

Mapping for the Masses

Accessing Web 2.0 Through Crowdsourcing

Andrew Hudson-Smith

Michael Batty

Andrew Crooks

Richard Milton

Centre for Advanced Spatial Analysis, University College London

The authors describe how we are harnessing the power of web 2.0 technologies to create new approaches to collecting, mapping, and sharing geocoded data. The authors begin with GMapCreator that lets users fashion new maps using Google Maps as a base. The authors then describe MapTube that enables users to archive maps and demonstrate how it can be used in a variety of contexts to share map information, to put existing maps into a form that can be shared, and to create new maps from the bottom-up using a combination of crowdcasting, crowdsourcing, and traditional broadcasting. The authors conclude by arguing that such tools are helping to define a neogeography that is essentially “mapping for the masses,” while noting that there are many issues of quality, accuracy, copyright, and trust that will influence the impact of these tools on map-based communication.

Keywords: *network economies; web-based services; map mashups; crowdsourcing; crowdcasting; online GIS*

Crowdsourcing (Howe, 2008) is the term often used for methods of data creation, where large groups of users who are not organized centrally generate content that is shared. In this article, we explore how crowdsourcing and the related concept of crowdcasting, facilitated by developments in web 2.0 technologies, is leading to the phenomenon of “volunteered geographic information” which, in turn, is revolutionizing the way in which maps are created and used (Goodchild, 2007). This leads us into a foray into contemporary ideas about neogeography (Turner, 2006), which might be defined as “geography for the masses” or even as mapping for the masses. In telling the story of how maps can be created and disseminated as web-based services, we will draw on various applications that we have developed, including GMapCreator, MapTube, and the London Profiler. Networks spin off from all these applications although to date there are very few applications that have specifically extracted networks from map usage and map users.

Authors' Note: Please address correspondence to Andrew Hudson-Smith, Centre for Advanced Spatial Analysis, University College London, London, United Kingdom; e-mail: asmith@geog.ucl.ac.uk

Crowdsourcing and Crowdcasting

The notion that there might be value in harvesting the knowledge of individuals is based on the observation that, although a large number of individual estimates may be incorrect, their average can be a match for expert judgment. Judiciously handled, randomly sampling the opinions or calculations of a large number of users might lead to data and information that is surprisingly accurate and that, in some cases, cannot be recorded in any other way (Surowiecki, 2004).

Essentially, this is called crowdsourcing when it is applied to the creation of content, which is then accessible and sharable as a web-based service. In one sense, enabling users to generate their own content is not a new idea, for it forms the basis of much group psychology and problem solving. What is new is the notion that this kind of group dynamics is uncoordinated and voluntary. In many contexts, the data simply accumulate and until interaction occurs between users, the only value added is because of the extent to which the data represent some phenomena of interest. The law of large numbers dominates in this instance. Where interaction begins to take place between users and where new derivative data come from these interactions and social networks emerge, then the value added depends on the wisdom of the crowd, albeit engineered in countless ways by the designers and managers of the web-based services involved.

Typical web 2.0 social networking sites like Facebook and communal databases like Flickr draw their users by word of mouth, using a whole range of media that potential users respond to, but there is little of a direct push factor from the owners or managers or designers of these sites. Where there is a more direct push digitally using online systems, this is referred to as crowdcasting, which is the union of broadcasting and crowdsourcing, with the potential crowd being “pulled” in as a direct result of the “push.” The incentives for developing such systems clearly depend on the value that the crowd sees in interacting in this way. Some crowdcasting systems actually introduce incentives for users to key in their own data by offering rewards, even “prizes,” often of monetary value (e.g., <http://www.threadless.com/> offers money for user-designed T-shirts) in much the same way social psychologists entice participants into laboratory experiments. None of our examples below involve such reward structures and, in fact, most of these depend on crowdsourcing, rather than crowdcasting. However, where we have invoked crowdcasting, the push factor is simply one of making potential users aware of our desire for them to participate, often through other web sites and related digital media such as e-mail.

By its very nature, it is impossible to discuss web-based services without invoking the notion that everything is potentially linked to everything else. This means that crowdsourcing or crowdcasting, tagging data, deriving patterns, and correlations such as social networks can all take place in different web environments, which are linked to one another. For example, one web site might broadcast to the world at large that users are required to key in data and/or respond to requests that are captured on another web site which, in turn, create data on a third. We do not yet have a good catalogue of such web-based services that would enable us to generate a useful typology and certainly it is not yet possible to create a series of ideal types. Thus, in the rest of this article, we will present examples based on our experience of developing and using such services with respect to geographic information, namely maps, which we take as instructive to the development of this field. However,

first, we need to present a brief history of how mapping for the masses has emerged as part of these web 2.0 technologies.

