# Evaluating Social Navigation Visualization in Online Geographic Maps

Yuet Ling Wong[a], Jieqiong Zhao[a] & Niklas Elmqvist[a,b*]

[a]Purdue University, West Lafayette, IN, USA

[b]University of Maryland, College Park, MD, USA

**Abstract**

Social navigation enables emergent collaboration between independent collaborators by exposing the behavior of each individual. This is a powerful idea for web-based visualization, where the work of one user can inform other users interacting with the same visualization. We present results from a crowdsourced user study evaluating the value of such social navigation cues for a geographic map service. Our results show significantly improved performance for participants who interacted with the map when the visual footprints of previous users were visible.

## 1 Introduction

We humans invariably leave marks on the physical world as we move through it, and these marks are multiplied by each individual. Our footprints collectively coalesce into trails, our fingers wear down the faces of buttons involved in a passcode, and our car wheels leave tracks or even ruts in the road as we pass. While much of this impact on our environment may be detrimental—pollution, trash, and resource depletion comes to mind—such collective marks may also sometimes serve as mechanisms for asynchronous communication between individuals: they tell of the safest or shortest path from one location to another, they remind us of the passcode to enter the building, or they indicate which backroad our friends took in getting to the rental cabin. In short, these mechanisms enable action not based on spatial or semantic information, but on *social* information. However, while such asynchronous communication enabling so-called *social navigation* [8] is intrinsic to the physical world, it is not a natural property of the digital world, where interacting with an object typically leaves no tangible mark.

One specific form of digital artifact that can benefit particularly from adopting the concept of social navigation is **web-based visualization** on the Internet—here is why:

1. Sensemaking using visualization benefits from involving multiple collaborating users [20, 23];

*Corresponding author: 2117H Hornbake Building, South Wing, College Park, MD 20742, USA. Phone: +1 (301) 405-2038, e-mail: elm@umd.edu

2. Web-based visualizations attract a large potential audience, facilitating emergent collaboration [37];

3. The web ecosystem already includes the server-side infrastructure necessary to store social navigation data; and

4. Visual representations of social navigation can easily be integrated with or overlaid on the existing visualization.

While some work has already been conducted on this topic, such as for crowdsourced graph layout [42], information scent displayed on interface widgets [41], and visual footprints for information seeking [40], no general approach to integrating social navigation in web-based visualization has been proposed. In particular, no results from an evaluation of this idea exists in the literature.

In this paper, we begin to remedy this gap in the literature by conducting a crowdsourced user study of how visual representations of social navigation data can facilitate emergent collaboration between independent users of a web-based visualization. Instead of trying to address the entire field of web-based visualization, however, we delimit our study to a smaller subset of web-based visual applications: geographic map services, such as Google Maps and Bing Maps. Our motivation for this is twofold: (a) geographic map services are highly visual web applications that are familiar to a large segment of the population, and (b) navigation on a digital geographic map is easily visualized as spatial trails, similar to footprints on a beach. By showing such visual footprints on a geographic map, the user is effectively exploring the map together with a crowd of earlier users without ever having to communicate directly with them.

In performing this evaluation study, we are following in the footsteps of previous work such as Hotmap [10] and TrailMap [43]. However, whereas Hotmap used off-line user navigation data for the Microsoft Live Search map service and visualized it primarily for analysts, we capture interaction data in an on-line manner and feed back visual representations of this data to the end-users themselves. As for TrailMap, the idea with their system was to automatically capture spatial bookmarks for a single user, whereas we utilize this idea to support emergent collaboration [37] between multiple independent users. Our evaluation is a crowdsourced user study conducted using Amazon's Mechanical Turk service where more than 100 participants took part in a gamified multiscale search task looking for targets placed at different zoom levels on a world map. The objective of the task was to find as many targets as possible within a five-minute time limit. Our web-based platform logged the navigation behavior of each participant and visualized these as rectangular footprints representing the viewport position of each participant. The participants were divided into groups depending on whether they (a) had access to visual footprints, and whether (b) the shown footprints were collected from participants who had access to footprints or not. Our results show significantly better search performance (represented by the number of found targets of different difficulty levels) for participants who had access to visual footprints, particularly when those footprints came from participants who were not searching the map randomly but who in turn also had access to visual footprints.

## 2 Background

Our work in this paper lies at the intersection of the following research areas, which we review in detail below:

- **Navigation in Large Visual Spaces:** Supporting advanced interaction and navigation in visual spaces;
- **Visit and Edit Wear:** Recording visits and viewing of digital objects for conveying other user's activities; and
- **Emergent Collaboration and Social Navigation:** Methods for crowd-powered collaboration over the Internet.

## 2.1 Navigation in Large Visual Spaces

Navigating in large-scale, possibly multi-scale [30], visual spaces has long been a core HCI problem [13] imbued with several difficult challenges, including retaining overview while seeing details [11, 12], maintaining multiple regions of interest [34], and traversing multi-scale spaces lacking navigational cues [27]. As a result, much research has been focused on developing effective navigation techniques; examples include OrthoZoom [5] and SDAZ [22], which provides variable zoom rate control, and content-aware scrolling [25], which adapts scrolling based on visible content.

A digital map represents one example of such a large visual space. Geographic map services have lately become one of the most commonly used services on the Internet, with sites such as MapQuest, Google Maps, and Bing Maps numbering among the top 100 accessed websites on the Internet according to the Alexa ranking [2]. For this reason, much research has been devoted to improving the usability and efficiency of such geographic maps. Magic Lens [6] techniques such as DragMag [38], high-resolution magnification lenses [4], and Sigma lenses [31] provide methods for viewing and navigating multi-scale structures that have been applied to geographic spaces. Similarly, dynamic insets [16] support navigation by showing the context of off-screen destinations, and PolyZoom [26] creates a hierarchy of focus regions to track map navigation.

