. Participants were given instructions on how to access the tutorial and survey website (surveymonkey.com) and were instructed to complete the activities at the time of their choosing within a 2-week period.

In the remainder of this report we highlight the preliminary findings from our study. Task analysis feedback and quantitative results from the usability and utility survey are outlined in separate sections below.

2 Results

The following sections describe and summarize preliminary results from task analysis and usability/utility survey questions that our participants completed.

2.1 Task Analysis

The first portion of our user study asked participants to complete three tasks and provide qualitative feedback on their analytical findings and experience. These three tasks included basic search and exploration, comparing tweets with different types of location information, and identifying/fixing geocoding errors. The following sections briefly describe each task and the qualitative results we gathered from our participants.

2.1.1 Basic Search / Exploration Task

Task 1 required participants to explore geospatial social media from a recent event (Figure 1):

To get started, type "whooping cough" into the "Search for" box at the top left of the SensePlace2 interface. Press the enter key to search for Tweets about whooping cough.

After SensePlace2 has finished loading the results for your search, take a few moments to explore the geographic locations mentioned in Tweets about whooping cough. Remember, the heat map shows the overall prevalence of location mentions in every Tweet about whooping cough, while the purple proportional circles show the locations (and number of Tweets talking about each location) for the top 1000 most relevant Tweets about whooping cough.

Next, use the timeline to identify and explore which time ranges include the most activity for Tweets about whooping cough. You should see peaks in activity at certain times. Narrow the time range to explore each peak in activity.

We asked participants to respond to four questions in this task. Each question and two representative answers are highlighted below:

Q1: What geographic patterns do you see?

A1: "Tweets about whooping cough seem to be appearing predominantly in two places: the continental United States and countries bordering the English Channel. There are two spots in Australia and one in New Zealand that are also registering ""Whooping Cough"". This suggests, to me, some possible linkage from whooping cough to countries and people of Anglo descent.

When examining locations, tweets from locations in America tend to link to other locations in America. This pattern extends to other countries as well. One exception to this is a person geolocated in Nice, France who is talking about how their mother caught whooping cough when they visited Melbourne, Australia."

A2: "The geographic pattern changes by time, but overall, the prevalence occurs in the United States and England (UK). Some cases in Canada."

Q2: What temporal patterns do you see?

A1: The peaks of whooping cough appear at around end of August in 2012. I used 1 month and 3 weeks fixed range to explore.

A2: "Binned by week: From the timeline it's very evident that the majority of tweets on the subject happened early on, in roughly the second week of May. This was part of an initial burst of development in the month of May, followed by a lull through most of the summer, and then a resurgence in later August and September.




Figure 1: Example screenshot showing results for "Whooping Cough" in SensePlace2.

Binned by Day: When binned by day you can see that the majority of discussion about whooping cough occurred between 09 May and 15 May. A small spike occurred again between 27 May and 30 May. The summer had some passing mentions. Then whooping cough became a semi-regular discussion from 20 Aug onward."

Q3: What did you learn from the content of the Tweets themselves?

A1: "From the tweet list in default sorting (most to least relevant), I learned a lot of whooping cough cases report in the U.S. and UK, and some announcements of vaccine by health agents.

Somehow I'm under the impression that the format of tweet text is not right. It seems a lot of space between words is missing. It's also possible those are typos from the tweet authors."

A2: "I learned how the government coped with whooping cough in the way of recommending vaccines to infants and adults. I could guess the peaks from tweets mentioning "the highest level in U.S." or "HIT 10-year High" and also specific regions where whooping cough broke out."

Q4: Provide at least two questions that you would ask another analyst to explore after seeing these patterns.

A1: *"Why are there more outbreaks occurring in the fall? What factors are influencing the outbreak in the northeast?"*

A2: *"Is there any way we can filter out tweets that link to external sources on whooping cough, segmenting just the content that gives personal information on the subject?"*

Does the whooping cough have any association with a particular demographic, specifically those of Anglo heritage?"

2.1.2 Comparing Tweets About Places to Tweets From Places

Task 2 required participants to compare the places mentioned in Tweets to the places that Tweets were reported from:

To begin this task, first click the "Reset Query" button to reset the SensePlace2 application, and then type "earthquake" into the "Search for" box at the top left of the SensePlace2 interface. Press the enter key to retrieve Tweets about earthquakes.

Once SensePlace2 has loaded these Tweets, take a few moments to explore the patterns of locations that are mentioned. Next, using the timeline, narrow the dataset to show only those Tweets from August 2012 until the present time.

Take some time to explore the data from just this time range and see if you can identify key events that received the most attention across these dates. Once you've done this, click the checkbox under the "Search for" box to "Retrieve only tweets with "from" places." This will have SensePlace2 refine your search by one more step and only show those tweets that included a reporting location (e.g. assigned by a phone or other means to indicate where somebody was when they tweeted).

Use the "Current task" dropdown list to switch back and forth between each of these two queries and note similarities/differences in what you see in terms of Tweet content and the patterns of relevant locations.

We asked participants to respond to four questions in this task. Each question and two representative answers are highlighted below:

Q1: What geographic patterns do you see?

A1: *"From" tweets are more concentrated in Japan, Alaska and California, where earthquakes were happening, while "about" tweets are more distributed."*

A2: *"As would be expected, there are a lot earthquake reports in the ring of fire region. This is especially true in Southern California, Indonesia, Japan, and the Philippines.*

There is considerably more international linkage of tweets in this query than in the previous query on whooping cough."