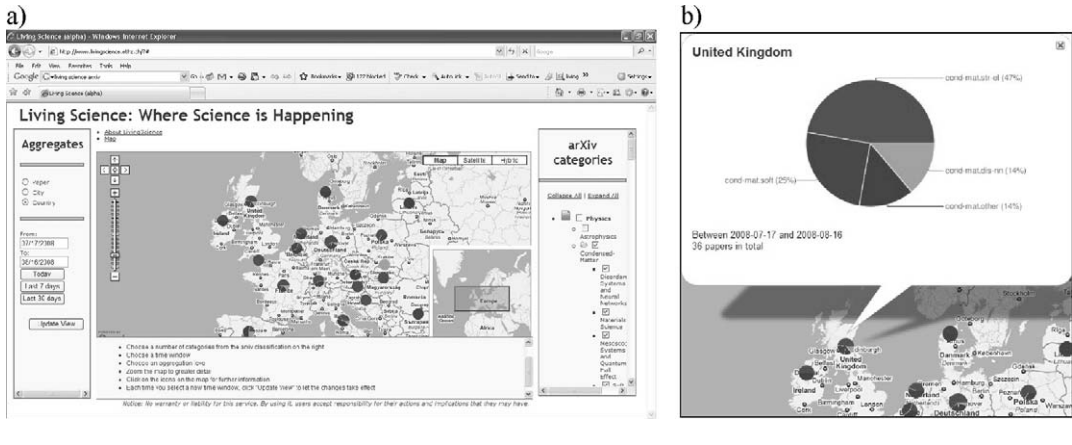
Web 2.0, Geographic Information, and Map Hacks

Web sites such as Multimap (<http://www.multimap.com>; and its many variants) have been around almost since the web itself began, but these represent expertly organized map data that is then accessed by lay users. It is organized in such a way that the user can query the map data to get back a relevant segment related to their locational request. Although most such web sites provide highly interactive user interfaces, they are largely passive in terms of the choice of data they allow users to have access to and their lack of support for collaboration between users. An early but still relevant example of the addition of specific functionality to such mapping sites is UpMyStreet (<http://www.upmystreet.com>), which contains information about property, crime, education, and so on for local neighborhoods, tagged, of course, to prespecified base maps (which in this case are from Multimap, Microsoft Virtual Earth).

Tagging information to maps became more popular than expected early in the development of the web. Many assumed that most users could not read a map and it was a surprise that mapping web sites took off. Map bases of more generic form were introduced by Google in their Google Maps in early 2005, where immediately users could add some content, albeit limited, which since then has increased in extent in various forms in particular through Google MyMaps, which lets users add polylines and other content such as embedded logos of locational features. Google Ride Finder and Google Transit enable users to plan trip services using taxis, public transport, and so on, where such data have been added and now Google Street View provides 360° panoramic street-level views of various U.S. cities. Google have made available an application program interface (API) that lets users embed their maps into third-party applications, and we make use of this facility extensively in our demonstrations below. This has enabled users to easily create mashups, that is, to combine two or more pieces of complementary web-based content or functionality to create a new web application. UpMyStreet, for example, is such a mashup, while Multimap is not. London Profiler, which we sketch below, uses the Google Maps API (see Haklay, Singleton, and Parker (2008) for a further review of web mapping applications and web 2.0 technologies).

Google Map mashups are now appearing everywhere as it is simple to include a Google Map into a web page. What is harder is to use the map as a resource to direct queries of a database that is locationally tagged, other than basic zoom, pan, and location queries, which are intrinsic to the map product itself. As the functionality increases in sophistication, the potential set of users narrows to more specialist and professional usage. A recent example is Living Science which is a Google Map mashup (<http://www.livingscience.ethz.ch>). This site enables a user to search a database of scientific papers deposited on the open archive arXiv (<http://arxiv.org>)¹ and then maps the data by city or by country. An example of its use is shown in Figure 1a, where we have queried the number of papers written in condensed matter physics in the last year from July 17, 2007 to July 16, 2008 in all countries of the world that feature within Google Maps. There are eight sub-categories that are displayed in terms of their proportions on pie charts, which are located

Figure 1
The Living Science Mashup: Relatively Sophisticated Functionality but Limited User Interaction



in each country. These pie charts are not mapped in proportion to their size, which is a limitation imposed on the mashup, but if one clicks on the map, then details about the number of papers emerges (see Figure 1b).

There are many other functions that could be added to such a mashup, such as providing data analysis facilities. For example, it is well known that size distributions with respect to numbers of papers produced scale with the size of the country or the city with respect to any disaggregation of the data into different categories. Thus, once a user selects a category and location type, then the scaling graph could be produced on demand (Carvalho & Batty, 2006). There are other properties of these spatial distributions that could be produced and this would represent a routine set of extensions.

