Kyndal A. Owens. Crowdsourcing the Census: Creating a Digital Database of U.S. Census Enumeration District Maps With Volunteer Participation. A Master's Paper for the M.S. in L.S. degree. July, 2012. 46 pages. Advisor: Richard Marciano

With the growth of the digital humanities and historical GIS technologies, scholars have begun to build national and urban databases of census and other records in order to trace the changes in boundaries over time and thus simplify the study of historical change. Because a project to digitize U.S. census enumeration district maps would require time, funding, and staff that most institutions do not have, the most sensible way to create such a database is through crowdsourcing. Other crowdsourced projects, both geographical and non-geographical, have succeeded owing to their creative design and competitive focus as well as the usability of the tools they provide to their volunteer users. Upon review of other crowdsourced mapping projects, recommendations are given for the creation of a GIS database for U.S. census enumeration district maps.

Headings:

Historical geographic information systems

United States – Census

Genealogy – Archival resources

Crowdsourcing

Digital libraries

Digital mapping

CROWDSOURCING THE CENSUS: CREATING A DIGITAL DATABASE OF U.S. CENSUS ENUMERATION DISTRICT MAPS WITH VOLUNTEER PARTICIPATION

by
Kyndal A. Owens

A Master's paper submitted to the faculty
of the School of Information and Library Science
of the University of North Carolina at Chapel Hill
in partial fulfillment of the requirements
for the degree of Master of Science in
Library Science.

Chapel Hill, North Carolina

July 2012

Approved by

Richard Marciano

INTRODUCTION

With the recent explosion in digital technologies, fields of scholarship long unassociated with technological tools have begun to experiment with digital methods. Today, the accepted term for this new, collaborative, interdisciplinary field is the "digital humanities," as codified by the National Endowment for the Humanities' 2007 establishment of the Office of Digital Humanities.[1] Because of the number of disciplines involved in digital humanities projects, as well as the newness of the field and the uncertainty as to what is possible in the digital age, "digital humanities" as a phrase "remains ambiguous and refers to a wide range of activities."[2]

One of the fields that has sprung up as part of the digital humanities phenomenon is historical GIS (Geographical Information Systems). Bailey and Schick (2009) point out that "historical research . . . demands the ability to synthesize processes in dynamic and nonlinear ways," often requiring the "ability to manage large arrays of data and variables."[3] In 1997, GIS was only just being introduced into the field of historical geography in order to better manage the types of data historians deal with.[4] According to Gregory and Healey (2007), GIS technologies provide three advantages to historians that are not afforded by traditional research methods: otherwise unrelated data may be linked through location; complex data visualizations such as maps and animations are possible; and data can be observed and analyzed spatially in previously impossible ways.[5]

Two major trends in the field of historical GIS are national GIS databases and urban development GIS.[6] The Great Britain Historical GIS, for instance, contains data from national sources like censuses and vital record registries, and focuses on the "core" issue of locating administrative boundaries as they moved and changed shape over time.[7] Indeed, the "usual approach" to national GIS development is to conduct "large amounts of research into where administrative boundaries were in the past and when they changed."[8] Urban historical GIS databases are smaller because they focus on a specific city like Tokyo or London, but like national GIS databases, their construction comes about because "a significant body of data exists that it is believed will provide a valuable research resource in GIS form."[9]

In the United States, the "primary source" of information on the nation's geographical transformations over time is the census.[10] As with other national historical geography projects, investigation of geographical change is hindered by the changing boundaries of census wards and districts over time (in the U.S., 1790-2010).[11] To study these changes by traditional means, researchers must create their own maps, as a result of which most studies of this sort are limited in number and in scope.[12] Consequently, a great need among U.S. geographers and historians is "a compatible set of electronic maps that describe the location of each geographic unit tabulated by the census."[13]

For urban geographers and historians (and historical geographers), a digital database of census unit maps could be used for numerous types of research, including but not limited to urban decline and renewal, settlement patterns of immigrants and ethnic groups, suburbanization, transportation routes, political district creation, areas of poverty and affluence, crime, public health and sanitation, and the segregation of residences.[14] Another group of researchers often neglected in the GIS literature is professional genealogists and amateur family history students, who frequently use census and vital records in their research.[15] Over all, therefore, access to digital census maps would be a boon to research in many areas, both scholarly and personal.

Presently, maps of U.S. census enumeration districts—the most common geographic level searchable on genealogical websites like FamilySearch.org—are available on microfilm at the National Archives.[16] Defined as "an area that could be covered by a single enumerator (census taker) in one census period," enumeration districts can be entire counties (where population is slim, such as rural areas), or a few city blocks if the population is dense.[17] The maps held by the National Archives provide boundary locations as well as, on occasion, other geospatial data like wards, precincts, congressional districts, and incorporated/unincorporated areas.[18]

The release of the 1940 census in digital form in April 2012 included scanned images of enumeration district maps, but the quality of the scans varies, making the maps frequently difficult to read and, moreover, split into several different pages with no digital search capability.[19] The original maps contain instances of coloration to indicate different types of boundaries, but the scans are all black-and-white, thus making much of the information inaccessible.[20] Even if the maps were scanned in higher quality, with color maintained, their usability would still be limited to a manual search through pages of scanned material with no way to visually compare changes in district boundaries over years and censuses.

A project to fully digitize available U.S. census enumeration district maps (in other words, to scan and georeference them for overlay on modern interactive maps) would require time, funding, and staff that most institutions do not have. Even large repositories like the British Library have only just begun to make their map collections digital and accessible to scholars and interested library users.[21] More importantly, they have done so not by hiring more staff or focusing more of their institutional funding on map digitization projects. Rather, a number of high-profile institutions have opted to rely on interested volunteers to digitize maps through a process known as crowdsourcing. While the institutions have to provide the tools for such projects, the projects themselves are carried out by unpaid users who access the tools digitally and freely spend their time and effort creating databases that will be accessible to anyone online. In the next chapter, I explore the concepts behind the notion of crowdsourcing and review some recent successful projects that have utilized this method.

CROWDSOURCING: WHAT MAKES IT SUCCESSFUL?

With the success of Wikipedia and similar free, user-built reference tools, historians and geographers (and other groups) began to contemplate the value of data created by "the masses." According to Andrew Hudson-Smith et al., crowdsourcing can be defined as "methods of data creation, where large groups of users who are not organized centrally generate content that is shared."[22] Like other Web 2.0-inspired projects, crowdsourcing relies on technologies that allow users to create and upload their own data and share it with others via the internet. Through these Web 2.0 technologies, users can build mash-ups (websites with material from other sites embedded within) and collaborate with other users in ways that were impossible in the internet's early days.[23]

The usefulness of crowdsourcing is, of course, dependent on the accuracy of the content produced by the users that make up the "crowd." In theory, crowdsourcing works because of the truism "two heads are better than one"; that is, an individual, or several individuals, may be wrong, but with enough users contributing, their average result will be correct.[24] Moreover, crowdsourced users may even arrive at results that could not be achieved through other methods.[25] Through the "uncoordinated and voluntary" agency of non-expert participants, expert-level results are frequently obtained, a phenomenon often referred to as "the wisdom of the crowd."[26]

One popular way libraries, museums, and other cultural institutions have successfully employed crowdsourcing is online transcription projects. The New York Public Library's website "What's on the Menu?" is perhaps the most well-known of these volunteer opportunities, but others exist as well. By taking a look at some of these projects, we can observe their practices and techniques and determine which are the most successful. I have selected three projects of current interest, which I review and analyze below.

What's on the Menu?

"What's on the Menu?" was launched in April 2011 with the object of producing a fully-searchable database via line-by-line menu transcriptions by users. In two days, users had transcribed over a thousand dishes from digitized menus in the New York Public Library's menu collection. Within a year, that number grew to 866,636 dishes.[27] The project has been so successful that it was awarded the 2011 Roy Rosenzweig Prize for Innovation in Digital History (sponsored by the American Historical Association and the Center for History and New Media at George Mason University). [28] In addition to providing transcription assistance, users have also volunteered suggestions, asked questions about menus and the transcription process, and commented on the project to NYPL staff.[29]

Remarkably, this tremendous public response to "What's on the Menu?" was a surprise to its creators.[30] Although the Library had gone to great lengths to design an effective transcription tool that would draw users in and foster involvement and interaction among users, librarians had "no idea" what the end result of their planning and publicizing efforts would be.[31] Besides the initial preparation of digitizing menus and updating their metadata, NYPL staff and interns promoted the project on Twitter and other social media sites.[32] The transcription interface itself was designed in conjunction with “leading game designers” with the purpose of encouraging user participation by making menu transcription an “engaging and competitive activity.”[33] The final result, it was hoped, would be a “culinary research database of extraordinary power,” which would have been “prohibitively expensive” for the Library to build on its own.

Upon visiting “What’s on the Menu?”, the user is faced with an invitation to begin transcribing. A short, readily-comprehensible tutorial is provided, and then the transcriber selects a menu from the Library’s database, based on labels classifying the menus as complete, partially transcribed, or untranscribed. The interface design prompts a highlighted region to move from line of text to line of text automatically as the user completes a dish. Then, as soon as a dish is completely transcribed, it takes its place along the side of the page as a hyperlinked tag. By clicking the tag, the user can see other menus with that dish, or, alternatively, he or she can visit the “tag cloud” to see dishes other users have transcribed, with type size increased according to the popularity of the dish.[34]

While we can only speculate on the cause of this project’s success, there are three elements that stand out: the interesting nature of the material, the game-inspired design, and the NYPL’s promotional efforts. As already stated, the Library publicized the project by giving it its own Twitter account and using this medium frequently and consistently.[35] Furthermore, the project was further supported with its own “What’s on the Menu?” blog, which joined the ranks of the Library’s many other blogs accessible from its homepage.[36] Clearly, a large part of the project’s success can be attributed to the Library’s efforts to increase public awareness of the need and the opportunity.

Secondly, the Library designed the transcription tool on gaming principles in the (subsequently-justified) hope that users would find participation fun and rewarding, even addictive. While there are no explicit rewards for transcribing, such as a user ranking system or recognition on the Library's website, the immediate gratification users receive through the tagging feature may constitute a reward in itself. Rather than having to wait days or weeks to see the results of their work, users are able to visualize their contribution as soon as it is made and thus have more motivation to continue transcribing.

Finally, the place of food and dining in popular culture certainly boosted the project's popularity. The following user-submitted comments from the "What's on the Menu?" blog are a sampling of responses generated by participation in the project:

> *As a long time menu reader and food history hound, this service project is perfect for me!* (Karen S.)[37]

> *This project is a lot of fun and I have enjoyed browsing the menus. I teach Culinary Arts and my Food and Culture class will get a kick out of this.* (Jessica)[38]

> *This project is so much fun! Thank you for opening this up to the public,*
> *I really enjoy merging my inner food geek with my inner editing nerd,*
> *and all for a great cause--just love it!* (Elizabeth)[39]

From these reviews, it is apparent that at least some users participate because of a native interest in food and food history, whether professional or amateur. Also of note as a motivation is the ability for users to browse the menus, even though they are yet to be fully searchable. The combination of this public fascination with the other two aspects—a game feel and strong promotion—made this crowdsourcing project viable and gave the public access to a collection that might otherwise have remained offline until time, staff, and funding allowed the Library to perform the work itself (if ever).