Q2: What temporal patterns do you see?

A1: "Given that we are investigating earthquakes, I was trying to see whether FROM PLACE tweets happened before other tweets regarding a specific earthquake or not. But I had a hard time finding an appropriate one, as many of the tweets with FROM location are tweets from NEWS MEDIA or SEISMIC monitoring centers, and not from ordinary users. It seems that the time range was not long enough to find an interesting case."

A2: "The peaks in timeline are roughly the same for both "about" tweets and "from" tweets. In both timeline, there are two peaks which are symbolized by the dark orange blocks in heat map and two or three stack bars. The first occurred between Aug.22–Sept.3, 2012 while the second peak occurred between Sept.19–Oct.2, 2012.

Note: Because of the slow response from the server, it took about 15-20 seconds to switch between "about" and "from" tweet views, which made the comparison much more difficult. I did a screenshot of the "about" tweet timeline and then switch the timeline in UI to "from" tweets to make the comparison."

Q3: What did you learn from the content of the Tweets themselves?

A1: "About tweets are typically very general, "WTF happened" kind of statements or exaggerations, emotional, little geographical detail, and from individual twitterers. From tweets are very specific providing only details, little sentiment or emotion, and are often from "official" sources like news agencies or "earthquake watch" kinds of accounts."

A2: "Many people are simply reporting the occurrence of an earthquake and its magnitude. Some refer to specific lat/lon coordinates, but many simply refer to a city name."

Q4: Provide at least two questions that you would ask another analyst to explore after seeing these patterns.

A1: "It would be interesting to overlay fault lines on the map. How does information temporally and spatially spread from the earthquake epicenter in digital space (how are ideas spreading)? How does population affect the findings of this distribution?"

A2: "Are there any emergency management tweets within this dataset, and how do we filter those out? Are emergency management agencies in Southeast Asian countries using Twitter, or putting out advisories?"

2.1.3 Fixing Geocoding Errors

Task 3 asked participants to explore Tweets from a recent event to identify and suggest corrections for geocoding errors:

To get started with the third and final task for this evaluation, first click the "Reset Query" button to reset the SensePlace2 application, then type "fire" into the "Search for" box at the top left of the SensePlace2 interface. Press the enter key to retrieve Tweets about fire.

Once SensePlace2 has loaded these Tweets, your task is to identify tweets that have locations that exhibit one or more of the following problems:

1. The locations in a Tweet were not correctly highlighted (e.g. they include a placename that does not show a blue background in the Tweet list)
2. The locations in a Tweet are misread by our system (e.g. SP2 highlighted Lafayette but did not include the word "Lake" before it, when the original Tweet says "Lake Lafayette")
3. The locations in a Tweet are misplaced by our system (Using the map, you determine that the circle referring to the place mentioned in a Tweet is in the wrong place)
4. The locations in a Tweet are mistaken by our system (SP2 has highlighted a placename that you know is not in fact a real placename)

You should identify at least one example of each problem from the results shown for the search for Tweets about fire. For each problem you discover, hover over the Tweet in the Tweet list and click the "Geocoding errors" link to launch our geocoding correction interface. Choose the right type of error you wish to report for each example and submit your report when you are ready.

Please identify and suggest fixes for at least 10 geocoding errors you discover in the Tweets. When you are finished doing this, continue through the rest of this survey to record your feedback.

We asked participants to respond to three questions in this task. Each question and representative answers are highlighted below for the first two open-ended questions. The third question asked for a multiple-choice response and we provide an overview of those responses (Figure 2):

Q1: SensePlace2 allows you to make corrections for a range of geocoding errors. Are there other error types that should be fixable that are not currently supported?

A1: "I don't know how it is treated in the background, but the "misplaced" one is too general. The difference between Washington State and Washington DC is not the same as two completely unrelated locations (say in alphabets say in location) and they are yet categorized together, that makes a systematized approach a bit difficult."

*Only 6 of 8 participants responded to this question and four of the answers said that there were no other types that should be supported.

Q2: What functionality would you add (or take away) from the SensePlace2 interface for handling geocoding errors in Tweets?

A1: "Linking to the tweet on the Twitter main page. This would allow the user to try and derive a location."

A2: "I think it would be nice to be able to fix errors in a batch by excluding tweets that meet certain criteria. For example, when you explore the tweets about an earthquake in Kobe, Japan. You might want to exclude tweets that contain keyword "Kobe" from Staples Center, Los Angeles (Michael Jordan has a similar problem :-)."

Q3: In your opinion, what is an acceptable proportion of results having location accuracy or precision problems when working with social media in a tool like SensePlace2?




Figure 2: Participants self-evaluated tolerance for location accuracy or precision problems when working with geospatial social media in a tool like SensePlace2.

2.2 Usability / Utility Survey Metrics

Users rated their experience with SensePlace2 along common usability metrics as well as specific utility metrics that we developed to assess SensePlace2’s capabilities to support space-time analysis, situational awareness, analytical reasoning, and geocoding error remediation.

