

How We Watch the City: Popularity and Online Maps

Danyel Fisher

Microsoft Research

1 Microsoft Way, Redmond, WA

danyelf@microsoft.com




Figure 1: A heat map of popularity of tiles over the US at level 12 (approximately 40 meters per pixel scale). The brightest points have on the order of half a million hits, while the dimmest visible points show closer to a thousand hits. Note that usage patterns at this scale seem to follow population. **Inset:** mapping imagery at the level (but not size) represented by one pixel.

ABSTRACT

One way of conceptualizing physical spaces is to look at where people notice, remember, or note them. Computer-assisted methods give us new tools based on implicit, rather than explicit, data about how users have examined and travelled online through cities. “Hotmap” is a tool that visualizes how people have used maps.live.com, an interactive mapping service, looking at what parts of the maps they find most compelling.

Author Keywords

Maps, information visualization, server logs, mashup

In submission: 2007 CHI workshop on Imaging the City

ACM Classification Keywords

H.5.2 User Interfaces; I.3.0 Computer Graphics

INTRODUCTION

To “image the city”, as the workshop title has it, we must draw on some knowledge about the city. There have been a wide variety of ways of collecting this information: Kevin Lynch [2] interviewed people within the city, for example, asking them to sketch and draw parts of the city that they were familiar with. William H Whyte [3] dramatically photographed their patterns of interaction, looking at them from a remove in order to understand their movements within the city.

IMPLICIT METHODOLOGIES FOR DATA COLLECTION

To abstract a step further from Whyte’s method, we might imagine extracting even more implicit information: information on how people *use* the city. These techniques are familiar within HCI circles: techniques like monitoring “read” and “edit wear”, for example [1], form the basis of collaborative filtering and social navigation. In those areas,

systems make metadata visible showing how previous users have accessed or modified data in the past. This basic feedback loop gives users cues to better understand the space they are exploring.

These techniques are less frequently used in social sciences circles, as “implicit data collection”. In the social science context, we might imagine monitoring real use of the city: looking at what sidewalks accumulate the most shoe stains (an effect of heavy use), for example, or which roads have the most accidents (an effect, also, of heavy traffic.)

These methods have the advantage of allowing non-obtrusive data collection, but (therefore) the disadvantage of not knowing precisely what the participant was doing. Worse, these signals can be hard to interpret: are there more accidents at this intersection because it is a poorly-designed intersection, or because it is a popular bit of road? Are the sidewalks more stained because of differences in shoe styles in different parts of town?

This paper does not attempt to resolve these issues, but, rather, will set them aside for the moment. Instead, it merely notes that while these proxies are imperfect, they can still be valuable, and the more we learn about the space, the more usefully we can interpret these results.

In this paper, I present “Hotmap”, which uses implicit information about how users interact with a live mapping service in order to generate heatmaps of where they find “most interesting.”

INTERACTIVE MAPS

Hotmap is based on usage data from Microsoft’s Virtual Earth, one of a new generation of interactive mapping services. Not unlike Google’s and Yahoo’s interactive mapping services, it provides both a main mapping site (<http://maps.live.com>) and a Javascript API for developers to embed the maps in their own site. It provides symbolic maps as well as aerial imagery of the relevant areas, down to (in urban areas) foot-per-pixel quality. Users can pan and zoom around the map smoothly, drilling from world-wide views down to very zoomed-in views in moments. Since its launch in late 2005, it has improved its imagery, maps, and coverage rapidly.

Both the Javascript API and the front page work in similar ways: the browser shows a bounding box and the system downloads appropriate imagery that goes within that box. In the case of Virtual Earth, this is divided into “tiles”: 256 x 256 images that together show the world. Each of these tiles is requested by the code, and provided by a cluster of tile servers, dedicated machines which provide the imagery to users as they pan and zoom around.

In the current system, users’ machines will cache tiles for a fairly short period of time (on the order of days, by default); the system requests only tiles that are within or very near the bounding box. (The software attempts to download a

buffer around the current screen to pre-cache when the user scrolls or moves.)

Thus, should we see a list of tiles that a user has downloaded, we can be fairly certain that they have looked at the place that corresponds to those tiles at least once, although we have little information about whether they returned or how often.

While it is some external application have downloaded individual tiles—testing scripts, for example—the vast majority of tile downloads come from users specifically choosing to look at some part of the world.

Processing and Visualizing Tile Logs

Because this is an online application, we can benefit from the server logs. Tile server logs are standard IIS logs: they store an IP address, a date and time of access, and a URL, as well as several other data points that are not relevant to this discussion.

We aggregate those logs together, producing database records that record only the **number of hits on each day for each tile**. Because we know where each tile is, we can then visualize that usage as a choropleth of tile downloads over location. Note that tiles are arranged by zoom level; rather than trying to combine information from zoomed-in and zoomed-out data, the maps can conceptually be thought of as distinct. Thus, we will refer to the “map at level 10” to mean the server logs describing the usage of the level 10 (approximately 150 meters per pixel, or m/p, scale) tiles.

For this project, we obtained copies of the Virtual Earth tile server logs. In this paper, we present results with eight months of tile server logs, from January of 2006 through August of that year. The sample of tile servers we recorded shows several billion hits across a few hundred million tiles.

We refer to the number of hits at a particular point as its “popularity,” as a more popular tile is one that has been downloaded more. The term is not precise, but stands as a useful concept to how the results can be interpreted.

RESULTS

Choropleths from the system representing cities are shown in figures 2 through 5 at various zoom levels. The color scales, running from white (high) to red (low) are consistent within a zoom level but not between zoom levels, and are scaled logarithmically, so that dim colors have only a tiny fraction of the hits that bright colors do.