In contrast to web environments that take prespecified maps from proprietary sources (even if the map data is copyright free), web 2.0 technologies are enabling users to create the data themselves. The notion of users creating the data themselves is fraught with controversy, but some data that are generic to all users such as our experience of the world around us can benefit from this approach. In the spirit of crowdcasting and the wisdom imparted from large numbers of individuals creating and improving content, there are now several sites where users can contribute their efforts to the creation of maps. Just as Wikipedia is a resource whose contents are continuously improved by individual contributions, so map systems created from the "bottom-up" evolve. Why users do this is a mixture of pride, the notion that if something is wrong then it can be improved, but most of all, it is based on the idea that the world around us should not be "copyrighted" as many mapping systems created by specialists are. OpenStreetMap (OSM) is a good example of user created content in mapping, broader than other alternatives such as WikiMapia (which is based on Google Maps but with much less user functionality to add content).

Users essentially put data in OSM using GeoPositioning Systems (GPS), which are supplemented by various text, voice records, digital pictures, and so on. There are rather strict

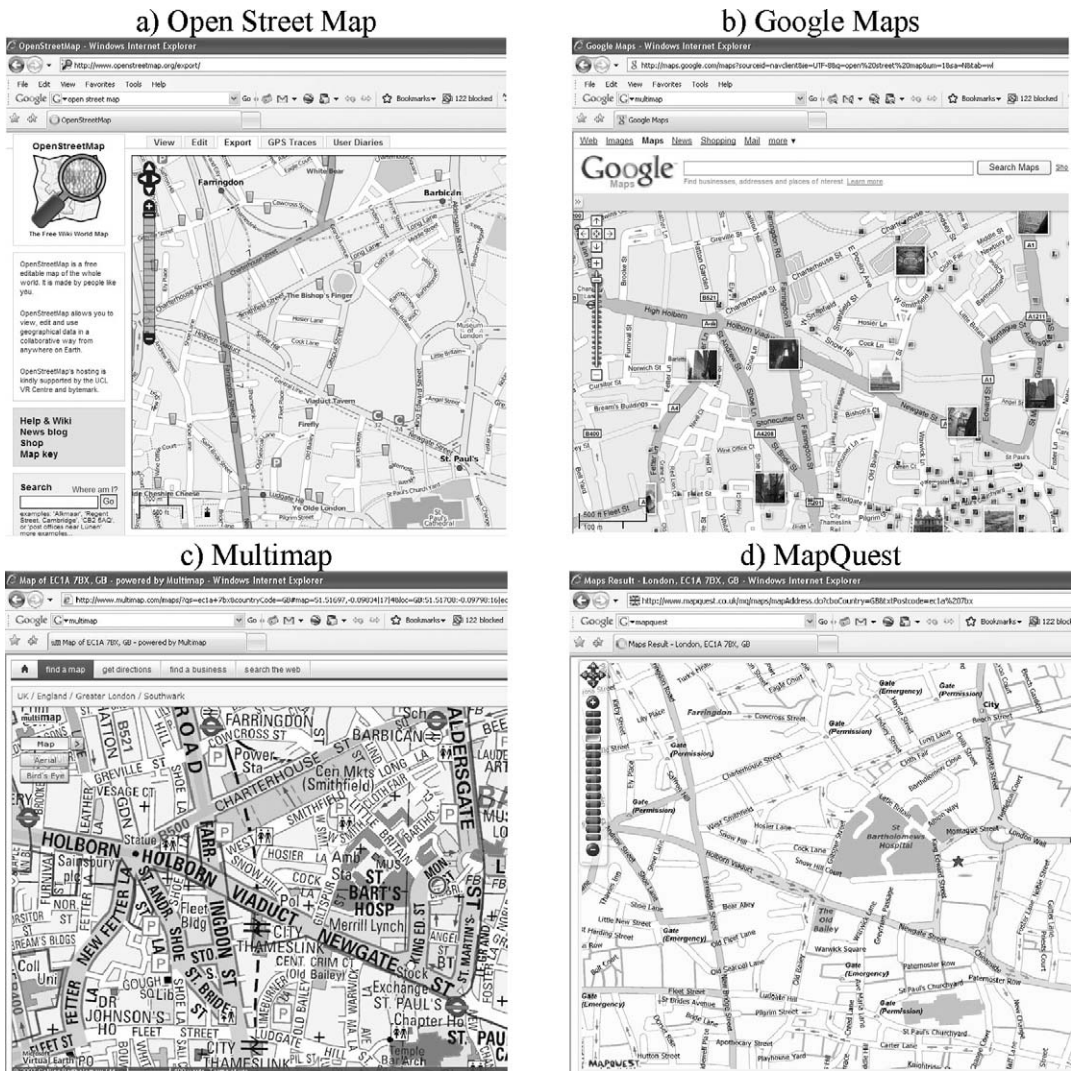
rules involved in creating content so that some consistency is maintained, unlike perhaps, Wikipedia where any user can add an entry. Mapping, notwithstanding its immediately usable form, is considerably more technical in content than might appear at first sight and thus rules are required for any user to create content that is then added to the evolving map. Nevertheless, OSM like Wikipedia, is a process of evolving a good product, not a complete product in itself because there is no end goal in sight as to what constitutes the best map (or the best entry in the case of Wikipedia). Users must register to input data and upload GPS traces but the traces and diaries of those producing content can be viewed in the public domain, which is one of the criteria that OSM insist users as creators must agree to. In fact, OSM also uses whatever data it can get on maps that official or commercial agencies have created and which do not have any copyright attached. Much of the basic data produced by users are checked against map data produced by government agencies where this is in the public domain.² In Figure 2, we show the basic map data from OSM, Google Maps, Multi-map, and MapQuest for an area of central London, which is the center of the postcode indexing EC1A.