## 2.2 Visit and Edit Wear

Capturing interaction in a user interface is a common method in human-computer interaction primarily due to the need to support undo and redo operations [1] as well as for managing navigation histories in web browsers [36]. However, more recent efforts have also started to focus on using interaction logs as graphical histories of the user's interaction [29]. Heer et al. [19] present a complete analysis of the design space of interaction histories.

Actually associating interactions with the digital objects they operate on has been referred to as *read wear* and *edit wear* [21], and is analogous to physical wear: interacting with the digital object will leave markings that track the usage. This is primarily for the user's own benefit, but exposing this wear to other collaborators enables a mechanism called social navigation (discussed below). These concepts have given rise to a range of ideas, including breadcrumbs navigation for web design, aiding the user's memory when reading documents, and supporting revisitation for large visual spaces [35]. The Footprints scrollbar [3] is a novel scrollbar widget based on read wear; it adds visual marks to the scrollbar region to oft-visited (and presumably important) parts of a document to aid revisitation.

Our basic idea in this paper is to apply visit wear to a geographical map to improve and guide multiple users navigating in the map. Most related to our work is TrailMap [43], which analyzes a user's mouse interactions to implicitly create spatial bookmarks in a geographic map. The prototype implementation uses the Bing Maps API and was evaluated using a one-week longitudinal deployment with 11 participants. However, whereas TrailMap is targeted at map revisitation tasks and associates implicit bookmarks with search queries, our work is aimed at social navigation settings, and was accordingly evaluated using a large-scale crowdsourced user study [18].

### 2.3 Emergent Collaboration and Social Navigation

Whereas the field of computer-supported cooperative work (CSCW) traditionally focuses on settings where the users know they are working together, Terveen and Hill propose the notion of *emergent collaboration* [37] as settings where the participants are not actively working together, yet their efforts can be combined as if they were. They quote the forming of paths through the woods as a motivating example, noting that the unified efforts of past wanderers serve as both a history of use as well as a resource for future users.

Combining ideas from emergent collaboration and usage wear gives rise to the notion of *social navigation* [8], where users are guided not by the spatial or semantic structure of the information space, but by an awareness of the actions of others. Wexelblat [39] discuss the use of interaction histories to enable social navigation and present several tools based on this idea. One popular theme in this research is the use of trails to aid social navigation, particularly for document reading and navigation [14, 15, 32, 40]. Social navigation has also been applied to search, yielding the concept of social search [7, 9].

Beyond document spaces, social navigation has also been applied to more visual representations. For example, Yuan et al. [42] propose a graph layout algorithm that combines the input of many users to create a merged 2D layout for a node-link diagram that is aesthetically more pleasing than an automatic layout. The collaborative visual analysis tool Cambiera [24] uses visual marks to indicate whether a particular document has been previously read and whether the other user is currently reading it or not. Similarly, scented widgets [41] overlay visual representations of data on top of user interface widgets; one potential data source for these visual representations are social activity metrics, such as typical settings for sliders, or the number of views for a specific option for a radio button or group of check boxes.

## 3 Social Navigation in Geographic Spaces

People often navigate the physical world based on the activity of others; for example, a crowd forming may indicate an event of interest [41], or a footpath through the woods may indicate the safest or shortest way through [37]. Similarly, such trails also give information for how a person may intentionally avoid the crowd and "get off the beaten path." In the context of computer systems, such navigation is often referred to as *social navigation* [8]. Our main goal with this work is to investigate appropriate and effective social navigation methods for casual geographic information spaces, such as the interactive map services provided by websites such as Google Maps, MapQuest, and Bing Maps.

What is the purpose of providing social navigation cues in online map services? Let us briefly review a couple of illustrative usage examples and motivations for this work. Geographic attention as a general concept has a wide variet of potential applications [10], including for map service developers wanting to optimize their service, designers wanting to understand and improve the visual design of the map, and researchers wanting to understand geographic attention with regards to cultural, political, or topological factors. However, introducing a social aspect to geographic attention means that the activity of map users should be fed back into the map itself. This mechanism can be used to aid navigation to areas on a map that is of current wide interest to a large audience (such as the site of major news story), as well as a source for serendipitous discovery of ongoing events based on geographic location (i.e., essentially the opposite of the former). A third potential use would be to help a viewer diverge from areas that hold the current interest of a large audience; essentially taking the "road less traveled."

[Insert Figure 1 here.]

Figure 1 presents an overview of a general social navigation framework for geographic spaces. As can be seen from the figure, the two most important components of the framework, besides the map service itself, is (1) the **navigation logging** component, which records user interaction to an event database, and (2) the **visual footprints**, which feed interaction data from previous users back to the current user.

[Insert Figure 2 here.]

## 3.1 Navigation Logging

A key aspect of social navigation for geographic spaces is determining *which* navigations should be logged. The purpose of the logging mechanism is to convey the collective navigation behavior of all of the past users who have utilized the system for the benefit of the current user.

### 3.1.1 Design Constraints

The above goal gives rise to two important yet conflicting constraints: that (a) *all important navigation events should be captured* to faithfully represent the crowd's behavior, but that (b) the *total amount of navigation events should be minimized* to avoid high complexity in the resulting visual output. In essence, this amounts to logging only vital interactions.

Another issue is to determine *what* to record. Since we are interested in higher-level navigation events, it makes sense to disregard low-level mouse and keyboard interaction, and instead focus on logging the viewport, which could include its position (the viewport footprint), its direction, or its speed.