Figure 1 shows a map of the United States as a whole at level 12 (40 mpp). Note that bright spots roughly track population: people seem most interested in looking at New York, Los Angeles, and other major cities.

Figure 2 shows a broad representation of the greater Los Angeles area and Orange County at level 17. The bright cluster of spots near the black “+” are downtown Los Angeles. Other bright spots to the west follow the

university district and the city of Santa Monica, and track along the coastline. Below, bright spots are visible at the city of Long Beach. Note that we are seeing places where usage does not match population: the neighborhood of Compton (very near the plus sign) is comparatively uninteresting. The final bright spot worth noting is the city of Anaheim, California, east of Long Beach. Located in the middle of Orange County, Anaheim is interesting because of the presence of Disneyland at its center.

Figure 3 shows the Pacific Northwest at zoom level 13, showing that users follow roads and coastlines and pick out cities at that fairly broad level. Figure 4 zooms in on Las Vegas, showing the popularity of the linear Strip at the map's center.

Last, Figure 5 shows the city of Seattle at zoom level 19. (For reference, figure 6 shows how detailed zoom 19 is). On the west side of Lake Washington, bright spots follow downtown, with its popular Space Needle and sports stadia; duller colors trace out populated neighborhoods. On the east side, the downtown part of the city of Bellevue again is bright.

One additional point of high interest is visible on that map: a small, very bright point on the shore of Lake Washington points out Bill Gates' house.

CONCLUSIONS

We have seen several apparent motivating factors for the use of maps. Collectively, users seem to look at particularly unusual imagery, at roads, borders, and edges; and at homes and neighborhoods (presumably their own). The vast majority of hits are focused on a fairly small area, following population, suggesting that largely, users currently find use of the tools for looking at natural scenery less compelling than they do city imagery. Last, the tools here suggest that visualizing the *use* of online spaces can provide valuable insight into the space.

ACKNOWLEDGMENTS

I thank the members of the Virtual Earth team, who supported this effort, and especially Greg Strodel, who provided the raw data. I also thank the members of the MapCruncher team, who helped build the visualization infrastructure, and Paul Johns and the members of the Community Technologies Group for technical and research help.

REFERENCES

1. Hill, Hollan, Wroblewski & McCandles. "Edit Wear and Read Wear," *Proceedings of CHI'92*. ACM Press, 1992.
2. Lynch, K. *The Image of the City*. Boston: MIT Press. 1960.
3. Whyte, W. H. *City: Rediscovering the Center*. New York: Anchor. 1990.

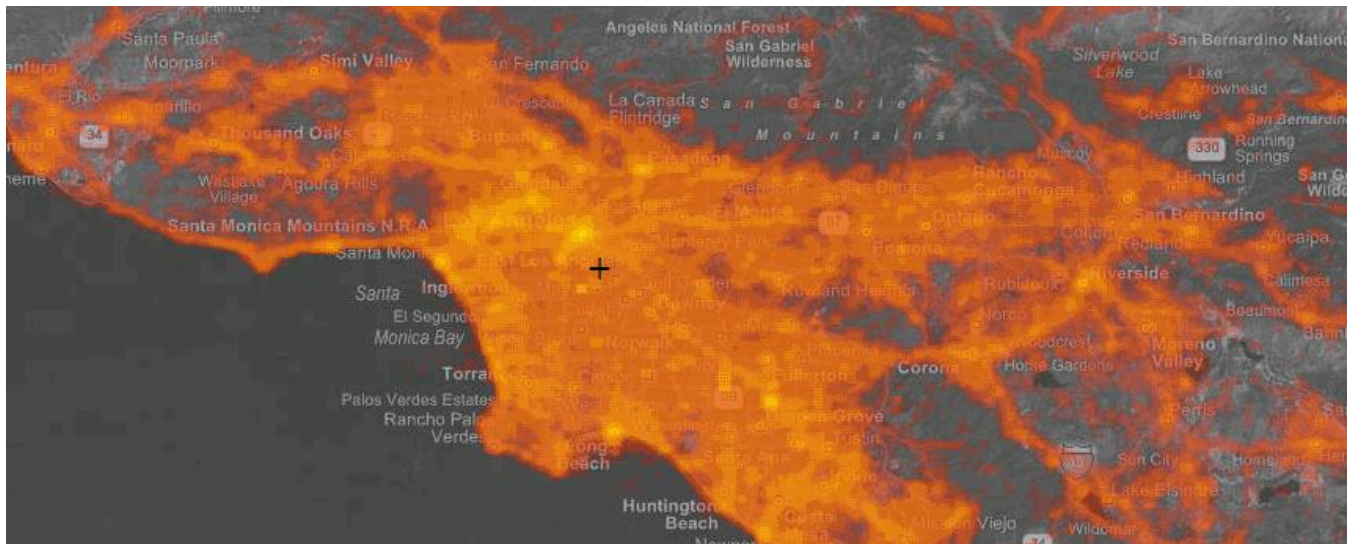


Figure 2: Picturing the City. The greater Los Angeles and Orange County area, at zoom level 17. Note that while the population of the area is very broadly spread out, social focus is far more limited. Interest follows population and coastlines, and points of interest, including Disneyland.




Figure 3: The US Pacific Northwest, at zoom level 13 (17 meters per pixel). Interest follows populations, as well as roads and shorelines.




Figure 4: The greater Las Vegas area, at zoom level 18 (0.6 mpp).




Figure 5: Seattle downtown and suburbs at zoom level 19 (0.3 mpp). Interest follows local popularity. See text for discussion.



Figure 6: Sample imagery at zoom level 19, for reference.