The highest ratings from this portion of the survey were for SensePlace2’s overall integration, which participants generally agreed was well-conceived. The lowest ratings concerned SensePlace2’s ease of use and the likelihood that most people would be able to learn how to use SensePlace2 quickly. In terms of its basic usability, our participants generally gave average to below-average support when asked to rate SensePlace2 along a range of common usability metrics concerning appeal, learnability, simplicity, intuitiveness, and ease of use (Q1 – Q10 in Figure 3).

In terms of its utility for supporting space-time analysis (Q1-Q3 in Figure 4), participants agreed that SensePlace2 is capable of revealing spatial, temporal, and topical aspects of social media information. Strongest support however was shown for its spatial capabilities, with slightly weaker support for temporal analysis, and the lowest rating given to its ability to reveal topic information.

SensePlace2’s support for situational awareness (Q4-Q6 in Figure 4) was rated positively (above the mid-point) when it comes to perceiving key components of and understanding relationships between space, time, and attribute information. Support for the third component of situational awareness, which concerns prediction, garnered weak agreement from participants (Q6).

Participants did not reach consensus on whether or not SensePlace2 would be helpful for generating a report during a crisis situation or if it would help someone tell a compelling story about a crisis situation. Ratings for both questions (Q7-Q8 in Figure 4) yielded average scores around 3.0 which neither agrees nor disagrees with the statement that SensePlace2 supports either design objective.

Finally, participant ratings yielded an average score of 3.0 in terms of its support for easily identifying geocoding errors (Q9 in Figure 4). This score signals no consensus in support for the assertion that

SensePlace2 can help easily identify those errors. In contrast, once an error has been discovered, participants generally agreed that SensePlace2 allows one to easily suggest a change (Q10 in Figure 4).

Usability Metrics




Figure 3: Usability survey results.

Utility Metrics




Figure 4: Utility survey results.

3 Conclusions

Our preliminary findings reveal that participants view SensePlace2 as having the capability to integrate and analyze geospatial dimensions of social media, but that the execution of the interface has many limitations related to ease of use and support for efficient analysis.

Qualitative feedback from our tasks shows that users were able to generate good answers to our task prompts in most instances. However, users frequently mention that their answers were difficult to generate and that they were uncertain about the quality of those answers. This further supports the overall finding that the key mechanisms may exist to support solid analysis, but that the means for interacting with these mechanisms require significant further refinement.

A preliminary review of qualitative feedback to identify major bugs and ideas for new features provides us with a clear set of goals for further refinement going forward.

Summary of Identified Bugs:

- The Timeline tool has to be reset to show the full available time range after it has been narrowed once. Users found this confusing and some were unable to reset it.
- Some Tweets appeared to have incorrect text formatting.
- When scrolling the Tweet list, the “promote” and “demote” bar often does not follow the mouse cursor.
- System performance issues make it very hard to filter/search productively and efficiently
- The timeline cannot be used to browse the dataset since it begins a new query with every interaction – most users try to narrow it by both sides and then want to move that window around the available time frame. A nested approach is needed.
- If the phenomena of interest are happening in the Pacific Ocean (e.g. earthquakes), it’s impossible to zoom and see that Ocean area all at once.
- The Geocoding interface is too complicated, un-styled, and the error categories are not easily understood.
- The Geocoding “does it have a GeoName” field does not allow the user to input any text.
- In general the interface lacks “polish” and sensible interactions between/across views are not supported (selection using bounding boxes in the map or tag cloud, for example).

Summary of New Feature Requests:

- Need a method for representing and filtering retweets so that individual reports of interest can be easily detected and evaluated – many queries end up showing lots of nearly identical reports as “most relevant,” making it hard to find the real conversations of interest.
- Given one or more keywords, help the user by suggesting which other entities constitute relevant subsets. E.g. if a user searches for “Whooping Cough,” display suggestions for other keywords that might be interesting.
- Make it possible to identify places where “About” and “From” placenames are co-located. This could help the user identify up front which places are trending.
- The temporal analysis controls should allow time units small than 1 day, to include hourly or smaller increment analysis. Additional notions of time should be supported to be able to

summarize patterns by the day of the week they occur, by season, and other points of reference.

- It should be possible to view both “About” and “From” tweets in the same search, and toggle both layers on and off to compare their patterns more directly.
- It should be possible to integrate related spatial datasets, for example, to show fault lines on the basemap when working with an earthquake scenario.
- Enable filtering of the tweet list to quickly identify official sources of information apart from unofficial general Twitter users (e.g. news orgs and gov’t accounts vs. regular citizens).
- Make it possible to fix geocoding errors in a batch (e.g. every instance of Kobe mentioned with Los Angeles).
- Link each tweet directly to the original Twitter page so that more information about a particular user can be easily retrieved.
- Support animation in the interface to explore the dynamics of “About” and “From” tweets.
- Provide a dedicated view to help explore and analyze Tweet contents. Content analysis was commonly viewed as a weak point in the current interface.
- Clicking on a Tweet should re-center the map as well as apply a salient highlighting method so that the relevant places are easy to see.
- Predictive tools to help identify emerging places and topics of interest are needed to fully satisfy situational awareness objectives.

Following the next round of SensePlace2 development revision to implement as many of these changes as is feasible, we plan to re-evaluate the system to determine if we are able to improve on our previous efforts.