### 3.1.2 Logging Policies

Based on our design constraints, we derive a set of possible policies for tracking the user's viewport in geographic space:

- **Time-interval logging:** Perhaps the most basic of event logging policies, this policy will log viewport characteristics at constant time intervals (for example, every half second). Logging should pause if the user is inactive, such as when no interaction has been detected for five seconds.

- **Space-interval logging:** Instead of using regular time intervals, logging can be tied to spatial intervals so that the viewport is logged every time the viewport has been moved a certain screen distance, such as 100 pixels. Special care must be taken to log zooming, such as when the zoom level has changed two or more zoom levels.
- **Salient event logging:** A more intelligent logging policy takes the semantics rather than the syntax of multiscale navigation into account. Such a policy would tie logging to salient events that are representative of the user's navigation behavior, such as when the user pauses, attempts to acquire a target with the cursor, or drills down on a target.

The choice of logging policy (or any other conceivable logging policy) to use depends on the application and tasks.

#### 3.1.3 Storage and Aggregation

Finally, all logged data is transferred back to a server-side interaction database (Figure 1) where it is stored for the benefit of future users. During this stage, it is possible to prune or aggregate the interaction data to make it more efficient for presentation. This includes eliminating fully enclosed viewport records or removing duplicates. Furthermore, this is also the stage for introducing temporal decay where the effect sizes of events are gradually decreased as the event ages, only to eventually be pruned entirely from the database.

### 3.2 Visual Representations

The second major component of social navigation for geographic maps is the visual representations of the navigation event database, which we call *visual footprints* in Figure 1. The purpose of these footprints is to communicate the contents of the interaction database back to the current user, thereby closing the loop on the feedback loop. Based on the visual representations proposed by TrailMap [43], the most useful technique is in-situ visual marks—called *viewprints*—representing the viewport position recordings in the interaction dataset. This has the benefit of being both non-intrusive, since the original map visualization is not altered, only augmented with visual marks, and presumably self-explanatory, since the marks appear as "ghost viewports" on the map itself. Furthermore, by choosing a discrete visual representation as footprints, the effect size can be communicated as the number of footprints in an area accumulates.

Figure 2 show three different concrete visual designs for our viewprint marks. These designs all have their relative strengths and weaknesses; for example, viewprint corners (Figure 2(a)) add the least amount of new visuals to the map, but may yield ambiguous representations, whereas viewprint areas (Figure 2(c)) eliminate ambiguity at the cost of higher visual complexity. Bounding box viewprints (Figure 2(b)), in contrast, add less graphics to the map than colored areas, yet provide less ambiguity than corners. Significantly, while colored areas come close, we avoid the heatmaps used in Hotmap [10] because they are space-filling and thus have a heavy impact on the underlying visual representation.

## 3.3 Implementation Notes

We have implemented a prototype geographic map service that incorporates visual cues for social navigation activity. The prototype is a hybrid system consisting of both a thin JavaScript client and a server-side PHP web service. For the client side, the tool uses the Google Maps JavaScript API v3 (`https://developers.google.com/maps/`) to render a world map augmented with visual footprints downloaded from the PHP web service. An event manager automatically collects the user's viewport manipulations (based on the logging policy) and periodically uploads the recorded events every 10 seconds. Viewprints are drawn based on a client-side setting and uses temporal decay so that older footprints have lower opacity. Each viewprint is automatically scaled so that they do not become smaller than $3 \times 3$ pixels in order to ensure visibility. Furthermore, viewprints that are larger than the current size of the viewport are hidden to minimize visual complexity.

For the server side, our PHP web service accepts asynchronous (AJAX) connections from JavaScript clients to receive interaction data as well as to transmit aggregated viewprints data. The service uses a PostgreSQL database to store all recorded interaction events. We also use a maintenance service that periodically prunes the database of old events.

# 4 User Study

The purpose with our work is to study how visualization of social navigation [8] in geographic map services can lead to emergent collaboration [37] between individuals by transparently providing visual footprints of many users interacting with a geographical space on the Internet. To this effect, we conducted a controlled experiment on search performance of multiscale targets in a geographic space for human participants with and without visible footprints.

[Insert Figure 3 here.]

## 4.1 Design Decisions

We initially performed a large-scale pilot study with a total of 93 participants using Amazon's Mechanical Turk. The pilot consisted of using our prototype tool to search for targets dispersed around a world map. The purpose was to fixate design aspects of the study, find suitable values for parameters of the system, and answer several of the open design questions we had about our experiment. Below we review the most important design decisions we derived from this pilot:

- **Crowdsourced:** Our pilot testing confirmed that a crowdsourced evaluation of this phenomenon is possible. Supported by findings from Kittur et al. [28] as well as Heer and Bostock [18], crowdsourcing is thus a cost-effective and practical approach to evaluating social navigation in maps with a large number of participants.

- **Gamified task:** Several crowdsourced participants never finished the simple map navigation task in the pilot. This indicates that the study task must be compelling and not too arduous. For this reason, we chose to design the study task to be a competitive game that was limited to five minutes in duration.

- **Salient event logging:** Several of our event logging policies yielded large footprint datasets; for example, temporal and spatial interval logging both resulted in a large number of footprints. This, in turn, yielded high visual clutter and noticeably slowed down the rendering performance for our prototype system. Instead, we decided on using a salient event logging policy, where interaction data (the viewport position) is only recorded for salient events (such as when the user stops to click on a target).

- **Bounding box:** Informal feedback and performance results indicated that our *bounding box* representation of the viewport position at the time of the logged event was most preferred by our pilot study participants, and also constituted a good compromise for minimizing clutter and complexity when visualizing large footprint datasets. We also found out that the visual representations of the bounding boxes captured at higher zoom levels were almost invisible at lower zoom levels. To counter this, we maintained the sizes of these bounding box to be constant untillits original zoom level. We also chose to not show the bounding boxes above their original zoom levels.

