Citizen science through old maps: Volunteer motivations in the GB1900 gazetteer-building project.

Paula Aucott [1] *

email: paula.aucott@port.ac.uk, OCRID: 0000-0002-1637-2672

Humphrey Southall [1]

email: humphrey.southall@port.ac.uk, OCRID: 0000-0002-4406-1425

Carol Ekinsmyth [1]

email: carol.ekinsmyth@port.ac.uk, OCRID: 0000-0001-8789-5027

Affiliation:

[1] Department of Geography, University of Portsmouth

Postal Address: Department of Geography, Buckingham Building, Lion Terrace, Portsmouth, PO1 3HE, UK

* Corresponding author

Abstract

The GB1900 project transcribed almost all text on 1:10,650 mapping covering Great Britain, published *circa* 1900: 2.6 million geo-referenced text strings, so possibly the largest specifically historical gazetteer. Nearly 1,200 volunteers made 5.5 million transcriptions, including "confirmations". This paper describes the project's interaction with online volunteers and then presents their experience, as recorded through the online system itself, six in-depth interviews and 162 responses to an online questionnaire. We find that, unlike volunteers in physical science "citizen science" projects, they were motivated by personal interest in the maps, in places that held meaning for them, and in how places had changed. These conclusions enable us to offer suggestions for volunteer recruitment and retention in similar future projects.

Citizen science through old maps: Volunteer motivations in the GB1900 gazetteer-building project

Introduction

The GB1900 project crowd-sourced transcription of almost all text strings on the six-inch (1:10,560) Ordnance Survey second edition County Series mapping published 1888-1913, creating not just a gazetteer but an intricate geographical portrait of Britain. It launched in October 2013 as Cymru1900, led by the National Library of Wales and the Royal Commission on Ancient and Historical Monuments of Wales (Ell, Hughes, and Southall 2016). It re-launched as GB1900 in September 2016, using mapping of the whole of Great Britain digitized by the National Library of Scotland, and software revised and hosted by the University of Portsmouth. All partners contributed substantial amounts of staff time and various existing resources but the only external funding was from the Welsh Assembly, to develop the original Cymru1900 software in 2012-13.

This history and the software are detailed elsewhere (Southall et al. 2017), so this paper is entirely focused on the project's volunteer transcribers, and on our experience working with them from the re-launch to the site's closure in January 2018. A third paper will describe the resulting data sets and exemplify their use.

The project was entirely online so volunteers were spread throughout Britain and beyond. We have four sources of information about them and their motivations. Firstly, they were asked their gender, age, and how they heard about the project when they first registered, but low response rates meant results were problematic. Secondly, every transcription included an internal user ID we can link to an e-mail address, plus a geographical coordinate from the map, so individual transcription histories have been assembled: how many, when, and where on the map. Thirdly, 162 volunteers participated in an online questionnaire, and their individual responses have been linked to their transcription histories. Finally, in-depth interviews were carried out with six of the most active volunteers, again linked to and based on transcription histories, including maps of their transcriptions week by week.

GB1900 was a "citizen science" project, in that it mainly asked volunteers to work with sources we provided, rather than contribute their own knowledge. The next section reviews the experience of other citizen science projects, mainly in the physical sciences, focusing on volunteer motivation. We then describe our own work recruiting and maintaining contact with volunteers, creating resources additional to the transcription web site inherited from Cymru1900. Three sections then present results: volunteer behavior within the site, then information from the interviews and questionnaires about who the volunteers were, and finally their

motivations. The concluding section draws a series of lessons for other historical citizen science projects.

The extensive digitization and online publication of historical maps in recent years has made them enormously more accessible to a wide audience, expanding and democratizing usage (Dodge 2017), but this covers both passive browsing and a deeper engagement. This paper analyses one form of deeper engagement, exploring how contributors based their selection of areas to transcribe on a personal knowledge of, or family association with, particular places; and then deepened their engagement with these places through the transcription process. Because we can draw on both very detailed data on individual's working patterns within the system and depth interviews "in real life", this paper explores the motivations of map users in unparalleled detail.

Review

The most tedious aspect of much historical research is the transcription of source materials. Few academic historians can employ research assistants, so having transcription done for free is enticing; but why would anyone do this for us? This section reviews existing research into volunteer motivation, and one central question is whether motivations differ between physical science projects and the three historical examples.

Like the present paper, all studies reviewed here were relatively small scale investigations ancillary to large scale crowd-sourcing, using a variety of research tools. These include a link on a website homepage, email invitations, direct interviews or a post in an online forum connected to the project, or some combination of these methods. In some cases the volunteer could directly benefit from participation through a gift card (Reed et al. 2013) or the chance to win a book in a draw (Raddick et al. 2010). Despite incentives all surveys received a low response rate or dealt with low group numbers.