Appendix A

SensePlace2 Interface Mini-Guide

Alan M. MacEachren, Alexander Savelyev, and Scott Pezanowski
GeoVISTA Center, Pennsylvania State University

SensePlace2 is one of the flagship research projects at GeoVISTA Center and is currently under active development. We expand SensePlace2 functionality quite often and are in the process of performing user studies to obtain feedback on the user interface (UI) features. This mini-guide outlines the key capabilities of the current version of SensePlace2 and provides a short tutorial on their proper use. SensePlace2 also has a built-in legend, accessible through a link at the bottom right corner of the project's web interface, which provides an abridged summary of the information found in this guide.

1. Access and Performance

The current version of SensePlace2 can be accessed using the following web address:

<http://www.geovista.psu.edu/SensePlace2/app/>


When prompted, enter the username and password you have been provided with.

System performance is only partially optimized, and is strongly dependent on the number of matches a given query has in the database. Thus, queries that use popular search terms (e.g. "fire" or "protest") will take longer than queries based on more esoteric keywords. Currently, some of the more popular queries can take as long as 1 or 2 minutes to complete. SensePlace2 caches query results in an aggressive fashion, which means that identical queries will return much faster on the second run than on the first. Further performance optimization is one of the goals for the Fall 2012 development effort.

In order to keep the user posted about the progress of the latest query, a status message is displayed at the top of the screen. Some of the status messages are directed at SensePlace2 users, while other are meant for the development team and can be somewhat cryptic. We are currently working on building a set of status messages that would be best fit for the general crowd. A typical status message would look roughly like this:

Processing (Tweet list search, heatmap)...

Once the query is complete, the status message disappears, tweet list is populated with matching tweets and point symbols on the map appear. This indicates that it is now possible to interact with the display or initiate a new query.

In an unlikely event of a catastrophic UI failure, try and use the  button. This will preserve the changes you have made to the UI, and will likely fix all of the outstanding problems. If all else fails, use the "Reload" button in your browser.

2. Search Controls

A sample query based on terms “flee” and “Syria” will be used throughout the rest of this guide to demonstrate the functionality of SensePlace2 UI. A screenshot of the entire web interface, taken upon the completion of the query, is shown in Figure 1 below.




Figure 1, Sample SensePlace2 query.

Search controls are located in the purple zone at the top left corner of the web interface (shown closely in a figure below). Search controls can be used in three different ways.

This close-up shows the search controls in the purple sidebar. It includes a 'Search for:' input field containing 'flee Syria'. Below it are two checkboxes: 'Retrieve only tweets with "from" places.' and 'Cluster with univariate k-means'. The 'Cluster with' dropdown is currently set to 'univariate k-means'.

2.1 Free-text query

Current version of SensePlace2 is primarily driven by user queries. The “Search for:” input field allows users to insert one or more query terms of interest. You may currently search for single- or multi-word phrases (e.g. *football riots*), as well as an exact phrase by using double-quotes (e.g. *“football riots”*).

It might be possible to have slight mismatch between the data shown in different parts of the UI. For example, the heatmap may plot a single square on the map indicating a query match, while the tweetlist will show no results. This is due to the fact that we process queries using a combination of relational database search tools

(PostgreSQL with PostGIS extensions) with a high-performance text search engine (Apache Lucene). This combination allows us to run sophisticated queries in near-to real time, yet sometimes results in slight mismatch as described above. This mismatch is minimal and only becomes apparent when few to none query matches are found. Similar to the overall performance of the system, the consistency of query results is one of the areas of our current work.

2.2 Working with “from” places

One of the main features of SensePlace2 is that we extract geographic information from two sources. First source is the *body of the tweet itself* (i.e. the names of locations that are mentioned in the tweet). We refer to information coming from this source as “*about*” locations, as people talk “about” them. Second source is the metadata associated with the tweet that often has explicit geographic coordinates in the form of latitude and longitude. We use the term “*from*” locations to refer to this kind of information, as people send tweets “from” them. “From” tweets are distinguished visually on the map by being shown as green rather than purple circles.

The number of tweets that have “from” locations is quite small (typically about 1.5% of all tweets, somewhat higher in crisis situations), and they tend to be drowned in the stream of relevant tweets with locations of the “about” kind. The checkbox labeled “Retrieve only tweets with “from” places” enables users to only bring back the tweets with “from” locations associated with them.

2.3 Clustering similar tweets


The “**Cluster with**” option allows users to apply one of two text clustering tools that group tweets into a small number of clusters. The resulting clusters are shown at the bottom of the tweet list using a few frequent terms that occur in tweets within the cluster. Clicking on a cluster in this display will bring the tweets from that cluster to the top of the tweet list. Clicking again will return to the default list. The cluster function currently works properly only when queries are limited to single terms (expanding this is on the development list).

3. Overview and Detail

SensePlace2 provides both overview and detail depictions in the timeline, map, and place-tree views, as described below.

3.1 Timeline

Timeline displays the changes in the density of tweets over time and matches the parameters of the user query. Color shaded bands represent the number of matches the given query had in the entire database, with dark red indicating the time span with highest number of tweets. The stacked black bars represent the number of query matches in the list of top 1000 relevant tweets.



Both color bands and the stacked bars use quantile-based classification scheme (quintile and tertile, respectively). By default, the width of the individual color bands is set to one week.