MapTube: The Promise of Neogeography

The problem in generating maps is that although nearly everyone understands a map in its generic form, the way a map is configured in digital terms is nontrivial for it involves a spatial database and some knowledge of geocoding. Thus, for users to add value to maps—for example, by overlaying them—or convert maps into a form that might exist on the web as pictures (as portrayed in all the images presented so far in this article), this involves some knowledge of the difference between vector and raster maps. A raster map is a map whose features and other content are captured and/or displayed on a fine grid—a raster—which often approximates to the pixilation of a computer screen. In vector maps, the data defining the map features are identified by polylines and other geometric attributes. For speed of processing, many map systems hold their data as a raster or a grid of “tiles” even though it may have been created in vector form. Google Maps, for example, load a grid of tiles from these data, whereas other maps load the data as a vector file that then becomes a raster for display on the screen. The latter is usually much more demanding of computer storage and processing time than the former.

To facilitate the creation of maps from user-sourced data, we have constructed a means of converting any vector-based map into a raster map, which can be overlaid on Google Maps. We call this software GMapCreator³ and what it essentially does is to take a vector file that is composed of boundaries and attributes tagged to polylines or polygons (as well as point data), first converting these data into the map projection used by Google Maps and then rasterizing the data as set of preconfigured tiles that can be overlaid on the Google Maps base. As Google Maps has 16 layers of zoom, the user must then choose the right range of zoom, trading off the number of tiles to be created against the map scale and zoom required. Moreover, the colors of the attribute range also need to be chosen while the original projection of the map data needs to be known. GMapCreator creates a Google Map layer of the output and generates a web page into which the Google Map and its layer are inserted. The data format for the vector map data is the Environmental Systems Research

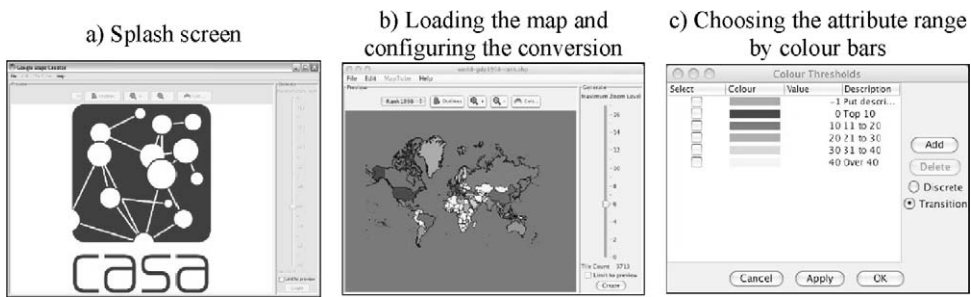
Figure 2
A Comparison of Different Public Domain Map Systems With Respect to Detail and Added Content



Institute (ESRI) proprietary shape file format, but there are numerous converters freely available in the public domain to convert other vector files to this format. We believe this software can be used to handle virtually any data format indirectly.

The details of how GMapCreator operates are presented in the MapTube web environment, which we describe below. The crucial issue is judging the level of resolution needed for the map display because, if all 16 levels are chosen (the slider in Figure 3b controls this operation), then this may be far too detailed for the data in question and also the number of tiles created might be so large that they take far too long to create. To an extent, this is a