Old Weather

Another project based on crowdsourcing is Old Weather, a collaboration between numerous British organizations including Oxford University and the National Maritime Museum.[40] The object of the project is to produce transcriptions of weather reports in historical Royal Navy ship logs so that they can be used to predict future climate change. A sidebar on the website's homepage tracks current statistics, which in June 2012 stood at 1,062,722 pages from the logs of 289 completed ships, with over 28,000 volunteers transcribing.[41]

The Old Weather project is different in several ways from "What's on the Menu?" but shares certain aspects like its gaming mentality. Unlike food and restaurant menus, weather does not have a particularly strong hold on the average user's attention span. Converting a topic that is traditionally viewed as the last resort in awkward small talk into a successful crowdsourcing project requires ingenuity, and Old Weather displays this in a number of ways. First, the site's introduction—contained on the homepage—clearly explains the justification for the project and tells users exactly why their help is needed and how their work will contribute to the ends of science, climatology, and history. Second, Old Weather provides information on the ships whose logs have been digitized for transcription, allowing users to become invested in the period and circumstances in which the logs were created. Finally, and most strikingly, the site furnishes users with distinct identities and rewards them based on the number of log pages they transcribe, creating a competitive aspect lacking in "What's on the Menu?".

To achieve this competitive game-like atmosphere, Old Weather requires that users sign up with a user name and password, which immediately requires greater effort and commitment from the user than the menu project does. However, by signing up, the user gains the ability to "join the crew" of a particular ship that he or she chooses.[42] Once assigned to a crew, under the leadership of a captain (typically a super-volunteer of sorts), the volunteer receives a rank and the opportunity for promotions after transcribing specific numbers of pages.[43] The interface therefore rewards (if only psychologically) users for greater activity, and eventually provides them with a more involved role and greater responsibility, if they stay with the project long enough. Although eighteenth-century naval captains' handwriting can be difficult to decipher at times, and the use of codes sometimes requires further investigation into (provided) keys, the community forum gives users a place to ask questions, advise other users, and generally communicate with the communities aboard their ships.[44]

Combine this competitive spirit with a relatively easy-to-use interface, and transcribing old weather logs becomes surprisingly addictive. Another exciting feature, especially as it pertains to historical geography and data visualization, is the embedded maps that display the paths of the ships as their logs are transcribed. Once a geographical location for a ship is transcribed—either coordinates or port locations, such as Santo Domingo or Cadiz, Spain—a marker appears on the map so that users can literally see the direction and distance traveled by their ship.[45] As with "What's on the Menu?", the immediate gratification users receive as their every act of transcription produces a visible change in the website is compelling encouragement to continue participating, even though participation requires a greater degree of investment and effort.

Indexing the 1940 Census

The final transcription project I will mention is the massive effort by the website FamilySearch.org to crowdsource the indexing of the 1940 U.S. census, so that genealogists and other researchers can search multiple fields without having to wait for a paid subscription service to make the census similarly accessible.[46] According to the June 28, 2012, indexing report on the FamilySearch blog, seventy-five percent of the census records had been indexed, arbitrated, and made fully searchable since the census' release on April 2, 2012. The entire project is expected to be complete at the end of July 2012.[47]

To participate in indexing on FamilySearch, users download a program from which they can load "batches" of records and transcribe according to provided instructions.[48] Each record is indexed by two volunteers and then a third volunteer arbitrates the record to resolve any discrepancies or disagreements between indexers. This system is meant to improve accuracy and conformity, particularly as census enumerators did not always follow enumeration guidelines and their handwriting is frequently only semi-legible. Once a user indexes one hundred records with a high rate of accuracy, he or she is invited to become an arbitrator.[49] Consequently, arbitrators are volunteers with experience and proven aptitude in indexing records, providing assurance of their probable good judgment while at the same time rewarding indexers for conscientiousness and continued participation.

The opportunity to become an arbitrator may encourage participation for some, but in large part users volunteer to index because they want to obtain the finished product. Many are genealogists or amateur family history buffs looking for their own ancestors or family members in the census. Comments on the June 28 indexing report mentioned above include:

> *I was so excited yesterday to see California orange [fully searchable]! I found my family easily and got a few new insights and clues.* (elzabel)[50]
>
> *Was so glad to see the state I needed was in orange.* (Elistariel)[51]

Despite the lack of reward provided by the indexing interface itself, participants are so invested in the results of their work that they willingly volunteer so that they and others can have access to the census. FamilySearch contributes to this interest by publishing a state's records as soon as they are completely arbitrated rather than reserving them all until the completion of the entire project.[52]

As has been seen, with the proper design, promotion, and reward system, crowdsourcing can be an effective way to quickly and inexpensively complete projects that would otherwise have been impossible or required much more time, staff, and funding. Censuses and other historical materials of genealogical interest are especially engaging for users because of their investment in family history and shared research. But is crowdsourcing a GIS project, even one involving the census, feasible? To determine whether it is, I will now review recent trends in geospatial crowdsourcing and consider the implication for historical GIS projects.

Crowdsourcing Geospatial Data

In a 2010 article entitled "Crowdsourcing Geospatial Data," Christian Heipke identifies seven sometimes-overlapping groups likely to contribute to a crowdsourcing project that involves mapping.[53] Two of the groups, map lovers and experts, are the most invested in the material, as they either study and/or create maps for a living or they have an already-existing love for and experience with maps and mapping. Another three groups—casual mappers, media mappers, and open mappers—are less consistently engaged and may be less interested in accuracy than in sharing their own data. (Media mappers are users who respond to promotions and competitions but may drop off after the promotion period has ended; open mappers are participants who believe in the philosophical ideal of openly-available online data.) Finally, the last two groups are passive mappers and what Heipke dubs "mechanical turks" after the Amazon.com marketplace of the same name. Passive mappers may provide geospatial data without realizing it, as when their cell phone's GPS system monitors their movements. Mechanical turks are paid volunteers who sign up to accomplish certain tasks for a fee.[54] Each of these groups constitutes a different level of commitment to mapping projects, different motivations, and different styles of participation. In order to achieve success in a crowdsourced geospatial data project, these groups must be taken into account and appealed to in their own right.[55]

While crowdsourcing a project relinquishes the security of expert data creation, it can in some sense add another kind of expertise. First, users may have local knowledge that mapping experts lack, especially if the dataset is an extremely large one. Second, the number of contributors producing, perusing, and correcting data lowers the possibility of error, both from a statistical perspective (the wisdom of the crowd versus the individual expert) and the checks-and-balance process crowdsourcing creates (a hundred people are more likely to catch a mistake than a handful of people are).[56]

After investigating the accuracy and completeness of two successful crowdsourced geospatial data projects, OpenStreetMap (active contribution) and

TomTom's HD Traffic system (passive contribution), Heipke predicts that the growth of new Web 2.0 technologies will continue to improve both the process and the results of crowdsourcing.[57] On the other hand, a problem arises in regard to the sustainability of open mapping projects and the difficulty of maintaining interest in a project "once the initial hype has passed."[58] (This problem is less of a challenge for non-continuous projects that have a definite end, such as indexing the 1940 census or transcribing a finite menu collection.)

According to Hudson-Smith et al., "the problem in [digitally] 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."[59] As a solution to this problem, Hudson-Smith and his co-authors created a tool called GMapCreator to convert user-generated maps into raster maps, which is the type of map necessary for overlay on Google Maps.[60] Other free online software can likewise convert data files into the type of file GMapCreator requires.[61] Once a user uploads a map, other users can overlay it on their own maps, giving them a means to visualize and compare data like geolocated crime statistics in a manner previously unavailable.[62] In the example Hudson-Smith et al. provide, the "visual correlations" are "of significance," revealing that user-sourced data "can add real value to . . . insights."[63]

Despite the potential they see in "mapping for the masses," Hudson-Smith et al. express concern about the inequality of access to the technologies needed for mapping projects, as well as unsolved questions of "copyright, accuracy, trust, and accessibility."[64] However, they conclude that crowdsourced maps "may rival or surpass" maps created in "more conventional ways" and will have the advantage of "continuous improvement," as users can continue editing them *ad infinitum.*[65]

Considering the success of numerous crowdsourced geospatial projects in which users upload their own maps and interact with them through GIS features like overlay, it appears that crowdsourcing is a viable method for creating useful data mash-ups and up-to-the-minute local geographical visualizations. On the other hand, historical map work requires a different skillset, one perhaps more likely found among participants in crowdsourced historical projects. The question now becomes whether or not the historical and geospatial crowds can overlap, and if so, what tools are the most effective for encouraging, enabling, and monitoring these crowds? In the following chapter, I will compare two high-profile projects that train users to georeference historical maps and rely on the crowd's participation to create new products of historical interest and significance.

GEOREFERENCING TOOLS FOR CROWDSOURCED USERS

Both the New York Public Library and the British Library have recently developed tools that allow users to interact with the libraries' collection of digitized maps in an unprecedented way. Rather than simply exploring and viewing scanned images of maps without reference to a larger and/or more current map, library users now have the opportunity to overlay old maps onto modern interactive maps (called "georeferencing") and, in the New York Public Library's case, perform tracing and annotational functions as well. In the following pages, I compare the success of both libraries' digital mapping projects and critique the usability and reliability of the tools underlying them.

The British Library's Georeferencer

According to the British Library's website (www.bl.uk), the purpose of its Georeferencing project is "to make a selection of maps fully searchable and viewable using popular online geotechnologies."[66] To accomplish this goal the Library is "crowdsourcing location data on Britain's historic landscape," particularly the landscape prior to the nineteenth century, when mass industrialization altered it nearly beyond recognition.[67] The project draws from two groups of maps in the British Library's collection: Ordance Surveyors' Drawings (1789-1840 maps of England and Wales) and the Crace Collection of London maps (1560-1860). These map collections were chosen for their thorough coverage and their usefulness for "overlaying and comparison with modern mapping."[68]

Beginning in February 2012 and running until January 2013, the georeferencing project aims to completely georeference the two collections via crowdsourcing, with the hope that this method may be extended to other map collections in the future.[69] After a successful pilot (in which over seven hundred maps were georeferenced by the public in a week), the Library has continued to add more maps from the selected collections and has made the already georeferenced maps accessible for viewing and further editing by users.[70] While this tremendous initial response bodes well for the overall success of the project as defined by the Library (one hundred percent of maps georeferenced), it does not say much about the usefulness of the georeferenced maps in terms of accuracy and reliability. Because the Library has not released details of the accuracy of the georeferencing done to date (or the degree of accuracy it expects and considers acceptable), an understanding of this aspect of the project requires personal investigation.

Having had some training in the principles of GIS and georeferencing with ESRI software like the ArcGIS suite, I was probably more familiar with the requirements and best practices for georeferencing than many users (assuming participation by more "casual mappers" than "expert mappers," as defined in chapter 2). That being said, I was very grateful for the detailed instructions provided on the Library website, as the Georeferencer was somewhat difficult to negotiate without guidance. Once a map is opened in the Georeferencer, it is displayed side-by-side with a modern map which may be enlarged, moved, or searched through using the provided gazeteer (more on this later). The default modern map is Open Street Map by MapQuest, but users may opt for another base map by choosing from a drop-down menu.[71] (I missed this feature initially, so I used Open Street Map out of what I believed to be necessity.)