3.2 Map

The map (as shown in Figure 1 above) uses a combination of heatmap and graduated point symbol displays. Heatmap displays the spatial density of tweets that match the term, time and place parameters of the user query using a quantile-based sequential color scheme. Tweet density is calculated using the entire database. Top 1000 relevant tweets are plotted on top of the heatmap using graduated point symbols. Tweets “from” and “about” a particular location are shown as purple and green, accordingly. The size of the point symbols represents the number of relevant tweets referring to that location, while their color density represents the aggregate relevance ranking of those tweets.

3.3 Place-tree

The place-tree highlights the locations that have been mentioned in the query results in a more structured fashion. Each of the nodes in the hierarchy is colored according to the number of matches the given query had in the entire database, whereas the stacked black dots represent the number of matches in the top 1000 tweets. Place-tree is currently populated down to the country level.



4. Temporal controls

The timeline can be manipulated in three different ways. First, timeline sliders can be manually adjusted on either end, as illustrated below.



Multiple fixed time ranges are accessible by clicking the [fixed time range](#) link below the timeline. When first clicked, this option will set the time range to a default width of one month, as shown below.



The width of the time range as well as its units ("one" and "month", respectively) are both adjustable. To adjust the number of months, hover over number [1](#) (a prompt saying [drag](#) will appear), click the mouse, hold the mouse button down, and drag. Timeline will adjust accordingly, as shown below.




You can also click on the label [months](#) to change it to week, day, or year. You can now either click on [re-run time query](#) link, or drag the time range to any part of the timeline to initiate the new query.

It is also possible to change the resolution of the timeline display. Use the radio buttons in the "Bin by:" control to the right of the timeline to set the width of the color bands to either one week (the default) or one day.

5. Tweet List

The tweet list (as shown in figure below) has two visual significations and three kinds of manipulation available to the user.



5.1 Visual significations

5.1.1 Tweet relevance

The narrow bar at the left of each tweet is color coded to indicate relevance to the query (in quartiles, dark depicting more relevant), as estimated by the search engine.

5.1.2 Tweet locations

Locations that the system has identified in each tweet are now highlighted; this is currently used to help users visually recognize tweets relevant to places they are interested in but also to help users identify miscodings of places (see section 5.2.3 below).


5.2 Manipulations

5.2.1 Sorting tweets

Tweets in the tweet list can be sorted by relevance rank, by timestamp, or by their location (currently labeled as “space”). The *location sort* is done based on distance of individual tweets from the current map center. So, at present, to sort tweets based on their proximity to the place of interest, it is necessary to center the map on that place first.

5.2.2 Promotion-demotion of individual tweets

Specific individual tweets can be temporarily promoted (moved to the top of the list) or demoted (moved to the bottom of the list). This is accomplished by first bringing the mouse over the specific tweet, at which point a line saying “[Promote](#) or [demote](#) this tweet. Geocoding [errors?](#)” will pop up, as shown in the figure below.




Clicking on either [promote](#) or [demote](#) will push this specific tweet to the top or bottom of the tweet list, respectively. The results of promotion and demotion will only be visible when the tweet list is sorted by relevance rank, and will be hidden when sorting by time or space.

5.2.3 User input on geocoding errors

This feature is not completely implemented as of the date of this guide. At present, the interface has been implemented, but the system does not yet save any data the user inserts. It is described here as a preview of the near-term features (and because the interface can be tried out). The logic of this feature is described below.

First, user brings the mouse over the specific tweet, at which point a line saying "[Promote](#) or [demote](#) this tweet. Geocoding [errors](#)?" will pop up. User then clicks the on the "[Geocoding errors](#)?" link that will bring up a pop-up window that has the geocoding report controls, as shown in the figure below.



At this point, the user is prompted to select any type of error they would like to report, by clicking on one of the highlighted links. For example, "taking something unrelated" for a location implies that a regular word was taken to be a valid place name. Once the user selects a particular type of error, UI will be automatically expanded to incorporate user input, as shown below.

Below is the original tweet, with locations (a total of 1) highlighted:



Spezzatino di maiale (pork stew) with polenta. Nice winter food, in May, which i 5/7
s feeling like winter. <http://t.co/RUD7nwFX> 2012

Which specific location is mistaken?

Nice

— [add another location?](#)



Include another [misread](#), [misplaced](#), [omitted](#) or [mistaken](#) location, or [submit error report](#).

More complex error cases will request user input concerning the original spelling of the location, its *referent* (e.g. “Georgia” might refer to one of the US states or to a country of its own), as well as a questionnaire-based explanation of what exactly went wrong. For the patient, a GeoNames ID lookup tool is also provided, activated by a click on the “[GeoNames ID](#)” link. An example of a more complex error report is presented below.

Below is the original tweet, with locations (a total of 1) highlighted:



Photo: A mountain airstrip in winter. Memmingen, Germany. <http://t.co/F0iHjLzq> 1/7
2012

Which specific location is misread?

Germany

— [add another location?](#)



How should we read it instead?

How is it spelled in the original tweet?

Memmingen, Germany

What place does it refer to?

Memmingen, Germany

Does it have a [GeoNames ID](#)?

2871992

Now, what happened, exactly?

Did we fragment a location name?

— [yes](#), [no](#)?

Did we omit any of the fragments?

— [yes](#), [no](#)?

Did we merge any of the fragments?

— [yes](#), [no](#)?

Thank you!

Include another [misread](#), [misplaced](#), [omitted](#) or [mistaken](#) location, or [submit error report](#).

