

The Impact of Hotmap

Danyel Fisher

Microsoft Research

ABSTRACT

In 2006, we released Hotmap internally at Microsoft. Hotmap illuminates how users have used Virtual Earth by visualizing the tiles that have been looked at by users. While the technical aspects and basic usage of Hotmap have been previously discussed [1][2], this paper discusses the next steps: the impact of Hotmap, the ways that Hotmap was used over the following years, and the fate of the tool three years later.

KEYWORDS: Geographical visualization, heatmaps.

INDEX TERMS: H.5.1 [Multimedia Information Systems]: Hypertext Navigation and Maps

1 INTRODUCTION

Understanding how users interact with online maps can allow map designers to figure out how to extend and modify their mapping tools. While many web sites can be monitored with traditional log management tools, Hotmap is a visualization of users' use of Virtual Earth (VE), a mapping system which is visible (among other ways) as Microsoft's Live Maps (<http://maps.live.com>). Using Hotmap, a user can see where web visitors have observed the site.

In 2006, the author had a series of discussions with members of the Virtual Earth development team, which is responsible for the Live Maps website and content. These discussions helped clarify that monitoring of Live Maps was based on techniques that, while state of the art for other web sites, were not applicable to an AJAX-based mapping solution. For example, it recorded users' usage of imagery to the nearest 0.1 degree, and did not record how tightly zoomed-in the user was. After some negotiation, several members of the team were able to release their server logs in order to allow a prototype to be developed.

Hotmap was released internally at Microsoft in 2006 as a research prototype; its use was monitored as it was developed further. A public edition of Hotmap (<http://hotmap.msresearch.us>) was released in 2007; the authors have collected feedback on that tool since.

2 VISUALIZATION DESIGN

The visualization is described in some detail in the previous publications [1][2]. The visualization is based on internal server logs from the Microsoft Live Maps site. Every time a user looks at a tile of Virtual Earth, a log file records the tile request. Tile requests thus contain their location (as an x,y coordinate), a zoom level, and the style of the tile (road, aerial, hybrid). When a user first looks at the map, their browser loads between ten and thirty tiles.

In the development of Hotmap, we did not build a regular process for data sampling; rather, during the prototype phase, we

improvised data delivery periodically. Hotmap is based on a several-hundred GB sample of Virtual Earth traffic (adding to less than 1% of total traffic); as even Virtual Earth was still in development, constant operational changes meant that the sample never quite represented the same subset of data.

Hotmap is built as a mashup over Virtual Earth itself; as such, it is able to take advantage of the pan and zoom features built into Virtual Earth.

3 INSIGHTS AND PROCESS

In previous publications, we have discussed some of the insights that our users found from Hotmap. In this section, we review some of those conclusions, and then add several additional scenarios that come out of later work.

Hotmap was used to allow the Virtual Earth team to learn what existing imagery VE users liked to look at. This is not a mere question of entertainment: aerial imagery is expensive to collect, as it must be custom-photographed from scheduled airplane flyovers. These flyovers cost per mile; thus, it makes sense to photograph the smallest useful area of a city. Hotmap was able to demonstrate that several aspects of cities were frequently of interest: users' aggregate attention was focused on universities, downtown cores, scenic sites, and airports. This last is important because airports are often located out of town, and so can be more expensive to photograph. In contrast, users seemed comparatively less interested in suburbs and areas further out of town. This allowed the team to budget their later imagery.

Hotmap also allowed the Virtual Earth team to decide where to purchase future imagery. By looking where users had zoomed in on the map, hoping to see detailed imagery of their favorite spots, the Virtual Earth team could gauge relative demand. VE users also tried to zoom in on beaches and coastlines, inspiring the VE team to purchase additional beach imagery even when those spots were out of city cores. This added up to a substantial cost optimization, as the team was much more able to figure out where to spend their fixed budget.



Figure 1. Hotmap, showing all trails. A pixel is red if at least one user has looked at it. Note the frequent horizontal and vertical lines from users holding down scroll buttons.

danyelf@microsoft.com

Microsoft Research

Hotmap acted as a way to catch errors. A surprising mass of users looking at the Atlantic Ocean near the cape of Africa turned out to be bad calls to the Virtual Earth API (visible in figure 2f). The point that was impacted turned out to be (0,0) in the latitude/longitude coordinate system; the team was able to document the failure mode more clearly to developers.

Hotmap also allowed the VE team to understand how users were moving through the Hotmap site. In Figure 1, we see all tiles that users downloaded even once. The horizontal and vertical streaks come from individual users continually scrolling out across the ocean. This behavior was surprising—Virtual Earth has no data there—but was a useful reminder of users’ interaction with the site.