The map I selected from those already georeferenced was an 1840 map of the southern bank of the Thames in Southwark, London.[72] Metadata for each map is provided in a section of the Georeferencer called “This map,” accessed through a tab at the top of the tool (see Figure 1). The information provided (the map’s creator, a description, and the year of creation, among other details) is probably more useful for a historian or a cartographer studying the map, as envisioned by the British Library, than it is for an amateur interested only in georeferencing the map. Nonetheless, knowing the background of a map does at least help locate it in time as well as space, especially if outside research is required (as I found to be true in my case).

My chosen map had been given three control points by its previous georeferencer, which situated it in the correct general area but did not seem to line up streets and other landmarks very well. Rather than relying on my own judgment, I visited the tab marked “Accuracy” to see how well the original georeferencer had done and what elements of the map were used to determine accuracy. I found that the accuracy measurement tool did not function without at least five control points, which would therefore be the intuitive minimum number of points required to render a map “fully georeferenced.” However, the Library’s instructions require only a minimum of three points to achieve fully georeferenced status (though they encourage more).[73] This discrepancy is puzzling, as the provision of an accuracy-checking tool seems superfluous if georeferencers are not required to prepare their maps to be compatible with it. Had I been viewing the map as a scholar rather than a volunteer, I would have had to add more points to the map in order

Figure 1. *"This map" tab on the British Library's Georeferencer tool. From* British Library Online Gallery—Georeferencing.

to see a measurement of its accuracy (a time-consuming process and one frustrating for a user expecting a completed map).

After accustoming myself to the zoom and mobility features on both the old and the modern maps, I began looking for shared features present in both maps. This process was facilitated by the gazetteer available for finding locations on the modern map. Once a street, intersection, or landmark is identified on the old map, it can be searched for in the modern map using a predictive-text search box (see Figure 2). If the location still exists and is in the gazetteer, the modern map will automatically focus in on the particular area. I found this function to be amazingly useful and it drastically shortened the amount of time it would have taken to find the corresponding area on the modern map manually.

Unfortunately, the passage of 172 years had greatly changed the landscape encompassed by my 1840 map, and many of the smaller roads present in 1840 no longer exist today, having been replaced by part of Charing Cross Station. More confusingly, some street names have changed and a number of landmarks (e.g. St. Thomas' Hospital) have since been relocated to other areas of London. Even with my basic GIS experience and knowledge of this area of London (I lived very close to the area of the map for two months), I still had considerable difficulty finding points on the maps that I could be sure were aligned properly. What was seemingly a simple task of studying two maps for matching landmarks became an extended search for information on the modern buildings located at the former site of St. Thomas' Hospital and the renaming of certain roads.

Once I had placed seven control points, I saved my changes and moved on to the "Visualize" tab to see the results. The visualization tool uses Google Earth to display the georeferenced map, with the option to export as a KML file. As in Google Earth, the

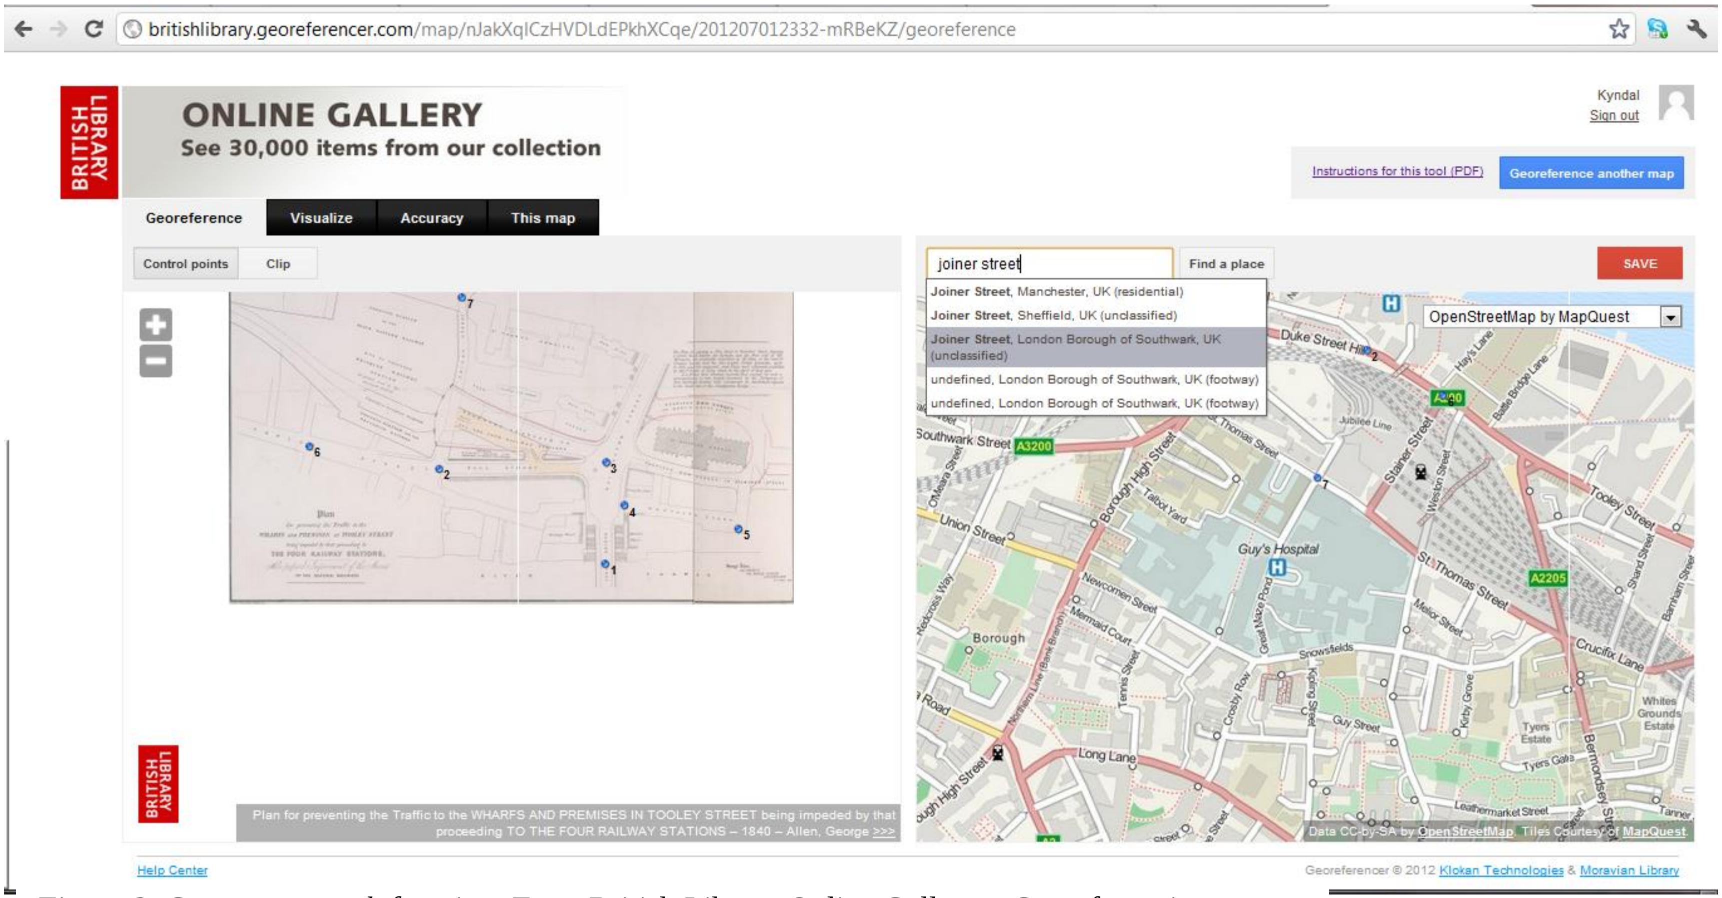

Figure 2. Gazetteer search function. From British Library Online Gallery—Georeferencing.

georeferenced map layer's transparency may be adjusted so that the modern base map is visible underneath, allowing a better comparison between old and new (see Figures 3 and 4). Satisfied with the alignment of the maps, I returned to the accuracy tab to learn how the measurement process works.

The measurement tool, MapAnalyst (http://mapanalyst.org), uses the control points to create a grid displaying the warp of the georeferenced map and distortion between control points, as well as isolines scale and rotation, the end result being "visualizations that illustrate the geometrical inaccuracy and distortion calculated from the georeference," thus helping to "identify wrongly assigned control points."[74] While undoubtedly a useful tool for map experts, the tool is less than intuitive for a GIS novice (see Figure 5). Even if an inexperienced user deduced that the red circles on the grid indicate potential inaccuracies (which the website does not make clear), there is little he or she can do to correct them. In my case, because of the limited zoom on the base map (at least in the Open Street Map default), it was difficult to place control points with the same degree of accuracy as on the historical map, causing a frustrating experience in which I knew the points were divergent from one another, but I could not find a way to make the points more accurate. Moreover, there was no apparent way to rotate the older map (oriented south) to face the same direction as the modern map (oriented north), so adjusting control points was frequently confusing by virtue of the maps' opposing orientations (see Figure 6).

In sum, the British Library Georeferencer tool has a roughly equal number of pros and cons. On the plus side, the tool is obviously attractive to users, as the public accomplished the initial goal of over seven hundred maps in less than a week. The project

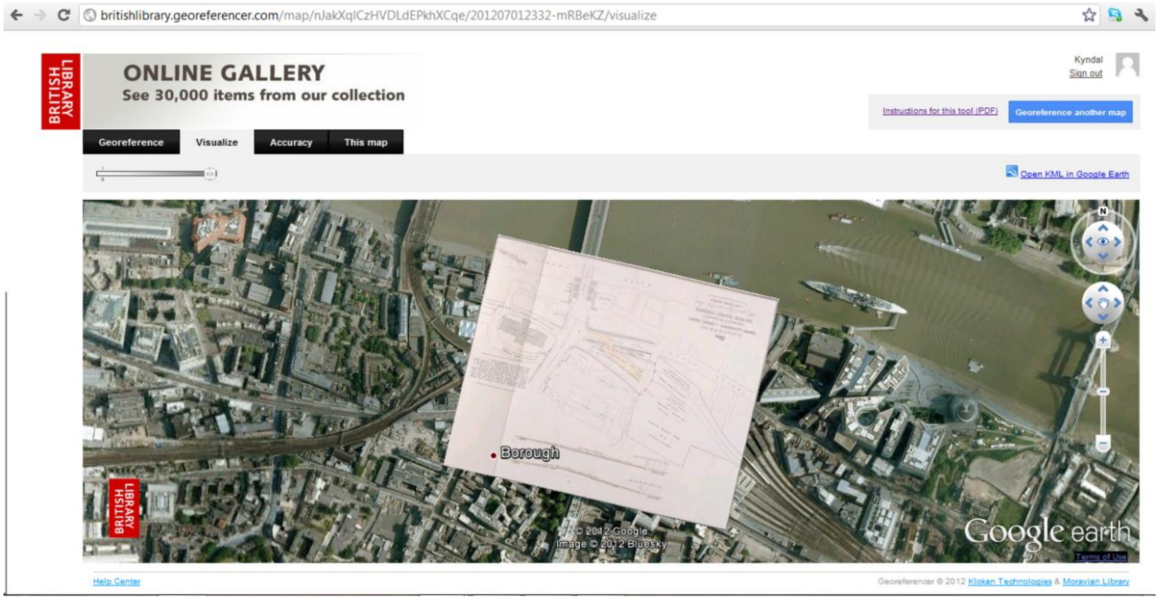

Figure 3. Georeferenced map in "Visualize" tab. From British Library Online Gallery—Georeferencing.