- **Guided visual search:** Our initial strategy of randomly dispersing search targets at different zoom levels on the world map turned out to be too difficult: our pilot study participants only found approximately 10% of the total number of targets in the geographical space. Therefore, we decided to limit target placement to specific areas (close to coastlines) and to a small set of zoom levels.

- **Off-line logging:** While logging and visualizing own navigation could help a participant remember which part of the map had already been visited, we found that this was often confusing for our pilot study participants. Instead, we chose to perform all logging in an off-line manner so that a participant's navigation did not interfere with their own performance. In fact, we decided to split our participants into disjoint groups whose footprints were logged and visualized separately.

## 4.2 Participants

We recruited 136 participants for our user study using Amazon's Mechanical Turk (MTurk) service. To eliminate participants who only accept work on MTurk to earn money with a minimum of effort and without paying attention to the task, we filtered out all participants who scored less than 10 points. Most of these filtered participants did not perform a single action in the task and received a score of 0. This is common practice for crowdsourced user studies [33], and caused us to remove 36 participants from our results. Our final participant number was 100 (35 females). Their ages ranged from 19 to 61 (mean 30.4, s.d. 7.8). No participant self-reported as color blind.

The baseline payment was $0.50 for completing a full treasure hunt task, and an additional $0.50 was given to the top 20% participants with the highest scores.

## 4.3 Task and Dataset

Based on our findings from the pilot study, we designed our task as a treasure hunt game where participants were asked to find as many targets as possible on a world map during a 5-minute time interval. Each found target added to the participant's total score, which were tallied at the end of the game session. A participant marked a target as found by double-clicking it, thereby adding its points to the participant's total score. Targets were represented using the default Google Maps marker icon. A found target was drawn using a green color and was always visible, regardless of zoom level. Targets that were not yet found were either red if visible but not yet clickable, or yellow if both visible and clickable.

To make the task sufficiently challenging, targets were placed on different zoom levels (Google Maps uses a total of 23 zoom levels, where zoom 0 corresponds to a map of the Earth fully zoomed out, and higher numbers represent increasing magnification of the view) and were only clickable when the viewport was zoomed to that level and onward. For example, a target at zoom level 8 could be clicked on (and was drawn in yellow color) starting at zoom level 8 and all the way to 22 (the highest magnification). Furthermore, a target was also visible (but not clickable) at an additional 5 zoom levels above its placement. In other words, the target at zoom level 8 was visible but not clickable (and thus drawn in red) already at zoom level 3. Finally, to further constrain the task and guide the participants, all targets were placed on coastlines, defined as within 10 kilometer from an ocean.

We used four different zoom levels for all targets, representing very easy (zoom level 5, so visible in the zoomed-out view), easy (10), medium (15), and hard (20). A total of 150 targets were placed on the world map: 10 very easy targets, 20 easy ones, 40 medium ones, and 80 hard ones. The number of points awarded for finding a target was proportional to its difficulty level, on the scale 1, 2, 4, and 8 points.

## 4.4 Platform

We used our prototype social navigation map service as the testing platform for the user study. The test platform automatically integrates with Mechanical Turk for worker allocation as well as payment. Furthermore, we implemented the time-limited game task within the prototype service.

In choosing a suitable event logging policy and visual footprint design, we used our findings from the large-scale pilot study as a guide. Due to the large number of participants in the user study, we opted for a salient event logging policy where interaction data was captured only for situations when the user paused the map navigation, such as to click on a target. Similarly, informal feedback from the pilot study participants suggested to use the *bounding box* viewprints method as the visual footprint because of its discrete appearance and low visual clutter. To further minimize the visual complexity, we chose to use a red color with high transparency. The intention was that a single viewprint would be barely visible, but that a large concentration of them in the same area would quickly begin to stand out. We also disabled temporal decay.

## 4.5 Experimental Conditions

We logged and stored the navigation behavior for all participants in the user study. However, to evaluate the utility of visual footprints for social navigation in geographic spaces, we distinguish between two separate aspects of our footprints: whether (a) the participants saw visual footprints or not, and, if so, whether (b) the footprints came from participants who in turn had seen footprints or not. This led us to split the participant pool into four groups (where all participants in the same group saw the same footprint data):

- **Group 1 – No knowledge:** These participants did not see any footprints at all and thus explored the map "blindly";
- **Group 2 – Unguided footprints:** These participants saw footprints from Group 1, which were in turn unguided;
- **Group 3 – Hybrid footprints:** These participants saw footprints from both Group 1 and 2, who were both unguided and guided, respectively; and
- **Group 4 – Guided footprints:** These participants only saw footprints from Group 2, who were guided.

Figure 3 shows screenshots of the maps shown to the four groups described above. Since target positions did not change across instances (a necessity for visual footprints to be of any use), we used a between-subjects design for the groups to avoid systematic learning effects. We also used MTurk's worker control mechanisms to ensure that no same worker participated more than once in the experiment.

## 4.6 Procedure

A single task session started with the participant reading an information sheet that served as a consent form. They clicked a button to give their consent and were then given an illustrated instruction screen. Clicking on this button in turn launched the treasure hunt game with the countdown timer initialized at 300 seconds, the total score at 0 points. A scoreboard which contained the highest score in record was shown in order to motivate participants.

During the treasure game itself, participants were able to navigate in the prototype system's view of a Google Maps world map. They were able to zoom and pan using normal Google Maps interactions, including dragging to pan, mouse wheel for zooming, double-click to zoom in on an area, or using the standard navigation control. Double-clicking on a yellow target also marked it as found, changing its color to green. A simple example of how to find a target and get scores is demonstrated in Figure 4.