Researchers have generally found it easier to study transcriber behavior than motivation. Transcribe Bentham has been running continuously since 2010, and aims to transcribe the hand written manuscripts of British philosopher and reformer Jeremy Bentham. In 2012 they had 1207 registered volunteers, of whom only 259 had actually contributed. Seven transcribers (0.6%) had done 70% of all transcription, and the three most active volunteers were spending over eight hours per week on the project (Causer and Wallace 2012).

Similarly, the Australian Newspapers Digitization Program (Trove) gives volunteers the chance to correct and tag text collected automatically via Optical Character Recognition technology from scanned images of historical newspaper articles. Over a three month period they detected a daily pattern of use, rising from a low point around 5am to highs at 3pm and 9pm, with a slight dip between 4pm

and 6pm, presumably as household activities took precedence, and a final drop after 11pm (Holley 2009).

On-line surveys enable citizen science projects to study their contributors at low cost, but response rates are usually low. 10,992 out of 174,764 Galaxy Zoo volunteers (6.3%) responded to their online survey: 82.1% were male, and the over-representation of males, relative to census data, increased in older age groups (Raddick et al. 2013). A larger study covering multiple Zooniverse projects found that most volunteers were employed and had a university or college education (Reed et al. 2013). Similarly 97% of the Transcribe Bentham volunteers had a degree level education and only around 10% were retirees, but unlike Galaxy Zoo two thirds were female (Causer and Wallace 2012).

Crowson and Fagnot (2008) conclude that most initial contributors to crowd-sourcing projects like Wikipedia are motivated by curiosity, but individuals become more involved when they find they agree with the ideology behind the project and acquire a sense of social obligation and associate themselves with a group membership. Whilst GB1900 did not have extended contributor facilities in the sense that Wikipedia does, it did have 'committed' volunteers who have contributed significantly more than the average sustained user. These committed volunteers interacted further with the project team and offered to be involved in follow up procedures following the end of the transcription phase.

Although the Transcribe Bentham example cited above is particularly extreme, the general experience is that a relatively small subset of volunteers do most of the transcriptions, so motivational surveys should arguably target them. Nov et al. (2013) included an online questionnaire on the homepage of a website collecting sightings and photographs of butterflies and moths in North America (BAMONA). They obtained 462 responses but focused particularly on the 73 who had been active in the previous month. They concluded that those who identified with the project most closely were encouraged to contribute by a mission statement included on the website homepage, but those only loosely linked were discouraged as it made them feel even less connection with the project. Despite over 3000 volunteers participating in this project, the pool of potential participants was skewed by the need for specific expertise on butterflies and moths.

Volunteers active in the thirty days prior to the survey were also identified on the Stardust@home project, which allowed volunteers to classify images from a NASA spacecraft (Nov, Arazy, and Anderson 2011a). On this project the top motivation was enjoyment followed by agreement with the project goals (Nov, Arazy, and Anderson 2011a). The majority of surveys, however, have analyzed overall motivations among their volunteers. Galaxy Zoo found the majority of their volunteers identified two or more motivations for participating (Raddick et al. 2010). Zooniverse volunteers in general have been categorized having three main motivations: satisfaction from assisting with the tasks, utilizing and enjoying the

website interaction, and social engagement from interacting with other contributors (Reed et al. 2013).

Galaxy Zoo researchers interviewed 22 volunteers (1.4%) of the 1,336 they invited, to explore the motivations felt by the overall volunteer community (Raddick et al. 2010). They then used an online questionnaire to find out which motivations mentioned by the interviewees mattered most to a larger group. Their highest ranked motivation was a desire to "contribute", but there was a gender divide on the importance of other factors. This desire to contribute to knowledge is an important outcome for citizen science generally (Raddick et al. 2013). For the historical projects an interest in the past, furthering their own research, contributing to the greater good, personal challenge and pleasure are all major motivations for their volunteers (Causer and Wallace 2012; Holley 2009).

One study did directly compare motivations between science and history projects. Romeo and Blaser (2011) studied volunteers in two Zooniverse projects, Solar Stormwatch where volunteers catalogued solar storms erupting from the sun, and Old Weather, where volunteers transcribed handwritten ships' logs from the nineteenth and early twentieth centuries. They concluded that motivations were not very distinct, but Old Weather focuses on climate change and lacks the potential local focus of GB1900.