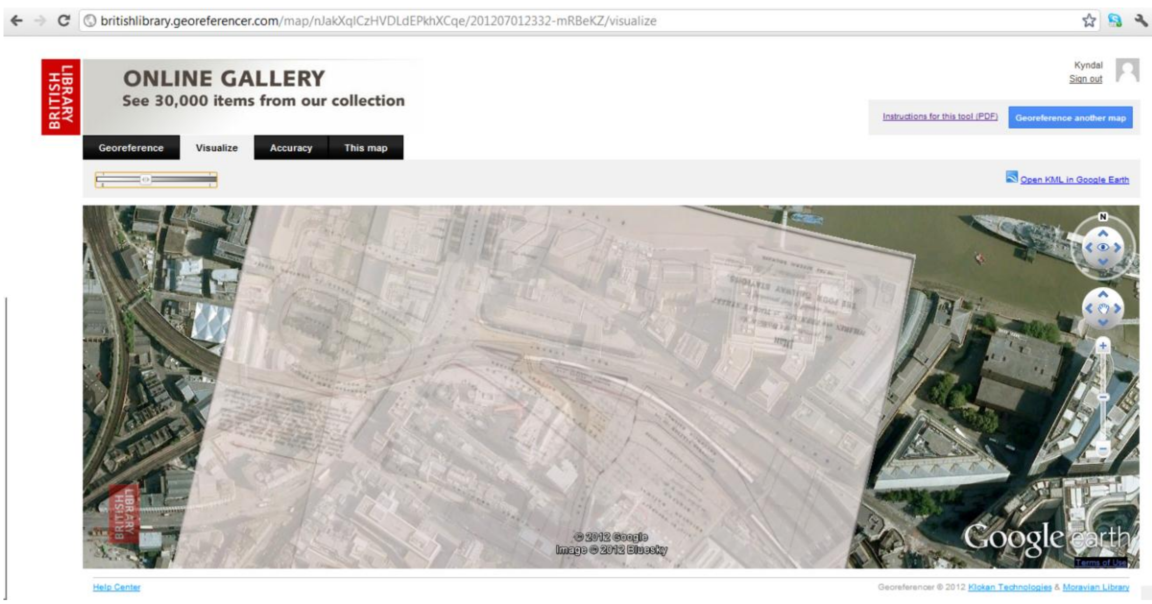

Figure 4. Closer zoom of georeferenced map, with transparency increased. From British Library Online Gallery—Georeferencing.

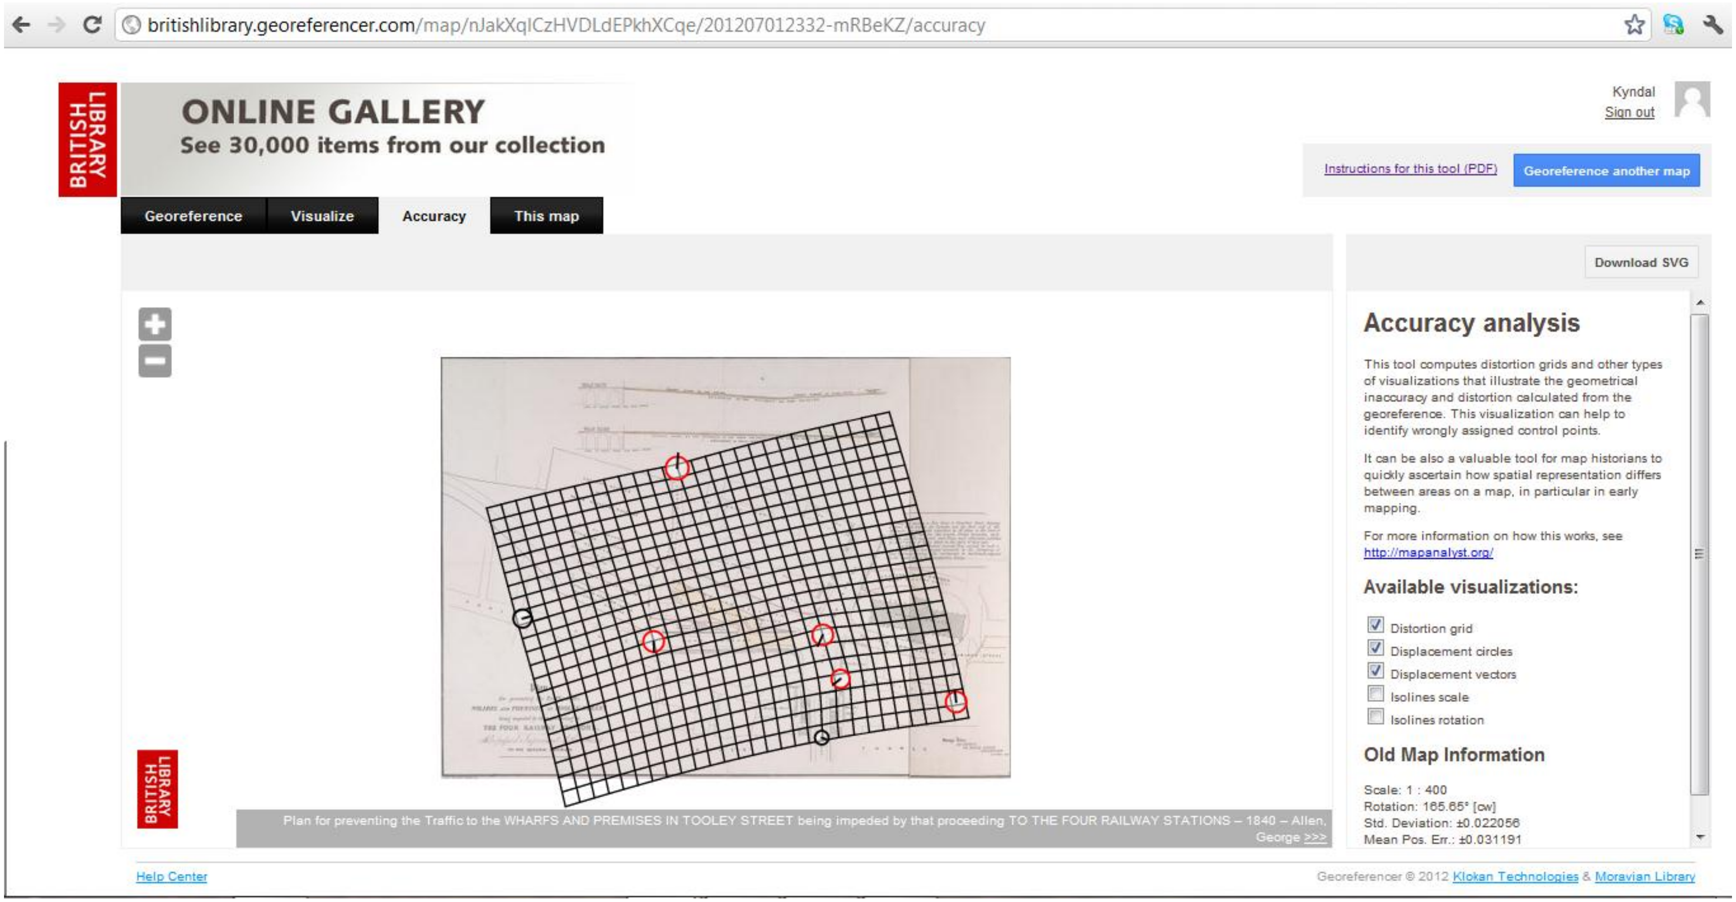

Figure 5. Distortion grid on "Accuracy" tab. From British Library Online Gallery—Georeferencing.

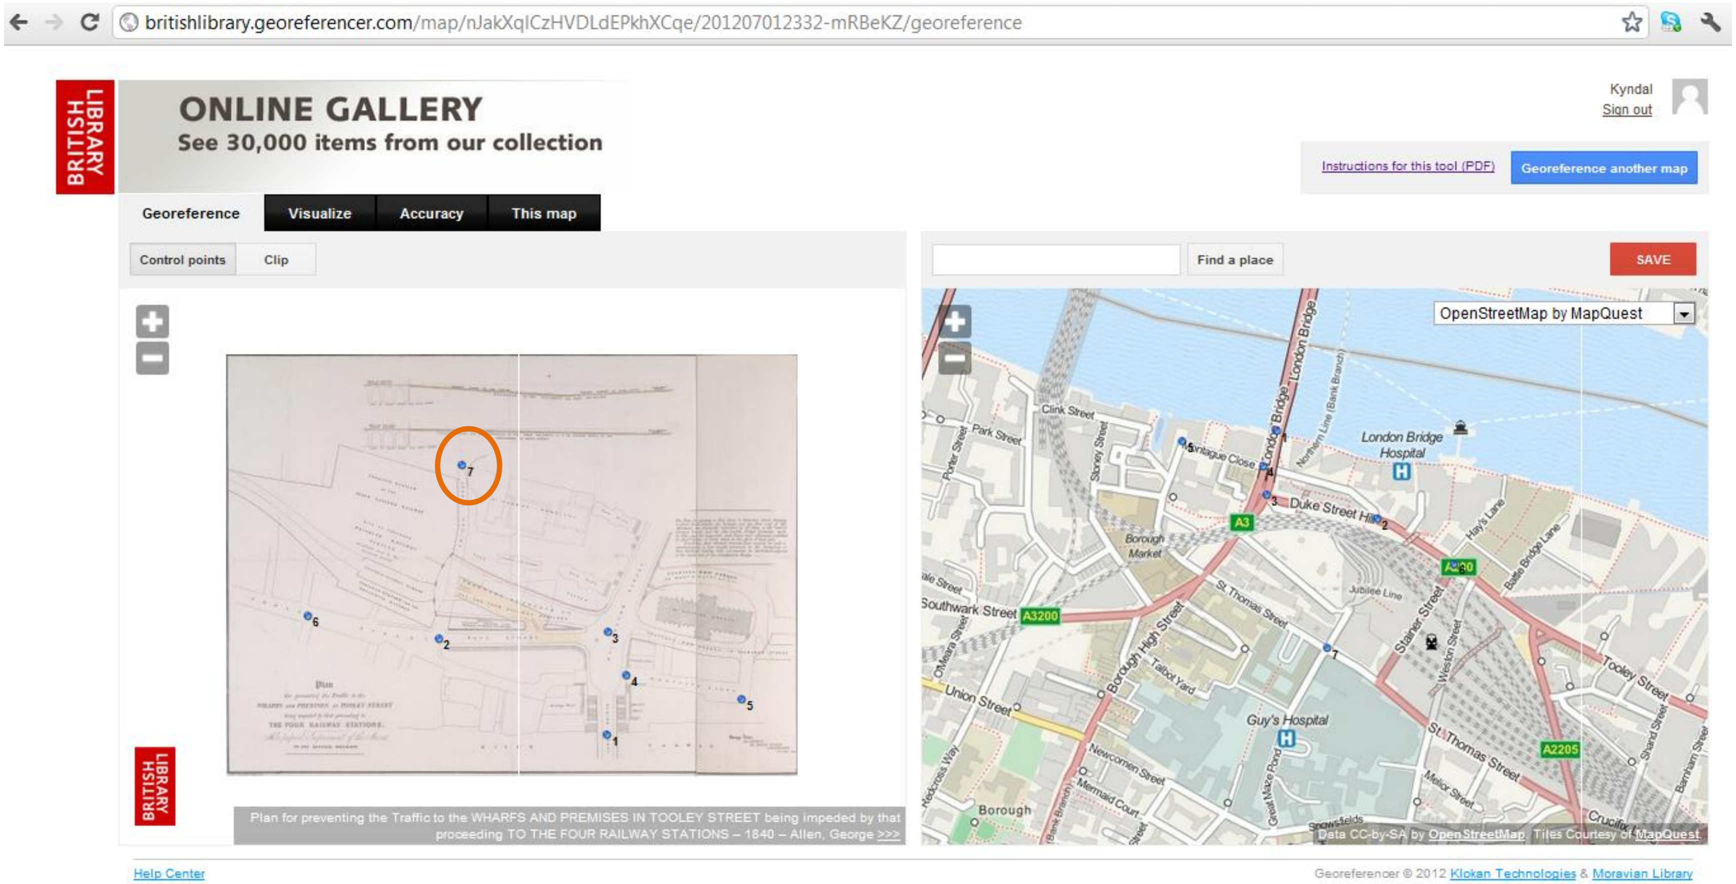

Figure 6.Old and new maps with control points (blue). In the modern base map (right), the river is at the top of the screen. In the historical map (left), it is at the bottom. Control point 7 is highlighted in orange on both screens to illustrate the difference in relation to other control points. From British Library Online Gallery—Georeferencing.