[Insert Figure 4 here.]

After the 5 minutes had expired, the test platform automatically advanced to a demographic questionnaire. Completing this provided the participants with the code used to redeem the payment for the task. Additional incentive rewards were then paid out within 24 hours to the top 20% highest-scoring participants for each group.

[Insert Figure 5 here.]

## 4.7 Hypotheses

Based on our user study design and the grouping based on visual footprints representation, we formulated two hypotheses on the outcome of our experiment:

**H1** *Visual footprints will result in significantly higher score.* This is the basic premise of our work; that visualizing the navigation of others will improve a user's own navigation.

**H2** *Guided footprints will result in significantly higher score than unguided footprints.* The quality of the footprints will also influence the navigation performance.

# 5 Results

Below we review the performance results from our study as well as the subjective ratings and qualitative feedback.

## 5.1 Performance Results

The performance metric for our user study was the total score for each participant at the end of the five-minute search task. Since the score is a discrete variable with a non-normal distribution, we choose to analyze participant score using a Kruskal-Wallis one-way analysis of variance (a non-parametric equivalent of a standard ANOVA). We found a significant effect of no footprints (Group 1) compared to seeing footprints (Groups 2, 3, and 4): $\chi^2(1, N = 100) = 10.028, p = .0015$. Figure 5 shows a bar chart of the average total score for the four different groups (representing the different footprints conditions). Furthermore, there was a significant effect of group on total score: $\chi^2(3, N = 100) = 14.364, p = .0025$ (Kruskal-Wallis test). We analyzed the pairwise differences between groups with a Mann-Whitney-Wilcoxon test using a Bonferroni correction to adjust for multiple comparisons (adjusted p-value .008). As indicated in Figure 5, we found significant differences for Group 1 (no knowledge) compared to Group 3 (hybrid footprints) (Mann-Whitney-Wilcoxon test, $W = 179, n_1 = 25, n_3 = 28, p = .0024$) as well as for Group 1 (no knowledge) compared to Group 4 (guided footprints) (Mann-Whitney-Wilcoxon test, $W = 134.5, n_1 = 25, n_4 = 23, p = .0016$). Both footprint conditions (Group 3 and 4) exhibited superior average score compared to the no footprint condition (Group 1). All other pairwise comparisons were not significant at the .05 level.

## 5.2 Subjective Feedback and Ratings

The subjective comments on the user study and the viewprints (red bounding boxes) were largely positive. One of the participants in Group 2 said that "It was not initially clear what the red squares did but I caught on intuitively after finding a few." This and other comments to the same effect (as well as the performance ratings above) indicate that the study participants were able to understand the meaning of the viewprints even if they received no training or instructions about them. Another participant in Group 3 commented that "The red rectangles give me some sort of indication of where to look. The entire world is a very large area to scope out to find small spots that need to be zoomed in on to even view them." Some participants even realized the relationship between clusters of rectangles and the corresponding score. As a case in point, one participant in Group 4 commented that "If I saw a group of small, dense rectangles then I would go for those because they were probably surrounding a pin that was worth more points."

In the post-test survey, participants rated how much they enjoyed the game on a Likert scale ranging from 0 (0 = high dislike) to 5 (5 = high like). The average rating for the game between 100 participants was 4.45 (s.d. 0.73). Participants also rated for the helpfulness of red rectangles on the scale scale 0 (0 = very misleading) to 5 (5 = very helpful). The average rating among the 75 users who had access to footprints represented as red rectangles was 4.11 (s.d. 1.00). Based on these ratings, it seems most participants enjoyed the game and found the footprints helpful in their treasure hunt.

# 6 Discussion

In summary, our user study yielded the following findings:

- Participants who saw visual footprints (Groups 2, 3, and 4) of past users scored significantly more points than those who did not (Group 1), confirming **H1**; and
- Guided footprints (Groups 3 and 4) from users who themselves had access to footprints did **not** yield significantly higher scores than unguided ones (Group 2), rejecting **H2**.

Below we explain and generalize these results. We then discuss applications of visualizing footprints beyond maps.

## 6.1 Explaining the Results

The overall results from our user study are not surprising and obey our initial intuitions: footprints that visualize past activity do help current users to find common targets in a multiscale geographic space. This is still true in the face of large numbers of footprints, yielding high visual complexity and overplotting. Furthermore, it is worth noting that we did not include any instructions on how to use the viewprints in our crowdsourced user study. Participants were still able to understand and utilize the footprints without any training.

However, there are several surprising facts and design implications to derive from our findings. For one thing, we did **not** find a significant pairwise improvement for Group 2 (unguided footprints) compared to Group 1 (no knowledge); Mann-Whitney-Wilcoxon test, $W = 289.5, n_1 = 25, n_2 = 24, p = .0776$). This means that the initial batch of footprints from the first 25 participants in Group 1 were not very helpful for the 25 participants in Group 2 since their navigation was not guided in the slightest. While this is a fairly intuitive result—emergent collaboration can only arise when the activity of one user can reinforce another—this can also be interpreted as an indication that footprints alone do not yield higher performance. In other words, perhaps we can only claim to partially confirm **H1** since there is no significant improvement in total score between Groups 1 and 2.

We were also surprised to see that Groups 3 and 4 (or at least Group 4, which had no unguided footprints) did not perform better than Group 2. Our intuition was that footprints derived from participants who themselves saw footprints (Groups 3 and 4) would yield better performance than those who saw footprints from participants with no guidance (Group 2). However, by the same argument as above, it is clear that emergent collaboration will only arise when one participant's activity feeds into the next one. Each successive snapshot of interaction data will yield better performance for the next group, but only in an incremental fashion.