Figure 3
GMapCreator: Creating a Google Map Layer From a Shape File

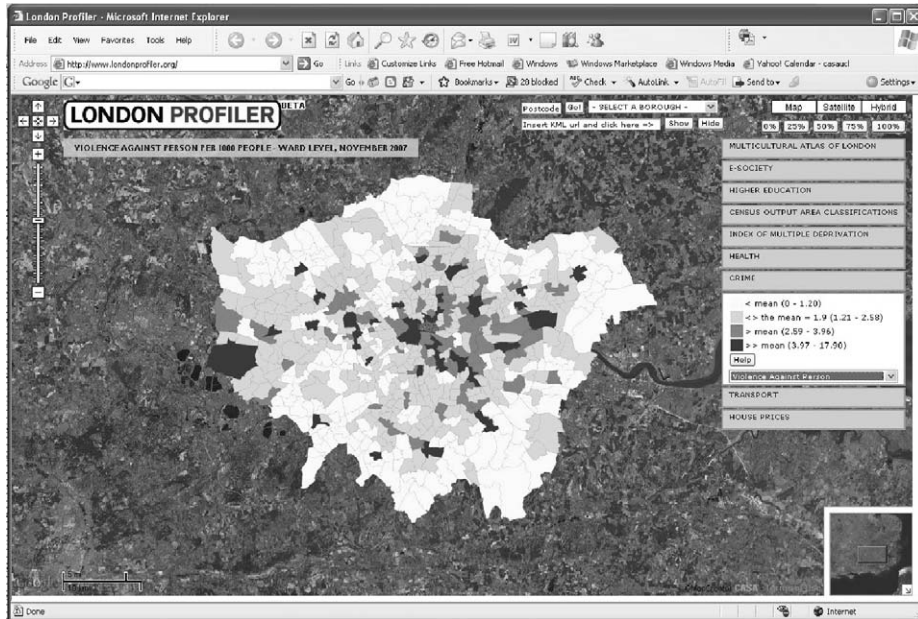


matter of gaining familiarity with the map interface and experimenting with the map data in question which is being converted.

A good example of the use of GMapCreator to create maps is the London Profiler (Gibin, Singleton, Mateos, & Longley, 2008),⁴ which contains a series of map layers that present data on various geodemographic attributes of the population at ward level in Greater London (see Figure 4). Information in the London Profiler consists of map layers taken from Census, Health Trust, and marketing company data, ranging from incidence of multiple deprivation to house prices. It is only possible to array one data set on the Google Map at a time, although one can flip between layers. It is possible to include map layers that are available as KML files (the Google Map format) from other sites by typing in the web address (URL) into the Profiler as any data that have been converted to KML format by any means can be accessed from within Google Maps. However, this environment is still “top-down,” with data being fashioned centrally and, in this sense, the site is not one where users can create their own content. They simply manipulate the content provided by ourselves. MapTube, on the other hand, is much more in the spirit of crowdsourcing, although limited by a modest learning curve with respect to using GMapCreator to produce maps, which can then be stored in its archive or accessed via its user interface.

MapTube (<http://www.maptube.org>) is designed as an environment where users can create maps, view them, and engage in simple manipulation using the time-honored method of map overlay. In fact, GMapCreator is now part of MapTube in that those who use the software to create their own maps are exhorted to share their maps through the MapTube site. The site is not the usual type of archive because of difficulties over copyright of source data. As users can use GMapCreator to create map layers of any geographical data, this might lead to the infringement of some third-party’s copyright. The site protects against copyright breaches by asking users not to share their map but their web address (URL) where they place the Google Map. In this sense, MapTube is more like Napster than YouTube in that the site stores pointers to other sites where the user-generated maps have been placed. Many of the maps archived on MapTube so far are our own and we exercise appropriate restrictions with respect to copyright, but if a user breaks copyright, then all we have is a link to their site, not to the map data itself. This is feature of the user interface that is complicated but necessary. Arguably, it is too

Figure 4
The London Profiler (www.londonprofiler.org)



Note: Showing the distribution of violence against the person measured as the ratio of reported crimes per 1,000 population; note that high values can indicate places of high employment and low population where crimes are committed because of the concentration of employment, not the location of the resident population as in the west at Heathrow airport and in the center in the borough of Westminster.

complex in that maps are not like CDs, and the notion of sharing where they are located rather than the data itself is problematic.

MapTube also enables users to create as many overlays as possible, although after three or four, their intersections can become unmanageable and of course unreadable. Currently, we are adding new functionality to make this feature more intelligible; but MapTube goes well beyond the London Profiler in that the overlay capability allows differential shading and related manipulation.

MapTube, like the London Profiler, allows users to link into other sites with KML files and this naturally extends to the MyMaps facility that is a local customization of Google Maps. Figure 5 shows how we can relate a series of local maps created in MyMaps to wider, more professionally created data. In the summer of 2008 (the time of writing), serious knife crime related to teenage murders has become a significant public issue in London. From casual reporting in the daily press, it is easy to extract the locations of these crimes and to compare these against national assaults in England, which had been uploaded to MapTube by another user. Data on such incidents are hard to come by from public sources as the police do not provide open access to incident locations, despite it perhaps being in the public interest. As such it is left to the “public at large” to fill in the gaps. In the case of teen killings, a user with the pseudonym *MapMan* has created a map of teenagers murdered in