includes comprehensible instructions and clearly differentiates between steps, with tabs for georeferencing, visualizing, checking accuracy, and learning about the historical map. The gazetteer search function is extremely useful, and the effort to provide an accuracy measurement makes the georeferenced maps infinitely more valuable to researchers and map experts than maps without any sort of moderation or review.

On the other hand, there is a discrepancy between the georeferencing instructions (aimed toward simplicity for users and thus requiring only three control points) and the accuracy-checking tool (built to compare five or more control points). Furthermore, the limitations of the base maps make accurate georeferencing difficult, even for experienced georeferencers. Finally, the usefulness of these maps for researchers (the stated goal of the project) is questionable when scholars' idea of a "fully georeferenced" map does not line up with the Library's definition of a fully georeferenced map.

NYPL Map Warper

Although not as popular as *What's on the Menu?*, the New York Public Library's map warper has still received a positive response from the public. As of June 13, 2012, approximately 2400 maps from 162 New York City atlases in the NYPL's collection had been warped (georeferenced) using the tool.[75] Considering that the website launch was announced in February 2010, this number (more than triple the number of georeferenced maps on the British Library site) is less impressive but still an indicator of success.[76]

The NYPL has the ambitious goal of georeferencing all of its scanned maps, which in January 2012 included all of the public domain New York City atlases in its collection (over ten thousand maps).[77] The warping tool was built not only to hasten this process through crowdsourcing, but also to allow users the world over to "virtually stretch old maps onto a digital model of the world."[78] Using the Warper, participants are "creating a new [map] copy that is not only aligned with spatial coordinates on the Earth, but normalized across the entire archive of old maps." The copy can then go through a series of further transformations, resulting in a systematic "workflow" of warping, cropping, sewing together ("mosaic-ing"), and tracing.[79] In the end, this "long-term, collaborative research project" involves staff, library interns, and volunteers "accret[ing] little bits of new information to a collectively held historical knowledge base, free and open to all."[80]

Unfortunately, the extensive justification and explanation of the project are not accessible from the Warper's website (maps.nypl.org), either as a division of the site itself or via a hyperlink connecting the site to the pertinent NYPL blog posts. The site contains a brief paragraph describing the project and an embedded video demonstrating how to use the Warper, but the detailed instructions for warping, cropping, and tracing (provided on the blog[81]) are nowhere to be found. Matt Knutzen of the Library's Maps Division gives an excellent and thorough introduction to the project, breaking it down into its elements and providing examples of outside projects that have used the Library's mapping tools.[82] However, this information is not readily available to users, as it requires a search of the NYPL website and some trial-and-error guessing to turn it up (unless, of course, one is a regular follower of the blog).

From the outset, therefore, the NYPL's georeferencing site functions largely as a database for digitized maps, with an invitation to users to participate in a scarcely-explained project (in contrast to the British Library's site, which functions as a project base and contains all the information pertaining to it). If a user opts to participate, he or she selects a map for rectifying/warping through a search bar with fields for title and description, as well as more technical options like catalog ID. Unlike the British Library's Georeferencer, in which maps are selected by location, NYPL users can search for any feature(s) they choose and then browse the results for a map that interests them.

Once a map is selected, the Warper (an interface with multiple tabs, similar to the Georeferencer) opens on the default "Show" tab, where the historical map is visible as a scanned image with metadata (see Figure 7). The user may then select the "Rectify" tab, which brings up a split-screen frame with both the old and the modern map (Open Street Map) side by side, ready to receive control points (see Figure 8). As with the Georeferencer, there are other base map options, but I found that only Open Street Map worked with the map I selected, as it was of Long Island rather than New York City. If I had chosen a map of downtown Manhattan, for instance, and zoomed in to a certain level, I could have used one of the other maps. Once again, my task was made more difficult by the absence of a base map with few other features besides streets, and I was again forced to turn to Google Maps once or twice to confirm a control point location.

Also lacking was the gazetteer search function that was such a noteworthy element of the Georeferencer. In order to negotiate the modern base map, I had to constantly zoom in to read street names (which must be at a certain magnification to appear), zoom out to orient myself better, then zoom in to another location, repeating these steps over and over for each point I set. While this method is workable, it requires a considerable investment of time and energy, particularly if one is unfamiliar with the area. On the other hand, as a New York City public library displaying New York maps,

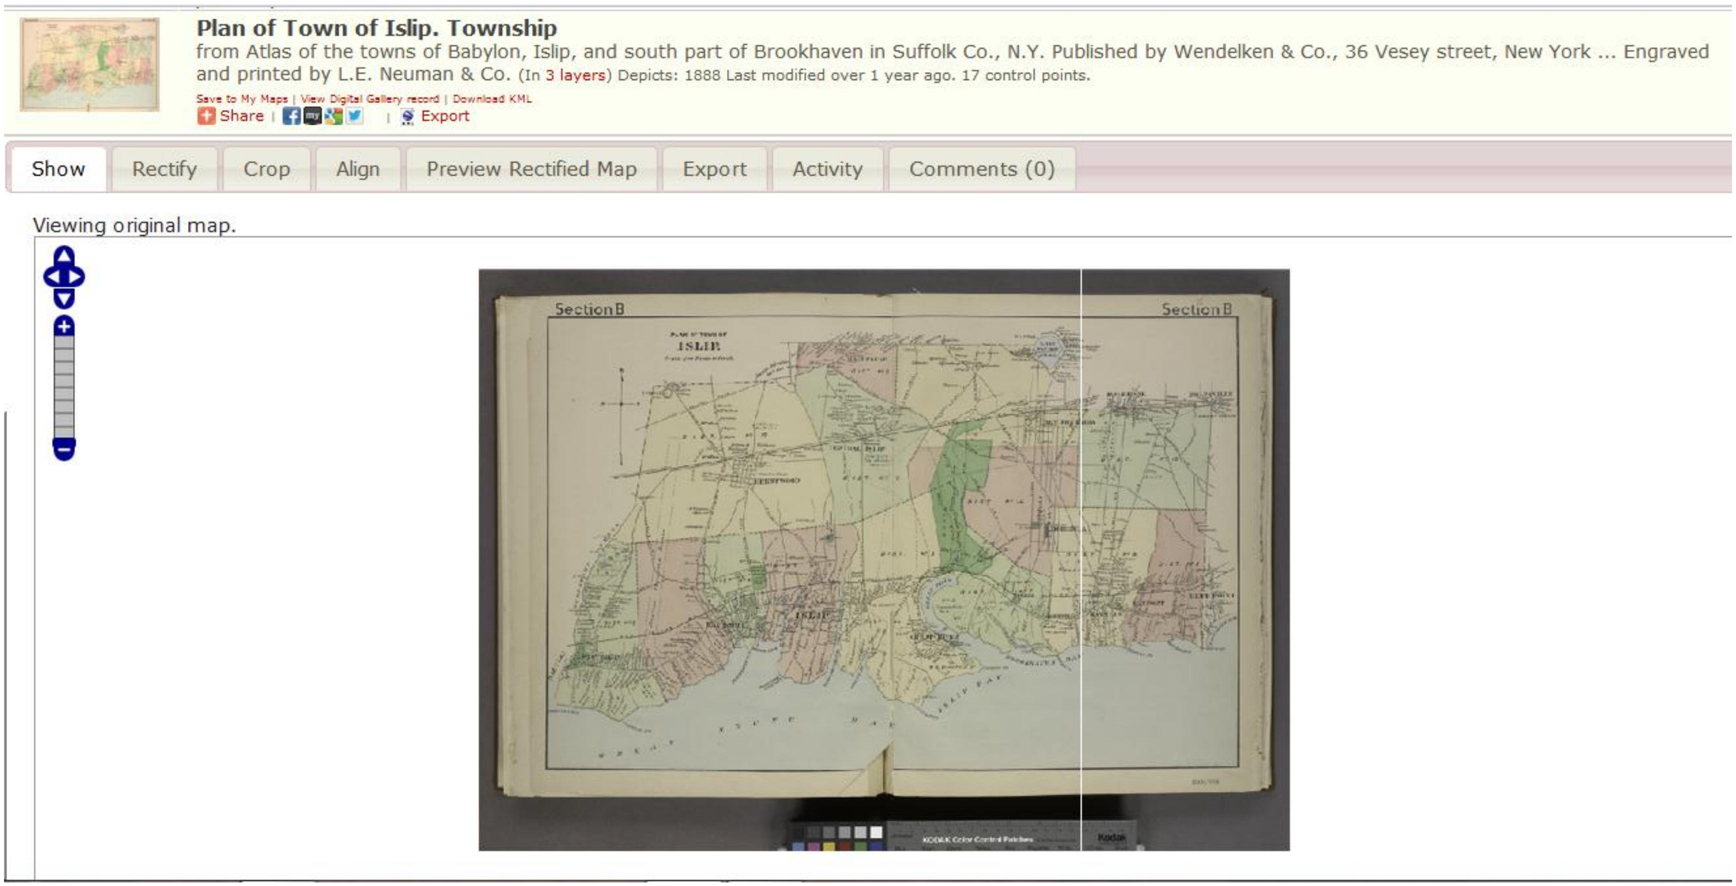

Figure 7. The "Show" tab on the New York Public Library's Map Warper. From maps.nypl.org.