## 6.2 Generalizing the Results

How can these results be applied to settings beyond the treasure hunt game used in our crowdsourced user study? Several caveats exist in generalizing our viewprints approach to other applications, including managing scale, visual complexity, and logging policies as well as avoiding echo chamber effects. Before we get into these issues, however, it is worth discussing some concrete applications for this idea.

While visualizing social navigation clearly excels at the type of treasure hunt game employed in the user study, it can be argued that these ideas are of limited utility for general geographic map services. To motivate the utility of this idea, we use the proliferation of the concept of "trending topics" in social media platforms. Twitter has long used trending topics—defined as a word, phrase, or topic that is tagged at a greater rate than other tags—as a method to help Twitter users to understand what is happening in the world. Similarly, Data Wrangling has since 2009 published `http://trendingtopics.org`, a website summarizing current viewing activity on Wikipedia, and Facebook has recently introduced hashtags and trending as part of their service. The method we present in this paper could be a complement to these ideas, giving a visual and tangible representation of collective user activity in geographic spaces beyond the keywords currently used by social media services.

Scale is clearly an important design aspect when visualizing geographic attention [10]. Realistic implementations of this concept would have to perform more aggressive summarization, aggregation, and filtering to maintain scalability. For example, we could use a semantic zooming approach that shows a select few viewprints on the world map level, only loading detailed footprints as the user zooms into a particular area. Similarly, temporal decay can be used to manage scale, particularly when it comes to faithfully capturing currently trending geographic topics, as discussed above.

The sequencing of footprint collection also bears discussion here. Our user study found that successively informed footprints yielded better performance for participants who were able to utilize them. This is a straightforward corollary of the general idea of emergent collaboration [37]: collective performance improves proportionally to how much users can build on the work by past users. This has several implications for how to collect and visualize interaction data for geographic maps. Specifically, while showing the user their own navigation may be confusing, it seems clear that footprints should be visualized as soon as they are collected, perhaps even in real-time for concurrently connected users. Every footprint becomes another step to progressively build on, so distributing fresh events is a key feature.

At the same time, trending topics in general are susceptible to a "tyranny of consensus" or "echo chamber" effect where some concepts in an information space are artificially amplified to the point of isolating its users [17]. For example, if such viewprints attract the geographic attention of its users, as our user study results seem to suggest, this new user activity in turn will feed into new viewprints being created in the same areas that were already indicated using viewprints from previous users. This, in turn, may yield exponential amounts of attention to a select few areas. Dealing with these issues is beyond the scope of this work, but this may be an argument against providing real-time updates of geographic attention across concurrent users.

### 6.3 Beyond Maps: Footprints for Sensemaking

While our work in this paper has been applied to geographic maps, we think these ideas are applicable to many other settings. General sensemaking is becoming increasingly collaborative [23], so harnessing emergent collaboration through social navigation cues is a fruitful area for future research. Existing work such as scented widgets [41] and crowdsourced graph layout [42] have begun to explore this area, but more disruptive designs are possible.

Recording and visualizing the viewport of many users is an indication of the collective attention of the crowd, but we think that there are many more opportunities beyond this. Our continued work in this area is still in very early stages, but we envision crowdsourcing layout, annotation, filter settings, clustering, and other formatting parameters and metadata across multiple users who are interacting with the same visualization tool. The purpose, of course, is to not waste any of the collective efforts of the users working with the tool. In fact, while our work so far has focused on asynchronous collaboration where only the efforts of past users will be fed back to the current user, we are also interested in putting all concurrent users of a tool in contact with each other to share data. Comments and discussions written about a visualization should also be collected and summarized, regardless of source. Finally, while viewport position is an acceptable approximation of the user's attention, accurate attention measurements should use large-scale eye-tracking across all of the users interacting with a visualization.

## 7 Conclusion and Future Work

We have presented results from a crowdsourced user study involving more than 100 participants evaluating the utility of social navigation visualization for an online geographic map service. Our findings indicate that using visualization to externalize the navigational actions of earlier participants helped to significantly improve user performance even with very little instruction. While the gamified task in our study represents only one aspect of geographic map attention, we think that our results show promise for the more generalized idea of visualizing social navigation for web-based visualization, which already includes a client/server ecosystem capable of supporting such emergent collaboration.

Our future research will continue to explore this phenomenon for other domains. On the technical side, we are working to create a general software framework for introducing social navigation visualization to any JavaScript-based visualization. On the applications side, we are looking at integrating social navigation cues in a wide range of tools such as multidimensional, graph, and text visualizations.

## Acknowledgments

This work was partly supported by the U.S. National Science Foundation grant TUES-1123108.

## References

[1] G. D. Abowd and A. J. Dix. Giving undo attention. *Interacting with Computers*, 4(3):317–342, 1992.

[2] Alexa Rank. `http://www.alexa.com/`. Accessed September 2013.

[3] J. Alexander, A. Cockburn, S. Fitchett, C. Gutwin, and S. Greenberg. Revisiting read wear: analysis, design, and evaluation of a footprints scrollbar. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*, pages 1665–1674, 2009.

[4] C. Appert, O. Chapuis, and E. Pietriga. High-precision magnification lenses. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*, pages 273–282, 2010.

[5] Appert, Caroline and Fekete, Jean-Daniel. OrthoZoom scroller: 1D multi-scale navigation. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*, pages 21–30, 2006.

[6] E. A. Bier, M. C. Stone, K. Pier, W. Buxton, and T. DeRose. Toolglass and magic lenses: The see-through interface. In *Computer Graphics*, volume 27, pages 73–80, Aug. 1993.