Figure 5
MapTube and Google MyMaps: Local Reporting of Murders Correlated With the National Pattern of Assaults

a) Murders in Greater London 2007-8



b) Distribution of assaults in England



c) Overlays of murders onto assaults



d) An overlay of regeneration areas



London since 2007. Created using MyMaps, the list has been compiled via various web sites (e.g., <http://www.capitalradio.co.uk/article.asp?id=532062>) with street names identified in related press articles. Actual positions within the streets are not likely to be accurate but the street names themselves are. Note the map relates to all murders, not just knife-related incidents and this is shown in Figure 5a. Using MapTube, the map can be overlaid with other data sets, such as the map uploaded detailing assault using a knife or sharp objects extracted from all 2007 hospital admissions, which are classified with code ICD-10 X99. The map excludes all codes that may indicate accidental injury ICD10—W25, W26, self inflicted injury ICD10—X78, and undetermined intent ICD10—Y28. This is shown in Figure 5b.

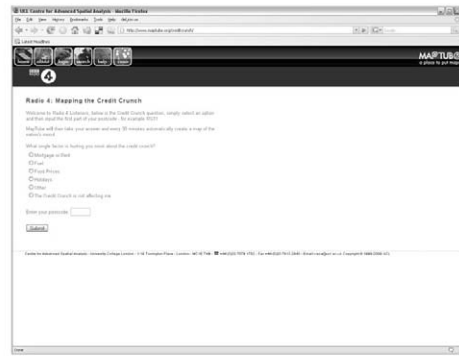
Figures are standardized by age per 100,000 population, while the actual counts were excluded from the map because of disclosure issues involving low numbers. By overlaying the two maps, one begins to get a picture of the extent of knife crime and the number of murders in London as we show in Figure 5c. Each link is clickable for more information. If one then uses MapTube to add in Regeneration Areas within London, a clear pattern between teenage murders and deprivation emerges. MapTube enables this layer to also be viewed against other indicators such as ethnic population density, barriers to housing,

Figure 6
MapTube and Crowdsourcing: The Credit Crunch Mood Map

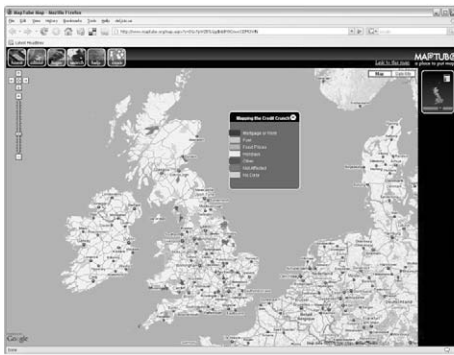
a) Radio 4 iPM web page on the mood map for the Credit Crunch



b) The web questionnaire



c) Early response distribution



d) After 23,000 user responses



or any of the other 47 maps relating to London in the archive. The visual correlations are of significance, thus illustrating how “professional” and “amateur” data⁵ can add real value to the sort of insights that these comparisons enable. This example illustrates how MapTube enables users to add value to maps in ways that their original creators cannot anticipate, and in this sense, it is driven from the bottom-up by whatever its users consider to be significant.

MapTube also has the capability for much more effective crowdsourcing and, indeed, crowdcasting. The way this works is as follows. Using broadcast media such as TV or radio, people are invited to log onto a web site, respond to a series of questions, and specify their location using a postcode. The data are then uploaded to MapTube, which converts it every 30 min, say, to a form where it can be displayed as a map.

This process was first used to create a mood map of the current economic recession in the United Kingdom, popularly known as the “Credit Crunch.” Working in conjunction with the BBC Radio 4 iPM show and BBC TV NewsNight (see Figure 6a), we created a short survey.

People were asked to choose from one of six options: mortgage or rent, fuel, food prices, holidays, other, or “the credit crunch is not affecting me,” which they considered the singly most significant factor hurting their personal finances and to enter the first part of their postcode (postcode sector, e.g., EC1A) so their responses could be geotagged (see Figure 6b).

No personal information was collected and participants were reassured that their actual locations could not be identified. This was enabled through the use of postcode sector rather than the postcode unit or building address, therefore preserving data confidentiality. Each response updated the database element of the underlying shapefile with GMapCreator running in the background to create a new map that was subsequently updated on MapTube as shown in Figure 6c. Over time, as more participants entered information, the map went from blank to varying shades reflecting the responses with respect to what people were worried the most about in the postcode sector (see Figure 6d).