The geocoding error report can be expanded at will – all the user needs to do is click on the type of error they want to add (“misread”, “misplaced”, etc.) once more, using the links at the bottom of the report.

Once you start working with the geocoding error tool, a line saying “This report will be filed anonymously. Would you like to [log in](#) instead?” will appear at the bottom of the report. If you click the [log in](#) link, you will have a chance to pick a pseudonym that can later be used to link your reports together into a cohesive story. The pseudonym you pick can be changed at any time.

6. Map

The map includes several actions that enable users to focus attention on places and regions.

Mousing over a point on the map will highlight the tweet(s) associated with the point in the tweet list (if they are visible). Similarly mousing over a tweet in the list will highlight the point symbols associated with it. If a location is connected to any other location or locations by joint mentions within a tweet (in the top 1000 tweets currently in the display), then connector lines are drawn to show the connected locations. The width of the line depicts quintiles of frequency for connections (bold lines represent more connections), as shown in the figure below.




Clicking on any point symbol on the map will bring the tweet(s) associated with that point to the top of the list (multiple tweets can be selected in succession). Clicking the same point symbol again will deselect that particular tweet, while clicking in a blank space will turn all of the promotions off and put the list back to its default. When a place is clicked that has connections, the connections remain visible as long as the place is highlighted. Thus, it is possible to click on a few places in succession (without clicking a blank space to clear the selections) to build up a network of connections from a few selected places.

An *Alt+Click* combination, when performed anywhere on the map, will launch a new query using the current query terms and a spatial constraint that brings back the 1000 tweets closest to the point clicked. Keep in mind that the 1000 most relevant tweets returned might have mentions of locations outside the desired region.

7. Place-tree

The place-tree (as shown in the figure below) has a number of user controlled features.




Users can toggle between a place-tree that shows the full set of locations (as described by the GeoNames hierarchy) and a “pruned” hierarchy that only shows places that match the user query parameters. The switch is performed using the **“Switch hierarchy”** button.

Users can select one or more places of interest using check boxes positioned next to them, which would highlight the tweets related to that particular location in the tweet list. Use **“Drop selection”** button to clear all of the check boxes set.

Finally, although the capacity to launch queries based on GeoNames IDs of the features selected in the place-tree has been put in place, the server side of this functionality is still under development. Thus, **“Query selected”** button is currently disabled.

8. Tag cloud

Tag cloud (as shown in the figure below) displays the list of locations that are most frequently mentioned in the query results. The size of the words in the tag cloud is proportionate to the number of mentions, and the words themselves can be clicked in order to filter the contents of the tweet list.



1. Informed Consent

The Pennsylvania State University, Title of Project: Geovisual Technology Use and Usability Assessment

Principal Investigator: Dr. Alan M. MacEachren, 302 Walker Building, University Park, PA 16802, (814) 865-7491; maceachren@psu.edu

*** 1. Purpose of the Study:** The purpose of this research study is to evaluate which visualization techniques are effective for identifying patterns in the display of geographic data and the changes that can be seen over time and space.

2. Procedures to be followed: You will be asked to participate in one or more of the following activities. The researcher providing this form will describe exactly which activities you will participate in and answer any questions you have about the procedures for each. Briefly the range of tasks include the following:

- Task support:** You will be given a brief demonstration of the intended use of a software application and then you will be given a set of instructions for completing several tasks with this software. Our interest is in how well suited the software is for performing a set of tasks. Though we would like participants to perform the tasks to the best of their ability, we are evaluating the software, not any one individual's performance.

- Online survey:** You will be asked to complete a short survey about your domain expertise and prior experience using geovisualization software.

3. Discomforts and Risks: The risks of participation are minimal. At any time during the experiment, you can terminate the session should you experience discomfort.

4. Duration/Time: You will be asked to participate in a session that is expected to last 60 minutes.

5. Benefits: Your input will result in improved usability and utility of geovisualization software developed at the Penn State GeoVISTA Center. Your input will also help to shape the next generation of geovisualization and geovisual analytics tools. You will also receive \$20 in compensation for your time.

6. Statement of Confidentiality: Your participation in this research is confidential. The data will be stored and secured at 211 Walker Building in a locked / password-protected file. In the event of a publication or presentation resulting from the research, no personally identifiable information will be shared. The Pennsylvania State University's Office for Research Protections and Institutional Review Board, and the Office for Human Research Protections in the Department of Health and Human Services may review records related

to this project. If the task you are completing involves the use of the internet, your confidentiality will be kept to the degree permitted by the technology used. No guarantees can be made regarding the interception of data sent via the Internet by any third parties.

7. Right to Ask Questions: Please contact Krista Kahler at (814) 865-9655 with questions, complaints or concerns about this research. You can also call this number if you feel this study has harmed you. If you have any questions, concerns, problems about your rights as a research participant or would like to offer input, please contact The Pennsylvania State University's Office for Research Protections (ORP) at (814) 865-1775. The ORP cannot answer questions about research procedures. Questions about research procedures can be answered by the research team.

8. Voluntary Participation: Your decision to be in this research is voluntary. You can stop at any time. You do not have to answer any questions you do not want to answer. Refusal to take part in or withdrawing from this study will involve no penalty or loss of benefits you would receive otherwise. Your grades or employment status will not be affected if you choose not to participate. You are welcome to contact someone other than the Principle Investigator with questions/concerns.