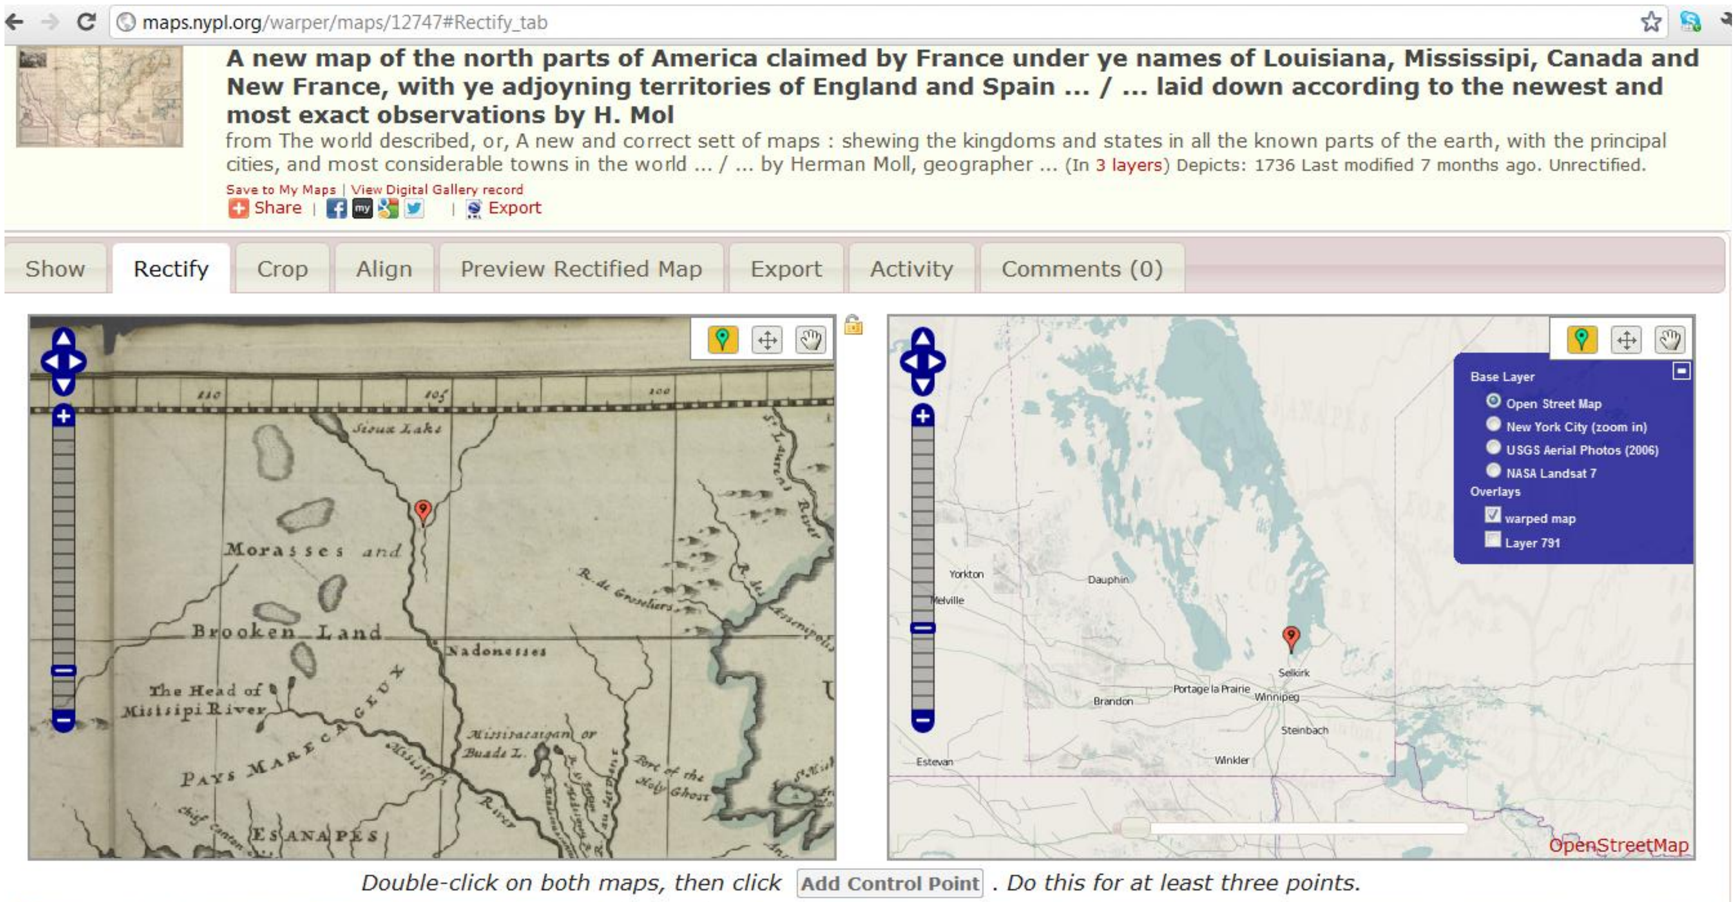

Figure 8. The Warper's "Rectify" tab with one control point. From maps.nypl.org.

the NYPL likely anticipates the project to appeal to residents interested in the historical geography of their own neighborhoods, thus increasing the value of the crowdsourcing method and decreasing the need for complicated search features. Regardless, the absence of a gazetteer or atlas index of any sort hindered my efficiency in georeferencing the map.

In terms of accuracy measurement, the NYPL tool is vastly superior to the British Library's, at least for casual/non-expert users. Although it suffered from a lack of explanation (like the rest of the site), I eventually figured out that the colored markers used to indicate control points are coded based on percentage of error. An optional display underneath the rectifying screen shows the locations of control points and their degree of error in red, yellow, green, or blue, depending on the number. This system requires little to no understanding of the principles of cartography or digital mapping to be coherent, and it is an immediate and easy way to see how accurate a certain point is in comparison to others and where a mistake may have been made. Unlike the Georeferencer's tool, MapAnalyst, the color-coding system is just as useful to the georeferencer as to the scholar viewing an already-georeferenced map.

Another interesting feature of the NYPL map gallery is that the metadata for each map includes whether or not the map has been rectified, and if so, how many control points it currently has. This is an especially useful bit of data for both volunteers and researchers, as it indicates how closely aligned a map is to its base map (accuracy/reliability) as well as giving users a way to judge whether a map needs further work, even if it is technically rectified (assigned three points). For instance, a map of Long Island that is rectified beyond the minimum number of points may still display unsatisfactorily when warped, as is evident in Figure 9 (a layer composed of several rectified maps, one noticeably skewed). In these cases, at least, the problem is immediately evident, as opposed to inaccurate rectifying in the Georeferencer, the evidence of which has to be intentionally sought out.

Finally, once a user assigns his or her desired number of control points to a satisfactory degree of accuracy, the “Preview Rectified Map” tab allows the result to be viewed as it would appear in a layer once digitized (see Figure 10). As in the Georeferencer, transparency may be adjusted, but the map is displayed on its original base map (in my case, Open Street Map) rather than in a Google Earth interface, though it may be exported as a KML file (among other options). The final rectified map, as seen in Figure 11, is certainly usable without further work, but in order to combine it with other neighboring maps and create an effective mosaic, cropping is required.

Cropping can be done simply with a tool that will be familiar to most users who have experience with a default photo editor, accessible under the “Crop” tab. The neighboring tab, “Align,” allows control points from one map to be dragged onto another. After completing a map warp to their satisfaction, or during the warp for that matter, users can see a history of their actions under the “Activity” tab, which is also visible to others interested in the warping process undergone by a particular map. This feature is reminiscent of those on Wikipedia and Google Refine in which editing history can be seen and mined for problems or previous versions. Overall, the Map Warper provides a broader range of tools than the Georeferencer, and is more consistent with the design of other open web projects and technologies.

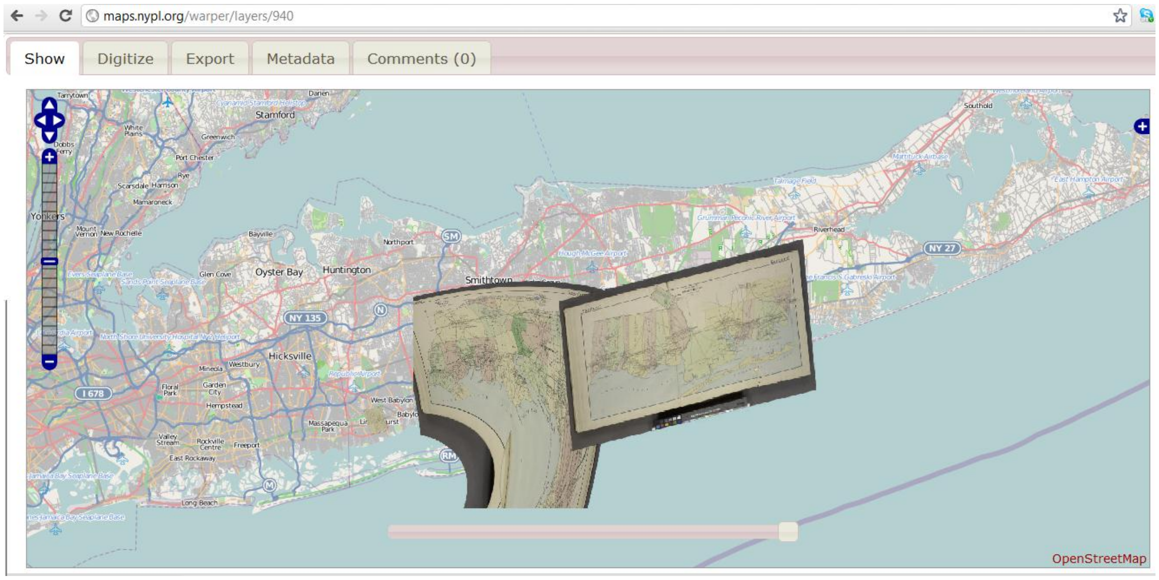

Figure 9. An inaccurately-rectified map (left). From maps.nypl.org.

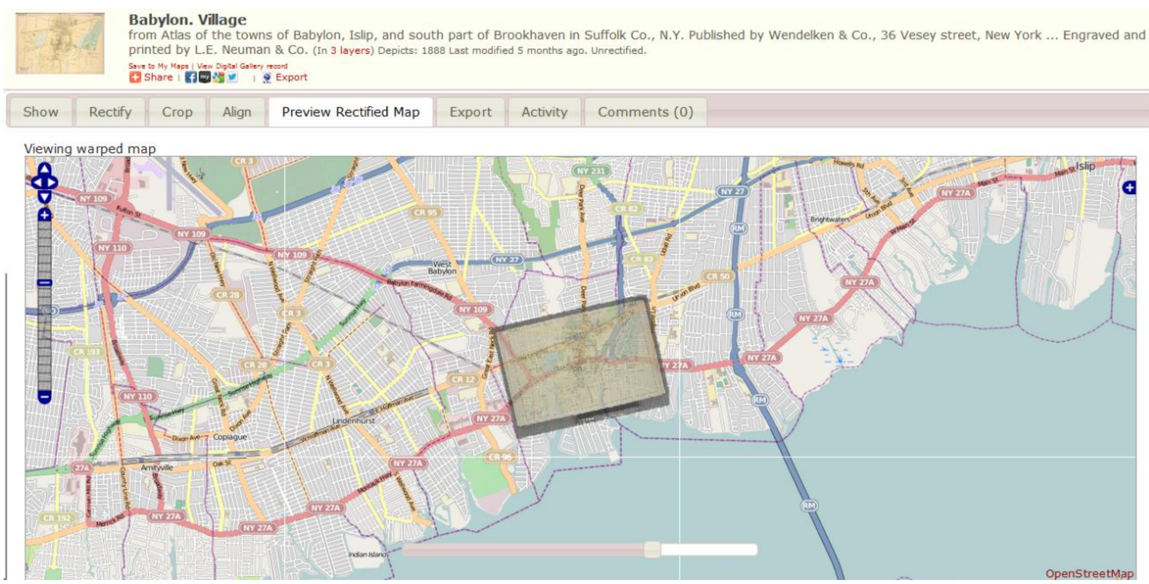

Figure 10. The "Preview Rectified Map" tab. From maps.nypl.org.