The potential of this approach for gathering and disseminating spatially coded information is enormous for social science research. For example, it could easily be used to gather other information such as fear of household burglary, the quality of primary school education, access to local health facilities, and so on. Mapping the Credit Crunch represents one of the first near real-time geographic surveys of a nation’s perceptions about a specific issue. As each response includes a time stamp, it allows the nation’s mood to be visualized in both time and space. In excess of 23,000 people took part in the survey over a 3-week period, creating a unique and interesting data set that is very much of its time. The Credit Crunch Map has since led to several other surveys, including BBC Look East, the nightly news program for East Anglia in association with BBC local radio, using the system to create a map of people’s perceptions of antisocial behavior. Using a similar data gathering technique, viewers of BBC Look East were asked to answer a series of questions on their views on antisocial behavior at a postcode district level. The survey at the time of writing is ongoing with 6,000 plus respondents to date. Figure 7 illustrates its use as part of a news segment on BBC Look East.

Discussion

In a “pre-Google” world, “mapping for the masses” would have been impossible because the license to use the base map and aerial imagery would have been prohibitively costly and the behind-the-scenes GIS would undoubtedly have been slow and cumbersome to use. It is the essence of neogeography:⁶ free, easy to use, and yet potentially very powerful in terms of its impact on geographical information sciences, social sciences, and its capacity to encourage a new partnership between professionals, lay practitioners, and the wider public.

However, a note of caution is in order. First, there is the issue of access to the technology. Analysis of the survey we carried out provides us with insights into the respondents. Demographically, these are mainly upper and middle income viewers or listeners, defined not only by the audience for Radio 4 and the BBC news programs but also by the differential take up of technological by different social groups (Longley & Singleton, in press). Second, it is important that we see the technologies for what they are. For example, Keen (2007) argues that web 2.0 should be used to encourage innovation and open

Figure 7
Crowdcasting: Mapping Antisocial Behavior in East Anglia Using MapTube on BBC TV Look East



communication and progress while simultaneously preserving professional standards of truth, decency, and creativity (rather than plagiarism). Tapscott and Williams (2006) see such technologies as a means for growing a new economy driven by mass collaboration, peer production, and sharing. Others such as Haklay et al. (2008) believe the role of such technology lies between the two, depending on the challenge being faced.

Furthermore, while the tools and mashups presented here allow researchers not only to map their results but also to make such results accessible to anyone who has an Internet connection, this potentially allows one to make “mashups of mashups,” removing the culture of silos of information seen within many areas of social research. How such technologies and tools evolve is not only dependent on advances in technology itself but also the users of such technology. While this article has portrayed such advances as a considerable achievement, there are issues of copyright, accuracy, trust, and accessibility, which need to be considered in much greater depth. All we can do here is to highlight these. There are also issues relating to quality of the data. As we discussed above, where large numbers of individuals are generating data by crowdsourcing can lead to either surprisingly accurate or inaccurate data. Moreover, these crowdsourcing methods are limited by their exposure and the time that the media make their listeners or viewers aware of the resource. They are subject to all the problems posed by online surveys in terms of interpretation. Crowdsourced maps may not be “fit” for purpose and may simply provide anecdotal evidence. The information provider who uses such methods may have agendas other than simply gathering and

disseminating information such as using it for target marketing, while the abuse of such systems through the provision of fake data either through sourcing or through the simple manipulation of the online map poses a severe problem.

The mashup itself might also be subject to volatility in the provision and persistence of the web sites which provide the services needed to create and archive them. If the Google Maps API were to change, sites such as MapTube may have problems in adapting to any new API. If the API is withdrawn, this kind of mapping for the masses could disappear overnight. There is a small learning curve that users who do not have knowledge of how geographical information is represented visually have to negotiate before they can fully use the power of tools such as MapTube. The same issues apply to professionals and policy makers who might use such mashups to communicate with their peers or their clients.

Concluding Remarks: The Scope of Neogeography

Thinking of relationships as networks is an old concept that has taken on a new urgency in an age when interactions are ever increasing and as new technologies allow us to communicate in ways that previously were impossible. Network science may have helped us articulate immediate relationships in a much more coherent way but it is still limited in that the major barrier to an understanding of the great kaleidoscope of social relationships is that our ability to deal with networks is largely limited to two-way relationships; many of them are in great clusters but they are nevertheless always conceived of as two-way traffic. There is no coherent theory that tells us anything at all about how networks intersect, interlink, how networks that are spatial coalesce with networks that are manifestly nonspatial or aspatial, how networks that are temporal correlate with the spatial and the nonspatial, and so on. We do not even know if the social networks that emerge when users share spatial data and have spatial associations are meaningful. The fact that we can extract these types of patterns does not mean we can necessarily understand them.

Online mapping and maps represent an important forum for user interaction and cooperation as we have illustrated by examples, which seek to sample opinion using forms of crowdcasting and crowdsourcing. Maps built in this way may rival or surpass anything produced in more conventional ways as they are likely to be more informative and benefit from continuous improvement in the style of other web 2.0 resources such as Wikipedia. Their promise, however, will only be realized when better ways of capturing locational information emerge. Nevertheless, crowdsourcing maps are in their infancy, and once they really begin to take off (and include much else about local environments), the prospect of delving into the patterns of association that underlie the social networks created in these environments can begin in earnest. This is an area that promises to inform social science and social action in the near future in ways that presently we can barely anticipate.