9. Participant Requirements: Participants must be 18 years of age or older to take part in this research study. You are also required to have adequate expertise in the domain of interest.

Completion of the tasks described above for this research implies your consent to participate. You should print a copy of this consent notice for your records.

☐ Agree

☐ Disagree

2. Your Background

What is your gender?

☐ Female

☐ Male

Which category below includes your age?

☐ 18-20

☐ 21-29

☐ 30-39

☐ 40-49

☐ 50-59

☐ 60 or older

Are you currently pursuing a degree?

☐ Yes

☐ No

3. Your Background

Which degree are you currently pursuing? (Degree Type & Major)

I consider myself to be knowledgeable in the following broad areas:

	Strongly Disagree	Disagree	Neither Agree nor Disagree	Agree	Strongly Agree
Geographic Information Systems	☐	☐	☐	☐	☐
Information Science	☐	☐	☐	☐	☐
Computer Science	☐	☐	☐	☐	☐
Social Sciences	☐	☐	☐	☐	☐
Arts and Literature	☐	☐	☐	☐	☐
Physical and Biological Sciences	☐	☐	☐	☐	☐

**Which keywords would best describe your research or professional work interests?
(please enter 3-5 keywords)**

Have you used social media for personal purposes? (check all that apply)

- ☐ I have/had a Facebook account
- ☐ I have/had a Twitter account
- ☐ I have/had a LinkedIn account
- ☐ I have/had a MySpace account
- ☐ I use another social media application; please specify:

Have you used social media for professional purposes? (check all that apply)

- ☐ I have used Facebook for professional purposes
- ☐ I have used Twitter for professional purposes
- ☐ I have used LinkedIn for professional purposes
- ☐ I have used MySpace for professional purposes
- ☐ I have used another social media application for professional purposes; please specify:

4. Task 1

Task 1: Explore Geospatial Social Media from a Recent Event

To get started, type "whooping cough" into the "Search for" box at the top left of the SensePlace2 interface. Press the enter key to search for Tweets about whooping cough.

After SensePlace2 has finished loading the results for your search, take a few moments to explore the geographic locations mentioned in Tweets about whooping cough. Remember, the heat map shows the overall prevalence of location mentions in every Tweet about whooping cough, while the purple proportional circles show the locations (and number of Tweets talking about each location) for the top 1000 most relevant Tweets about whooping cough.

Next, use the timeline to identify and explore which time ranges include the most activity for Tweets about whooping cough. You should see peaks in activity at certain times. Narrow the time range to explore each peak in activity.

What Geographic patterns do you see?

What Temporal patterns do you see?

What did you learn from the content of the Tweets themselves?

Provide at least two questions that you would ask another analyst to explore after seeing these patterns.

5. Task 2

Task 2: Analyze contrasting Geographic patterns in About/From locations in Tweets

To begin this task, type "earthquake" into the "Search for" box at the top left of the SensePlace2 interface. Press the enter key to retrieve Tweets about earthquakes.

Once SensePlace2 has loaded these Tweets, take a few moments to explore the patterns of locations that are mentioned. Next, using the timeline, narrow the dataset to show only those Tweets from August 2012 until the present time.

Take some time to explore the data from just this time range and see if you can identify key events that received the most attention across these dates. Once you've done this, click the checkbox under the "Search for" box to "Retrieve only tweets with "from" places." This will have SensePlace2 refine your search by one more step and only show those tweets that included a reporting location (e.g. assigned by a phone or other means to indicate where somebody was when they tweeted).

Use the "Current task" dropdown list to switch back and forth between each of these two queries and note similarities/differences in what you see in terms of Tweet content and the patterns of relevant locations.

What geographic patterns do you see in the Tweets About places versus those where Tweets were reported From places?

What Temporal patterns do you see in these two types of locations?

What did you learn from the content of the Tweets themselves?

Provide at least two questions that you would ask another analyst to explore after seeing these patterns.

6. Task 3

Task 3: Explore geocoding issues and evaluate interface methods for submitting corrections

To get started with the third and final task for this evaluation, type "fire" into the "Search for" box at the top left of the SensePlace2 interface. Press the enter key to retrieve Tweets about fire.

Once SensePlace2 has loaded these Tweets, your task is to identify tweets that have locations that exhibit one or more of the following problems:

1. The locations in a Tweet were not correctly highlighted (e.g. they include a placename that does not show a blue background in the Tweet list)
2. The locations in a Tweet are misread by our system (e.g. SP2 highlighted Lafayette but did not include the word "Lake" before it, when the original Tweet says "Lake Lafayette")
3. The locations in a Tweet are misplaced by our system (Using the map, you determine that the circle referring to the place mentioned in a Tweet is in the wrong place)
4. The locations in a Tweet are mistaken by our system (SP2 has highlighted a placename that you know is not in fact a real placename)

You should identify at least one example of each problem from the results shown for the search for Tweets about fire. For each problem you discover, hover over the Tweet in the Tweet list and click the "Geocoding errors" link to launch our geocoding correction interface. Choose the right type of error you wish to report for each example and submit your report when you are ready.

Please identify and suggest fixes for at least 10 geocoding errors you discover in the Tweets. When you are finished doing this, continue through the rest of this survey to record your feedback.

SensePlace2 allows you to make corrections for a range of geocoding errors. Are there other error types that should be fixable that are not currently supported?