In Figure 2, Hotmap was used to track the diffusion of Virtual Earth. As additional imagery was added, users began to look at more of the world: the initial release had only imagery of Great Britain and the United States; as such, only those areas were observed by users (Figure 2a-b). Almost immediately upon the release of extensive European imagery, people began to look at those locales, almost as often as they did at American locales (Figure 2f). The marketing team and planning team were both able to use this information to drive their plans for future rollouts and to understand the impact of promotions and advertising on the platform.

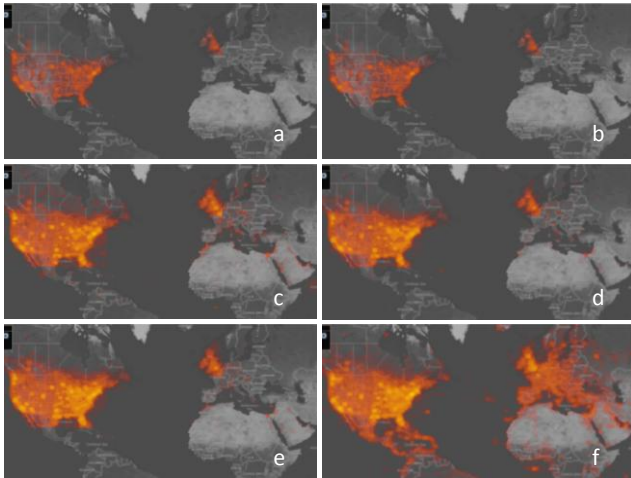


Figure 2. Hotmap for the first six months of Virtual Earth rollout.

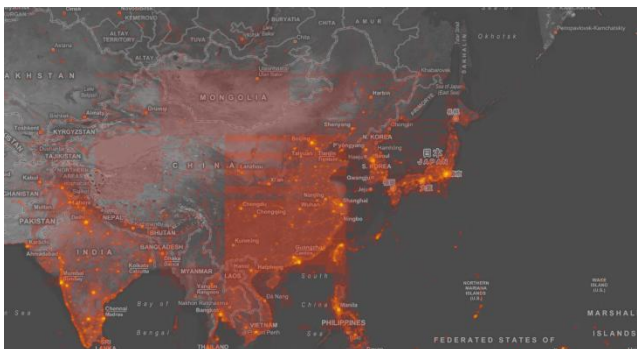


Figure 3. Hotmap in China. Note the faint highlight over greater China, and the brighter colors over the coastal area, suggesting repeated scraping of the data.

In Figure 3, a scraping attack is visible on Hotmap: the data over mainland China is being systematically collected. (This is visible as the uniform highlight across the top of the map). This sort of attack had not been expected by the Virtual Earth planners. They were now aware it was possible, and security measures were designed into later versions.

There were many different parts of the Virtual Earth team that were informed by Hotmap. The developer API, imagery planning, marketing, and usability groups all used Hotmap to understand better how VE was being used and developed.

After the public version of Hotmap was released, several outside groups sent email discussing how they used Hotmap. One notable user was from a geographical imagery vendor. This vendor was attempting to convince her customers that a smaller, targeted imagery collection might be as useful—and far cheaper—than a broader collection: after all, there were many places where users simply did not look. Hotmap was a useful demonstration: looking at the brightest spots helped her clients to know where to buy, and to know that they could afford to buy in a more precise pattern without losing the effect.

3.1 The Impact of Hotmap

Hotmap substantially changed the way that the Virtual Earth team refers to their system and data. Before Hotmap, traditional web monitoring tools allowed the team to list the tiles that were most downloaded, and to instrument additional data, but to do so at a fairly rudimentary level.

Many members of the Virtual Earth team did not use Hotmap frequently: they would find a few important observations, take screenshots to make their case, and go on to other things. None the less, those first screenshots were very relevant, as it allowed the teams to understand their system better.

Hotmap also became, for some members of the team, a sort of ‘mascot’ for the application. A Hotmap screen was proof that their website was being used, with real users taking advantage of the tools and images that they had worked so hard to create.

3.2 The Afterlife of Hotmap

It is worth observing that many of the insights do not really require continuous data updates. A one-time view of the data inspired many groups to think in new ways about their system: the additional insights to be gained from a data refresh were small compared to the initial value of the data. Product groups continued to refer to Hotmap over several years, but there was insufficient funding available to create a full-fledged internal application for the tool.

In early 2009, a product group became concerned that their view of the data was obsolete: Hotmap mainly referred to the use of the VE in its first six months, and VE is now a far more mature product, with more thorough coverage. That team has begun to adapt and update Hotmap’s data coverage, visualizations, and data.

REFERENCES

- [1] D. Fisher, "Hotmap: Looking at Geographic Attention," *IEEE Transactions on Visualization and Computer Graphics*, vol. 13, no. 6, pp. 1184-1191, November/December, 2007.
- [2] D. Fisher, "How We Watch the City: Popularity and Online Maps". Presented at the Workshop on Imaging the City at *CHI 2007*, May 2007.
<http://research.microsoft.com/apps/pubs/default.aspx?id=69421>