Notes

1. The preprint server was established at Cornell University in August 1991 and it lets users freely post their papers to this database. It is a preprint service that is widely used in the Physics world but covers a range of papers in Physics, Mathematics, Computer Science, Quantitative Biology, and Statistics.

2. In the United Kingdom, the national mapping agency (the Ordnance Survey) operates in the commercial marketplace and these data cannot be used but several government agencies in other parts of the world make mappable data “free” in some sense.

3. See <http://www.casa.ucl.ac.uk/software/gmapcreator.asp>

4. See <http://www.londonprofiler.org/>; these maps are now a subset of MapTube that are accessible in MapTube style from <http://www.maptube.org/london/>

5. We are well aware that the whole web 2.0 movement throws up the ambiguity between what is a professional and what is an amateur, see Shirky (2008).

6. Di-Ann Eisnor (2006), one of the founders of the mapping site www.platial.com, defines neogeography as “a diverse set of practices that operate outside, or alongside, or in the manner of, the practices of professional geographers”; see Turner (2006).

References

- Carvalho, R., & Batty, M. (2006). The geography of scientific productivity: Scaling in US computer science. *Journal of Statistical Mechanics*, 10, P10012, 1-11.
- Eisnor, D. (2006). What is neogeography anyway? Retrieved December 13, 2008, from http://platial.typepad.com/news/2006/05/what_is_neogeog.html
- Gibin, M., Singleton, A. D., Mateos, P., & Longley, P. A. (2008). Exploratory cartographic visualisation of London using the Google Maps API. *Applied Spatial Analysis and Policy*, 1, 85-97.
- Goodchild, M. F. (2007). Citizens as sensors: The world of volunteered geography. *GeoJournal*, 69, 211-221.
- Haklay, M., Singleton, A. D., & Parker, C. (2008). Web mapping 2.0: The neogeography of the geospatial Internet. *Geography Compass*, 2, 2011-2039.
- Howe, J. (2008). *Crowdsourcing: Why the power of the crowd is driving the future of business*. New York: Crown Business.
- Keen, A. (2007). *The cult of the amateur*. London: Nicholas Brealey Publishing.
- Longley, P. A., & Singleton, A. D. (in press). Linking social deprivation and digital exclusion in England. *Urban Studies*.
- Shirky, C. (2008). *Here comes everybody: The power of organizing without organizations*. London: Allen Lane at the Penguin Press.
- Surowiecki, J. (2004). *The wisdom of crowds: Why the many are smarter than the few and how collective wisdom shapes business, economies, societies and nations*. New York: Little, Brown and Company.
- Tapscott, D., & Williams, A. D. (2006). *Wikinomics—How mass collaboration changes everything*. London, UK: Atlantic Books.
- Turner, A. (2006). *Introduction to neogeography*. O'Reilly, PDF Publication. Retrieved November 27, 2007, from <http://www.oreilly.com/catalog/neogeography/>

Andrew Hudson-Smith manages the GeoVUE and GENeSIS Nodes of the National Centre for e-Social Science (NCeSS) at the Centre for Advanced Spatial Analysis (CASA) University College London, where he is senior research assistant. His research work is in digital visualization, 3D GIS and CAD, and in web 2.0-based interfaces for urban planning and design. He recently published a monograph *Digital Geography* (see <http://digitalurban.blogspot.com/>). He can be reached at asmith@geog.ucl.ac.uk

Michael Batty is the director of CASA and Bartlett Professor of Planning, University College London. His research work is in urban simulation models and their visualization. His most recent book is *Cities and Complexity* (MIT Press, Cambridge, Massachusetts, 2005). He is the editor of the journal *Environment and Planning B*. He can be reached at m.batty@ucl.ac.uk

Andrew Crooks is the GLA economics research fellow in urban systems at CASA. His expertise is on agent-based modeling, where he is working on urban applications to residential segregation and to embedding these models within virtual worlds using map interfaces pioneered under the GeoVUE project. He runs the GIS and

agent-based modeling blog (<http://gisagents.blogspot.com/>), which contain links to his publications in this area. He can be reached at andrew.crooks@ucl.ac.uk

Richard Milton is a research fellow at CASA, where he has developed the GMapCreator software and MapTube portal. He has worked as a programmer in various projects at the Meteorological Office and at Criterion before coming to UCL to work on programming applications connected to e-Science. His work on pollution has been featured on BBC TV and he has published a series of articles on this work which can be found at the CASA web site (<http://www.casa.ucl.ac.uk/>). He can be reached at richard.milton@ucl.ac.uk