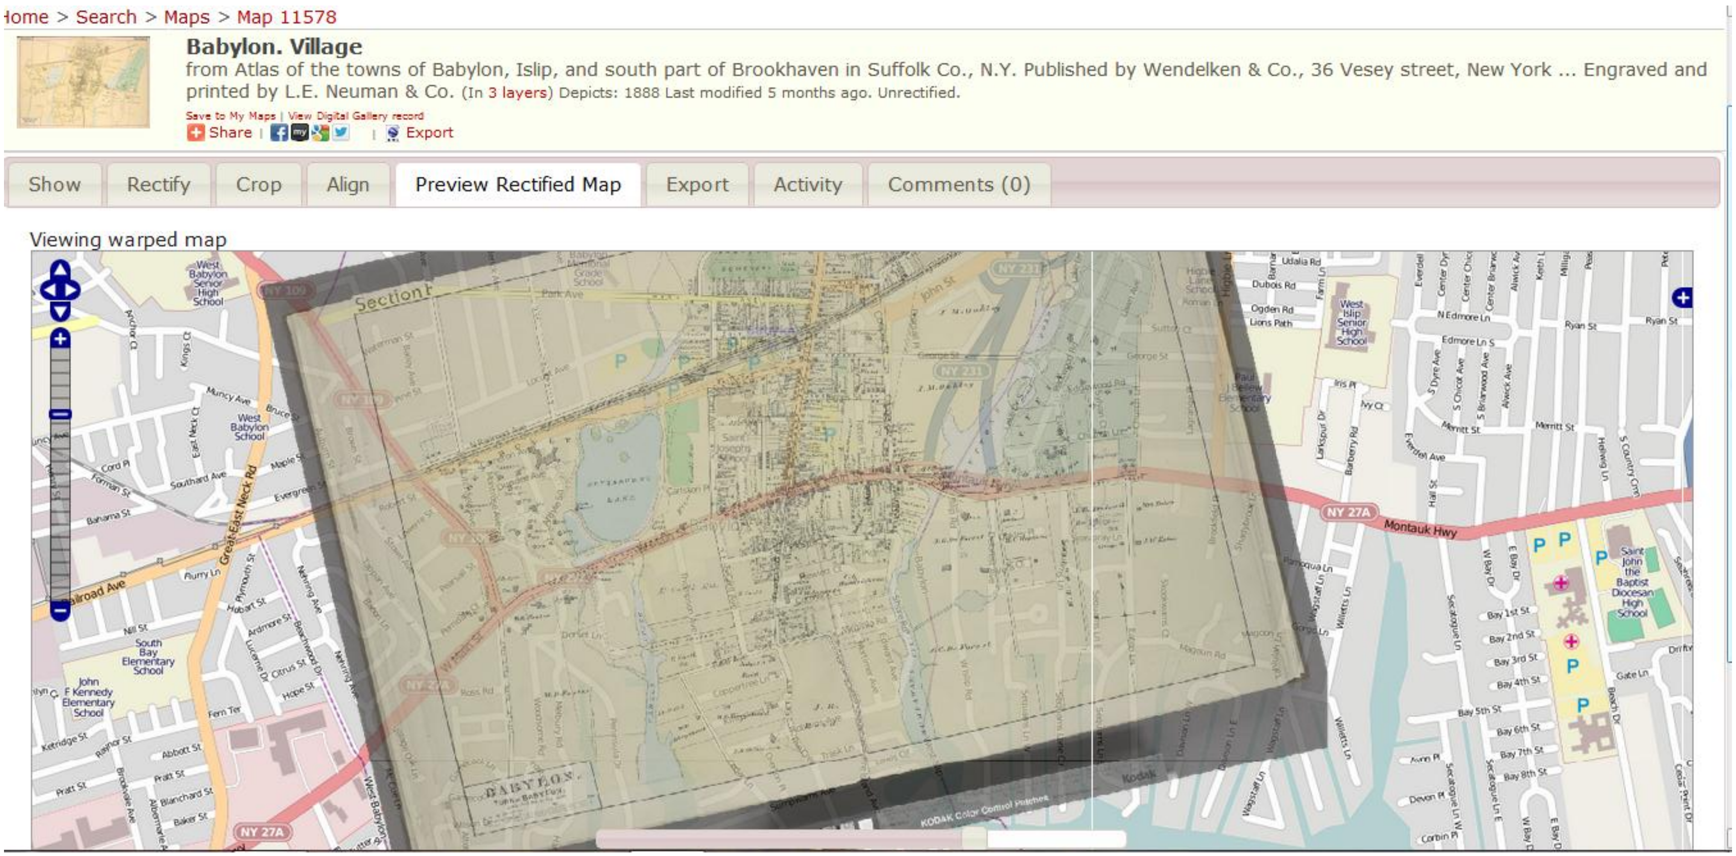

Figure 11.The transparency function on a warped map with uncropped edges. From maps.nypl.org.

The last component of the Warper is accessed by browsing for layers rather than maps on the main maps.nypl.org page. Once a map layer is selected (if complete, it will be a rectified, cropped, and mosaicked overlay of neighboring maps), the user has the option to "digitize." This is the point at which users can trace polygons onto the map, depending on the map feature of interest. For example, to trace buildings one selects "buildings" from a dropdown menu on the tracing tool and then traces the outlines of a school or other establishment. Once a polygon is complete, users can add attributes like name, building type, construction material, use, and others. This information can then be exported as a file type of the user's choice. While there is less data on the use and popularity of this aspect of the Warper, its availability adds even more value to an already powerful and useful tool.

Conclusions

From my investigation of two tools with similar purposes, the British Library's Georeferencer and the New York Public Library's Map Warper, I have gained a better understanding of the capabilities of geospatial crowdsourcing technologies as well as of the needs of users participating in these projects. Both projects involve the mass public georeferencing of maps in prestigious library collections and share a concern for usability, accuracy, and usefulness for researchers. However, the NYPL project design is more successful at mimicking other crowdsourcing projects and open web technologies. Furthermore, it adds value to the georeferenced maps produced by permitting further manipulation including cropping, creating multi-map layers, and tracing polygons with fields for attribute data. While the NYPL Map Warper could benefit from the addition of a gazetteer search function similar to that of the British Library Georeferencer, overall it is the more user-friendly and open tool available for rectifying digital maps.

So what are the implications for a crowdsourcing project to georeference historical census enumeration maps? First, of course, the original maps (on microform at the Library of Congress) would have to be digitized, a task unamenable to crowdsourcing. Once the maps are online, however, every subsequent step could potentially be delegated to public user-volunteers. As FamilySearch.org's census indexing project has shown, there is a vast user base interested in historical censuses and in family or local history, and this "crowd" could be tapped into to participate in a census georeferencing project. Such a project would additionally attract those interested in mapping in general, that is, the crowds that participate in the NYPL and British Library mapping projects, among others.

To best appeal to both of these user sets, a successful census-mapping project would need to combine elements of the NYPL mapping tool and the FamilySearch.org indexing system already familiar to genealogists and amateur users. While georeferencing a map takes longer than indexing, there are far fewer maps than there are total census records, so the two indexer-one arbitrator method could still work. Once two separate users had georeferenced a map or section of a map, the arbitrator could compare the two, possibly by aligning them as layers on a base map, and select the one with fewer errors and/or more control points. I would recommend requiring a much higher minimum than seen in either the NYPL or British Library projects, as three control points is not enough to make a map more than vaguely useful.

Furthermore, the NYPL method of cropping and building mosaic layers would allow users to recreate entire census wards using neighboring district maps. The tracing could remain optional, but it is certainly a valuable tool and one sure to be popular with genealogists and map lovers alike. Although family history researchers may not be as familiar with GIS and digital maps, they are frequent users of physical maps, which can pose numerous difficulties, particularly if scanned improperly into digital form. According to Michele D. Shular, genealogists "will seek out a variety of geographical materials to discover their family origins, locate ancestral residences, determine historical boundaries, identify name changes of villages and towns, as well as to examine ancestral migration trails and settlement patterns."[83] These materials may include "gazetteers, place-name dictionaries, historical atlases, and printed maps," and in some cases, "digitized images from scanned paper maps from a library's digital collection . . . or downloaded map images readily found on the Internet."[84] The ability to interact with a historical map in a manner similar to Google Maps would enable ease of map-reading and increase discoveries of historical and geographical interest among genealogists and amateur researchers. Furthermore, their experience with non-digital maps and with historical materials including censuses would also make them ideal georeferencers, whether or not they are well-versed in the process of georeferencing.

One major difference between the two library mapping projects and the FamilySearch.org indexing project is that the mapping goals are more open-ended while the indexing goal is finite. Even though the British Library and the NYPL do have reachable goals in terms of georeferencing their collections, users can continue interacting with the maps (adding control points, cropping, tracing) as much and for as long as they want. Both of these tools seem designed more for continual use and growth than for the somewhat limited goal of georeferencing a certain number of maps. Indexing, on the other hand, has a very definite end in sight, that being increased searchability of census records. The progress map showing the percentage of states complete and yet to be indexed, along with weekly status reports, keeps this goal in users' minds and pushes them further than if the project was continuous and never-ending. The speed with which records have been indexed and arbitrated is likely owing to the size, finiteness, and manageability of the goal.

How could a mapping project be made finite and manageable in the same way? In some sense, the usefulness of georeferenced maps depends on their being continuously available for manipulation and additions, rather than simply "read-only" in the sense of being merely viewable. On the other hand, getting all available maps georeferenced and thus prepared for further interaction is a goal unto itself and very similar to the indexing goal, considering the maps are also divided by state and could be calculated by percentage georeferenced. The same competitive push to complete the indexing of the 1940 census would therefore be achievable for georeferencing an entire census's maps.

Of course, such a project should not be begun without further investigation into the interests of genealogists and other users who participate in indexing. At the same time, it can be difficult to explain the concept and the benefits of georeferencing to non-geographers without providing an arena for them to try it themselves. There are numerous ways in which FamilySearch.org indexers could be introduced to georeferencing and other digital mapping elements in smaller projects than the georeferencing of an entire census. Whether starting with a state or even a ward, users could be eased into the process, and the potential for a larger project could be gauged simultaneously.

As georeferencing historical maps can be and has been crowdsourced effectively, and there is already an active group of volunteers interested in the census and in maps, a mapping project to georeference census enumeration district maps has the potential to succeed if managed effectively. As in other crowdsourcing projects, geographical or otherwise, a key element in success is to make the activity fun, competitive, and engaging, whether by openly acknowledging active users or by creating an environment that fosters self-fulfillment and community participation. As genealogists and family history researchers frequently rely on a community for information, references, and research tips, they are ideal candidates for voluntary participation in a census mapping project. They are interested in the material, and by georeferencing they are not only exploring and gaining familiarity with maps they will use in their work, but they are making these materials accessible for themselves and other genealogists in the future. Using a reward system like FamilySearch.org's, and developing an interface similar to that used by the New York Public Library for warping and tracing maps, a crowdsourced census mapping project is possible and likely to appeal to large numbers of conscientious users doing their own research. In partnership with organizations like FamilySearch.org and the NYPL that provide tools and materials freely and openly, a digital library could establish and maintain a project to make census maps accessible, usable, and free to a large population using online tools and volunteer participation.