[7] E. H. Chi. Information seeking can be social. *IEEE Computer*, 42(3):42–46, 2009.

[8] P. Dourish and M. Chalmers. Running out of space: Models of information navigation. In *Proceedings of the Conference on Human Computer Interaction*, 1994.

[9] B. M. Evans and E. H. Chi. An elaborated model of social search. *Information Processing and Management*, 46(6):656–678, 2010.

[10] D. Fisher. Hotmap: Looking at geographic attention. *IEEE Transactions on Visualization and Computer Graphics*, 13(6):1184–1191, 2007.

[11] G. W. Furnas. Generalized fisheye views. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*, pages 16–23, 1986.

[12] G. W. Furnas. A fisheye follow-up: further reflections on focus + context. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*, pages 999–1008, 2006.

[13] G. W. Furnas and B. B. Bederson. Space-scale diagrams: Understanding multiscale interfaces. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*, pages 234–241, 1995.

[14] E. Gams, T. Berka, and S. Reich. The trailTRECer framework - A platform for trail-enabled recommender applications. In *Proceedings of the International Conference on Database and Expert Systems Applications*, volume 2453 of *Lecture Notes in Computer Science*, pages 638–647. Springer, 2002.

[15] E. Gams and S. Reich. Following your colleagues' footprints: navigation support with trails in shared directories. In *Proceedings of the ACM Conference on Hypertext and Hypermedia*, pages 89–90, 2004.

[16] S. Ghani, N. H. Riche, and N. Elmqvist. Dynamic insets for context-aware graph navigation. *Computer Graphics Forum*, 30(3):861–870, 2011.

[17] E. Gilbert, T. Bergstrom, and K. Karahalios. Blogs are echo chambers: Blogs are echo chambers. In *Proceedings of the Hawaiian International Conference on System Sciences*, pages 1–10, 2009.

[18] J. Heer and M. Bostock. Crowdsourcing graphical perception: using mechanical turk to assess visualization design. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*, pages 203–212, 2010.

[19] J. Heer, J. D. Mackinlay, C. Stolte, and M. Agrawala. Graphical histories for visualization: Supporting analysis, communication, and evaluation. *IEEE Transactions on Visualization and Computer Graphics*, 14(6):1189–1196, 2008.

[20] J. Heer, F. B. Viégas, and M. Wattenberg. Voyagers and voyeurs: supporting asynchronous collaborative information visualization. pages 1029–1038, 2007.

[21] W. C. Hill, J. D. Hollan, D. Wroblewski, and T. McCandless. Edit wear and read wear. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*, pages 3–9, 1992.

[22] T. Igarashi and K. Hinckley. Speed-dependent automatic zooming for browsing large documents. In *Proceedings of the ACM Symposium on User Interface Software and Technology*, pages 139–148, 2000.

[23] P. Isenberg, N. Elmqvist, D. Cernea, J. Scholtz, K.-L. Ma, and H. Hagen. Collaborative visualization: Definition, challenges, and research agenda. *Information Visualization*, 10(4):310–326, 2011.

[24] P. Isenberg and D. Fisher. Collaborative brushing and linking for co-located visual analytics of document collections. *Computer Graphics Forum*, 28(3):1031–1038, 2009.

[25] E. W. Ishak and S. Feiner. Content-aware scrolling. In *Proceedings of the ACM Symposium on User Interface Software and Technology*, pages 155–158, 2006.

[26] W. Javed, S. Ghani, and N. Elmqvist. PolyZoom: multiscale and multifocus exploration in 2D visual spaces. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*, pages 287–296, 2012.

[27] S. Jul and G. W. Furnas. Critical zones in desert fog: Aids to multiscale navigation. In *Proceedings of the ACM Symposium on User Interface Software and Technology*, pages 97–106, 1998.

[28] A. Kittur, E. H. Chi, and B. Suh. Crowdsourcing user studies with mechanical turk. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*, pages 453–456, 2008.

[29] D. Kurlander and S. Feiner. Editable graphical histories. In *Proceedings of the IEEE Workshop on Visual Language*, pages 127–134, 1988.

[30] K. Perlin and D. Fox. Pad: An alternative approach to the computer interface. In *Computer Graphics*, pages 57–64, 1993.

[31] E. Pietriga and C. Appert. Sigma Lenses: Focus-context transitions combining space, time and translucence. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*, pages 1343–1352, 2008.

[32] S. Reich and E. Gams. Trailist - focusing on document activity for assisting navigation. In *Proceedings of the ACM Conference on Hypertext and Hypermedia*, pages 29–30, 2001.

[33] A. D. Shaw, J. J. Horton, and D. L. Chen. Designing incentives for inexpert human raters. In *Proceedings of the ACM Conference on Computer Supported Cooperative Work*, pages 275–284, 2011.

[34] G. Shoemaker and C. Gutwin. Supporting multi-point interaction in visual workspaces. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*, pages 999–1008, 2007.

[35] A. Skopik and C. Gutwin. Improving revisitation in fisheye views with visit wear. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*, pages 771–780, 2005.

[36] L. Tauscher and S. Greenberg. Revisitation patterns in world wide web navigation. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*, pages 398–406, 1997.

[37] L. Terveen and W. Hill. Evaluating emergent collaboration on the web. In *Proceedings of ACM Conference on Computer-Supported Cooperative Work*, pages 355–362, 1998.

[38] C. Ware and M. Lewis. The DragMag image magnifier. In *Extended Abstracts of the ACM Conference on Human Factors in Computing Systems*, pages 407–408, 1995.