What functionality would you add (or take away) from the SensePlace2 interface for handling geocoding errors in Tweets?

In your opinion, what is an acceptable proportion of results having location accuracy or precision problems when working with social media in a tool like SensePlace2?

- ☐ Less than 1%
- ☐ Between 1-5%
- ☐ Between 5-10%
- ☐ Between 10-20%
- ☐ Greater than 20%

Other (please specify)

7. Usability

The following questions ask you to rate the Usability of SensePlace2 now that you have completed three basic Tasks using the software.

I think that I would like to use SensePlace2 frequently.

	Strongly Disagree	Disagree	Neither Agree nor Disagree	Agree	Strongly Agree
Rate Your Opinion	☐	☐	☐	☐	☐

I found SensePlace2 to be simple.

	Strongly Disagree	Disagree	Neither Agree nor Disagree	Agree	Strongly Agree
Rate Your Opinion	☐	☐	☐	☐	☐

I thought SensePlace2 was easy to use.

	Strongly Disagree	Disagree	Neither Agree nor Disagree	Agree	Strongly Agree
Rate Your Opinion	☐	☐	☐	☐	☐

I think that I could use SensePlace2 without the support of a technical person.

	Strongly Disagree	Disagree	Neither Agree nor Disagree	Agree	Strongly Agree
Rate Your Opinion	☐	☐	☐	☐	☐

I found the various functions of SensePlace2 were well integrated.

	Strongly Disagree	Disagree	Neither Agree nor Disagree	Agree	Strongly Agree
Rate Your Opinion	☐	☐	☐	☐	☐

I thought there was a lot of consistency in the SensePlace2 interface.

	Strongly Disagree	Disagree	Neither Agree nor Disagree	Agree	Strongly Agree
Rate Your Opinion	☐	☐	☐	☐	☐

I would imagine that most people would learn to use SensePlace2 very quickly.

	Strongly Disagree	Disagree	Neither Agree nor Disagree	Agree	Strongly Agree
Rate Your Opinion	☐	☐	☐	☐	☐

I found SensePlace2 to be very intuitive.

	Strongly Disagree	Disagree	Neither Agree nor Disagree	Agree	Strongly Agree
Rate Your Opinion	☐	☐	☐	☐	☐

I felt very confident using SensePlace2.

	Strongly Disagree	Disagree	Neither Agree nor Disagree	Agree	Strongly Agree
Rate Your Opinion	☐	☐	☐	☐	☐

I could use all of SensePlace2 without learning anything new.

Strongly Disagree

Disagree

Neither Agree nor
Disagree

Agree

Strongly Agree

Rate Your Opinion

☐☐☐☐☐

8. Basic Capabilities

I was able to obtain information about the *places* relevant to each task I completed.

	Strongly Disagree	Disagree	Neither Agree nor Disagree	Agree	Strongly Agree
Rate Your Opinion	☐	☐	☐	☐	☐

I was able to obtain information about the *time periods* relevant to each task I completed.

	Strongly Disagree	Disagree	Neither Agree nor Disagree	Agree	Strongly Agree
Rate Your Opinion	☐	☐	☐	☐	☐

I was able to obtain information about the *topics* relevant to each task I completed.

	Strongly Disagree	Disagree	Neither Agree nor Disagree	Agree	Strongly Agree
Rate Your Opinion	☐	☐	☐	☐	☐

9. Situational Awareness

SensePlace2 allows me to easily perceive the key spatial, temporal, and attribute elements that are relevant to a crisis situation.

	Strongly Disagree	Disagree	Neither Agree nor Disagree	Agree	Strongly Agree
Rate Your Opinion	☐	☐	☐	☐	☐

Describe the reason for your answer.

SensePlace2 allows me to *understand the relationships* between spatial, temporal, and attribute data related to a crisis situation.

	Strongly Disagree	Disagree	Neither Agree nor Disagree	Agree	Strongly Agree
Rate Your Opinion	☐	☐	☐	☐	☐

Describe the reason for your answer.

SensePlace2 allows me to predict what may happen in future crisis situations.

	Strongly Disagree	Disagree	Neither Agree nor Disagree	Agree	Strongly Agree
Rate Your Opinion	☐	☐	☐	☐	☐

Describe the reason for your answer.

10. System Utility

SensePlace2 allows me to easily identify geocoding errors.

	Strongly Disagree	Disagree	Neither Agree nor Disagree	Agree	Strongly Agree
Rate Your Opinion	☐	☐	☐	☐	☐

Describe the reason for your answer.

When a geocoding error is found, SensePlace2 allows me to easily suggest a change.

	Strongly Disagree	Disagree	Neither Agree nor Disagree	Agree	Strongly Agree
Rate Your Opinion	☐	☐	☐	☐	☐

Describe the reason for your answer.

SensePlace2 has the right balance of tools to tell a compelling story about a crisis situation based on social media reports.

	Strongly Disagree	Disagree	Neither Agree nor Disagree	Agree	Strongly Agree
Rate Your Opinion	☐	☐	☐	☐	☐

Describe the reason for your answer.

SensePlace2 would be helpful when generating an analytical report to share with a decision maker to prompt actions before, during, or after a crisis.

	Strongly Disagree	Disagree	Neither Agree nor Disagree	Agree	Strongly Agree
Rate Your Opinion	☐	☐	☐	☐	☐

Describe the reason for your answer.