NOTES

[1] Neale Stokes, *A Study in Digital Geospatial History*, master's thesis, University of North Carolina at Chapel Hill, May 2011: 1 .
[2] Ibid.
[3] Timothy J. Bailey and James B. M. Schick, "Historical GIS: Enabling the Collision of History and Geography," *Social Science Computer Review* 27, no. 3 (August 2009): 292.
[4] Ian N. Gregory and Richard G. Healey, "Historical GIS: Structuring, Mapping and Analysing Geographies of the Past," *Progress in Human Geography* 31, no. 5: 638.
[5] Ibid., 639.
[6] Ibid., 640.
[7] Ibid.
[8] Ibid.
[9] Ibid.
[10] Catherine A. Fitch and Steven Ruggles, "Building the National Historical Geographic Information System," *Historical Methods* 36, no. 1 (Winter 2003): 41.
[11] Ibid.
[12] Ibid.
[13] Ibid.
[14] Ibid., 42.
[15] Michele D. Shular, "Turning Genealogists onto GIS," *Journal of Map & Geography Libraries* 5 (2009): 55.
[16] U.S. National Archives, "Enumeration District Maps," accessed July 4, 2012, www.archives.gov/research/census/1930/enumeration-districts-maps.html#maps.
[17] Ibid.
[18] Ibid.
[19] U.S. National Archives, "1940 Census," accessed July 4, 2012, 1940census.archives.gov/.
[20] U.S. National Archives, "Official Map of Charlotte, N.C., 1935," Office of City Engineer, Charlotte, N.C.
[21] British Library, "About this project," *British Library Online Gallery—Georeferencing*, accessed July 2, 2012, www.bl.uk/maps/georefabout.html.
[22] Andrew Hudson-Smith et al., "Mapping for the Masses: Accessing Web 2.0 Through Crowdsourcing," *Social Science Computer Review* 27, no. 4 (November 2009): 524.
[23] Christian Heipke, "Crowdsourcing Geospatial Data," *ISPRS Journal of Photogrammetry and Remote Sensing* 65 (2010): 551.
[24] Hudson-Smith et al., "Mapping for the Masses," 525.
[25] Ibid.
[26] Ibid.
[27] Rebecca Federman, "Happy Birthday to…Us! A Year of Menus," *What's on the Menu?* (blog), April 20, 2012, http://www.nypl.org/blog/2012/04/20/happy-birthday-to-us-menus.
[28] "AHA Award Recipients," *American Historical Association*, accessed June 28, 2012, http://www.historians.org/prizes/awarded/RosenzweigWinner.htm.
[29] Federman, "Happy Birthday to…Us!", http://www.nypl.org/blog/2012/04/20/happy-birthday-to-us-menus.
[30] Ibid.
[31] Ibid.
[32] Ibid.
[33] "May 2011 Sparks! Ignition Grant Announcement—New York," *Institute of Museum and Library Services*, accessed June 28, 2012, http://www.imls.gov/news/sparks_award_announcement.aspx#NY.
[34] New York Public Library, *What's on the Menu?*, accessed June 28, 2012, menus.nypl.org.
[35] Federman, "Happy Birthday to…Us!", http://www.nypl.org/blog/2012/04/20/happy-birthday-to-us-menus.

[36] New York Public Library, "What's on the Menu?: News, Histories and Culinary Findings from NYPL's Collaborative Menu Transcription Project," *New York Public Library*, accessed June 28, 2012, http://www.nypl.org/voices/blogs/blog-channels/whats-on-the-menu.
[37] Karen S., comment, *What's on the Menu?* (blog), April 21, 2011 (1:34 p.m.), http://www.nypl.org/blog/2011/04/19/doin-dishes.
[38] Jessica, comment, *What's on the Menu?* (blog), April 21, 2011 (10:46 a.m.), http://www.nypl.org/blog/2011/04/19/doin-dishes.
[39] Elizabeth, comment, *What's on the Menu?* (blog), April 20, 2011 (4:20 p.m.), http://www.nypl.org/blog/2011/04/19/doin-dishes.
[40] "About," *Old Weather: Our Weather's Past, the Climate's Future*, accessed June 28, 2012, www.oldweather.org/about.
[41] *Old Weather: Our Weather's Past, the Climate's Future*, accessed June 28, 2012, www.oldweather.org.
[42] "How to Take Part," *Old Weather: Our Weather's Past, the Climate's Future*, accessed June 28, 2012, http://www.oldweather.org/how_to_take_part.
[43] "*HMS Devonshire*," *Old Weather: Our Weather's Past, the Climate's Future*, accessed June 28, 2012, http://www.oldweather.org/vessels/4f961b11a2fc8e2def0024bb.
[44] *Old Weather Forum,* accessed June 28, 2012, http://forum.oldweather.org/.
[45] "*HMS Devonshire*," *Old Weather: Our Weather's Past, the Climate's Future*, accessed June 28, 2012, http://www.oldweather.org/vessels/4f961b11a2fc8e2def0024bb.
[46] "Worldwide Records Indexing," *FamilySearch.org*, accessed June 28, 2012, https://www.familysearch.org/volunteer/indexing.
[47] "FamilySearch Posts 5 New States: 1940 Census Index Report—June 28, 2012," *FamilySearch Blog*, accessed June 29, 2012, https://familysearch.org/blog/familysearch-posts-5-states-1940-census-index-reportjune-28-2012/.
[48] "Worldwide Records Indexing," *FamilySearch.org*, accessed June 28, 2012, https://www.familysearch.org/volunteer/indexing.
[49] "FAQs–Frequently Asked Questions," *FamilySearch.org*, accessed June 29, 2012, https://indexing.familysearch.org/support/faq.jsf.
[50] elzabel, comment, *FamilySearch Blog*, accessed June 29, 2012, https://familysearch.org/blog/familysearch-posts-5-states-1940-census-index-reportjune-28-2012/.
[51] Elistariel, comment, *FamilySearch Blog*, accessed June 29, 2012, https://familysearch.org/blog/familysearch-posts-5-states-1940-census-index-reportjune-28-2012/.
[52] 1940 census map, *FamilySearch.org*, accessed June 29, 2012, https://familysearch.org/1940census/.
[53] Heipke, "Crowdsourcing," 552.
[54] Ibid.
[55] Ibid.
[56] Ibid., 553.
[57] Ibid., 554-6.
[58] Ibid., 556.
[59] Hudson-Smith et al., "Mapping for the Masses," 528.
[60] Ibid.
[61] Ibid., 528-9.
[62] Ibid., 531-3.
[63] Ibid., 533.
[64] Ibid., 534-5.
[65] Ibid., 536.
[66] British Library, "About this project," *British Library Online Gallery—Georeferencing,* accessed July 2, 2012, www.bl.uk/maps/georefabout.html.
[67] Ibid.
[68] Ibid.

[69] Ibid.
[70] British Library, "Placing historic maps," *British Library Online Gallery—Georeferencing*, accessed July 2, 2012, www.bl.uk/maps/index.html.
[71] British Library, "Detailed instructions: georeferencing historic scanned maps via the British Library's online Georeferencer tool," *British Library Online Gallery—Georeferencing*, accessed July 1, 2012, www.bl.uk/maps/georeferencer_instructions.pdf, 1.
[72] George Allen, "Plan for preventing the Traffic to the WHARFS AND PREMISES IN TOOLEY STREET being impeded by that proceeding TO THE FOUR RAILWAY STATIONS," *British Library Online Gallery—Georeferencing,* accessed July 1, 2012, britishlibrary.georeferencer.com/map/nJakXqlCzHVDLdEPkhXCqe/201207012332-mRBeKZ/.
[73] British Library, "Detailed instructions," 2.
[74] British Library, "Accuracy Analysis," *British Library Online Gallery—Georeferencing,* accessed July 2, 2012, britishlibrary.georeferencer.com/map/nJakXqlCzHVDLdEPkhXCqe/201207012332-mRBeKZ/accuracy.
[75] Matt Knutzen, "The New York City Historical GIS Project," New York Public Library blog, June 13, 2012, www.nypl.org/blog/2012/06/13/nyc-historical-gis-project.
[76] Matt Knutzen, "Drawing on the Past: Enlivening the Study of Historical Geography at maps.nypl.org," *NYPL Labs* (blog), February 3, 2010, www.nypl.org/blog/2010/02/03/drawing-past-enlivening-study-historical-geography-mapsnyplorg.
[77] Matt Knutzen, "Unbinding the Atlas: Working with Digital Maps," New York Public Library blog, January 10, 2012, www.nypl.org/blog/2012/01/10/unbinding-atlas-working-digital-maps.
[78] Ibid.
[79] Ibid.
[80] Ibid.
[81] Ibid.
[82] Ibid.; Knutzen, "Drawing on the Past."
[83] Michele D. Shular, "Turning Genealogists onto GIS," *Journal of Map & Geography Libraries* 5 (2009): 55.
[84] Ibid., 55-56.

BIBLIOGRAPHY

Bailey, Timothy J. and James B. M. Schick. “Historical GIS: Enabling the Collision of History and Geography.” *Social Science Computer Review* 27, no. 3 (August 2009): 291-6.

British Library. “Detailed instructions: georeferencing historic scanned maps via the British Library’s online Georeferencer tool.” *British Library Online Gallery—Georeferencing*. Accessed July 1, 2012. www.bl.uk/maps/georeferencer_instructions.pdf. 3 pages.

British Library. *Georeferencing. British Library Online Gallery*. Accessed July 1, 2012. www.bl.uk/maps/.

FamilySearch.org. Accessed June 28, 2012. https://www.familysearch.org.

Fitch, Catherine A. and Steven Ruggles. “Building the National Historical Geographic Information System.” *Historical Methods* 36, no. 1 (Winter 2003): 41-51.

Gregory, Ian N. and Richard G. Healey. “Historical GIS: Structuring, Mapping and Analysing Geographies of the Past.” *Progress in Human Geography* 31, no. 5: 638-53.

Heipke, Christian. “Crowdsourcing Geospatial Data.” *ISPRS Journal of Photogrammetry and Remote Sensing* 65 (2010): 550-7.

Hudson-Smith, Andrew, Michael Batty, Andrew Crooks, and Richard Milton. “Mapping for the Masses: Accessing Web 2.0 Through Crowdsourcing.” *Social Science Computer Review* 27, no. 4 (November 2009): 524-38.

New York Public Library. *NYPL Map Warper*. Accessed July 1, 2012. maps.nypl.org/warper.

New York Public Library. *What’s on the Menu?* Accessed June 28, 2012. menus.nypl.org.

Old Weather: Our Weather’s Past, the Climate’s Future. Accessed June 28, 2012. www.oldweather.org.

Shular, Michele D. “Turning Genealogists onto GIS.” *Journal of Map & Geography Libraries* 5 (2009): 55-71.

Stokes, Neale. *A Study in Digital Geospatial History*. Master's thesis, University of North Carolina at Chapel Hill (May 2011). 40 pages.

U.S. National Archives. "1940 Census." Accessed July 4, 2012. 1940census.archives.gov/.

- - -. "Enumeration District Maps." Accessed July 4, 2012. www.archives.gov/research/census/1930/enumeration-districts-maps.html#maps.

- - -. "Official Map of Charlotte, N.C., 1935." Office of City Engineer, Charlotte, N.C.