[39] A. Wexelblat. Communities through time: Using history for social navigation. In *Community Computing and Support Systems*, volume 1519 of *Lecture Notes in Computer Science*, pages 281–298, 1998.

[40] A. Wexelblat and P. Maes. Footprints: History-rich tools for information foraging. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*, pages 270–277, 1999.

[41] W. Willett, J. Heer, and M. Agrawala. Scented widgets: Improving navigation cues with embedded visualizations. *IEEE Transactions on Visualization and Computer Graphics*, 13(6):1129–1136, 2007.

[42] X. Yuan, L. Che, Y. Hu, and X. Zhang. Intelligent graph layout using many users' input. *IEEE Transactions on Visualization and Computer Graphics*, 18(12):2699–2708, 2012.

[43] J. Zhao, D. Wigdor, and R. Balakrishnan. TrailMap: facilitating information seeking in a multi-scale digital map via implicit bookmarking. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*, pages 3009–3018, 2013.

# Figures

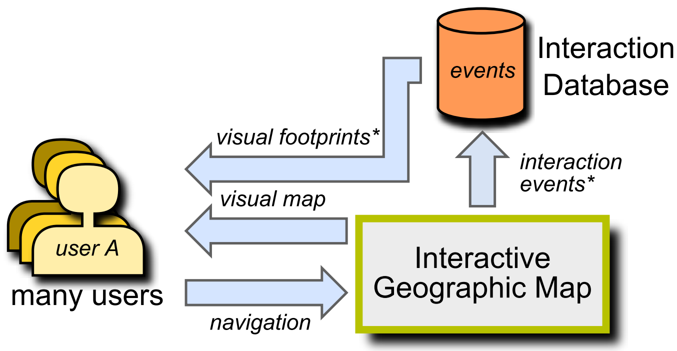


Figure 1: General data collection and visualization process for social navigation in geographic spaces.

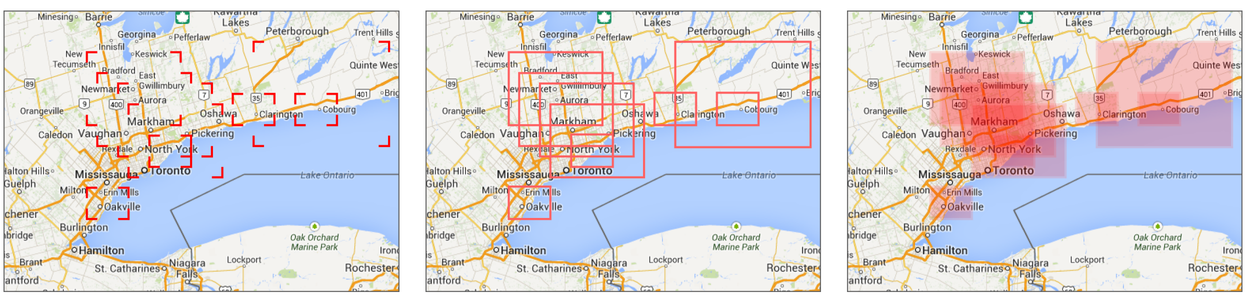

(a) Corners.

(b) Bounding boxes.

(c) Colored areas.

Figure 2: Different visual representations for viewprints.

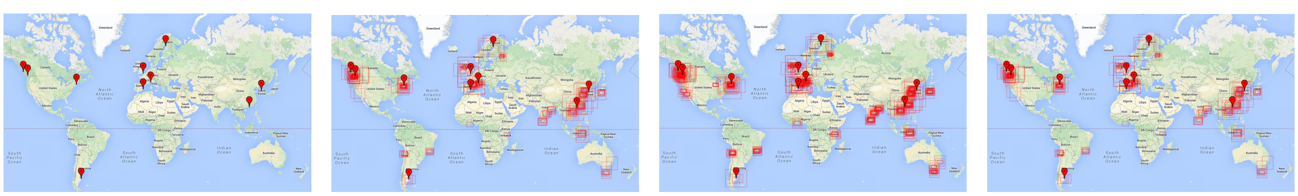

(a) Group 1 - no knowledge.

(b) Group 2 - unguided footprints.

(c) Group 3 - hybrid footprints.

(d) Group 4 - guided footprints.

Figure 3: Maps with social navigation cues visualized as bounding boxes shown to each of the four groups in our user study.


Figure 4: Instruction for scoring in the treasure hunt game.

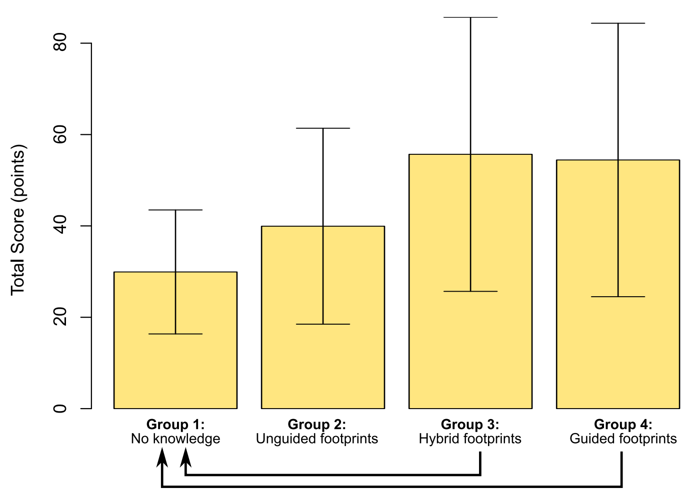


Figure 5: Average total score for participant groups (error bars show standard deviation). The two arrows at the bottom show significant differences ($p < .05$) between groups using a Mann-Whitney-Wilcoxon test with Bonferroni correction for multiple pairwise comparisons.