The GB1900 system followed a light-weight peer production model as defined by Haythornthwaite (2009) where a large group of contributors follow a set of rules to produce a common outcome without requiring either a large or a long term commitment. This is comparable to many Zooniverse projects, which follow a similar model, often on a larger scale. Based on the feedback from Old Weather volunteers obtained via an online survey, Eveleigh et al. (2014) recommend encouraging an irregular participation by many contributors over an extended period by allowing volunteers to be able to choose what they work on, how much and also enable them to track their personal progress. Nov and colleagues (2011a) argue that retention of volunteers is best achieved through regular communication and having a progressive levels of difficulty for volunteers as they become familiar with the project. It is also suggested many volunteers maintain registration to receive updates which may rekindle their interest, even when they are not necessarily actively contributing (Eveleigh et al. 2014).

Nov, Arazy and Anderson also suggest that projects need to "develop enjoyable contribution mechanisms" (2011b, 254). It has been argued a project should have a fun social interface, employing a gamified approach, as demonstrated in the early Solar Stormwatch project and developed further in Old Weather (Romeo and Blaser 2011). However, it should also be noted that volunteers are concerned about the quality of the work they do, so the tasks should not be trivialized. Eveleigh et al. (2014) recommend that a project provides a means for volunteers to learn through

test cases in order to gain confidence in the task methodology and the quality of their contribution.

The importance of the community aspect of projects varies according to the project content. Although there were opportunities for community interaction with other volunteers within Transcribe Bentham, via a forum, private messaging, twitter and Facebook, they were very little used. There is also little evidence of collaborative working, with three quarters of all manuscripts only having a single editor, despite there being a technical capability for multiple editors, suggesting the task was viewed by volunteers very much as an personal task and not a community task (Causer and Wallace 2012). This agrees with findings from Galaxy Zoo where "community" was rated lowest amongst the motivations identified (Raddick et al. 2013), but contrasts with the Trove project which found giving volunteers more control helped to build trust and commitment with their volunteers and encouraged them to feel part of a community (Holley 2009).

Supporting GB1900 volunteers online

The GB1900 web site launched on 22nd September 2016 and ran for sixteen months, closing on 30th January 2018. As explained in Southall *et al* (2017), the "site" was a single page containing a JavaScript app, inherited from Cymru1900 but modified to encourage confirmation of existing transcriptions, and to be more fun. The app included a registration form, a brief tutorial in English and Welsh and basic contact

details, but no real facilities to support volunteers, or report progress beyond basic transcription counts. Any change required a programmer.

Cymru1900 had gathered few "confirmations" so few transcriptions were ever finalized. In GB1900, the revised leader board based rankings on whichever was larger of a volunteer's counts of new pins and of confirmations, and pins changed color when a confirmation was added. Figure 1 presents weekly counts from both projects, using a four-times scale for Cymru1900, firstly showing that these small revisions succeeded in encouraging volunteers to balance the two tasks.

It also arguably shows that we succeeded in sustaining volunteer activity over a lengthy task through various activities. It is conspicuous that the creation of new pins on Cymru1900 fell away rapidly after week 18, reflecting growing difficulty finding un-transcribed text. The GB1900 time series shows signs of a similar decline starting then, but several upturns follow.

[Figure 1 near here]

The strongest upturn for GB1900 usage resulted from our being approached by the popular genealogy magazine, *Who Do You Think You Are*, to be one of six online projects which they would promote in a "Transcription Tuesday" event in January 2017: different magazine writers each promoted their favorite project and activity that day was then compared. We were the editor's choice and it became the day

GB1900 had the most transcriptions made, and the most volunteers participating. Figure 1 shows a sustained raising of activity following week 17.

The spike at around week 40 is unexplained, but the spike in week 55 was from a new scoring system we dubbed "Doubloons". An inherent defect of our leader board was precisely that it gave equal weight to new pins and confirmations, even though we inevitably needed some additional confirmations because either the initial transcription or the first confirmation was wrong: a year into the project, most "leaders" could improve their position only by finding new pins, when we mainly needed confirmations. In the final months we therefore allocated ten million notional Doubloons each week, divided equally between new pins and confirmations, and then allocated them to individual volunteers in proportion to the number of *confirmed* new pins and confirmations done that week. The "reward" for finding un-transcribed names therefore grew as the task became ever harder. The existing leader board was part of our implicit contract with our volunteers, so the Doubloon rankings ran in parallel.

Four additional support facilities were added to support volunteers, each necessarily working through a separate online system rather than through the limited GB1900 site.

Firstly, starting in March 2017 we began sending a monthly e-mail newsletter to all registered volunteers. These reported progress, gave tips and towards the end

included sample results and the Doubloon ranking. We used the free MailChimp service, and only 85 volunteers unsubscribed from it, suggesting most retained at least a passing interest in the project.

Secondly, we used the Zoho platform, again free, to create a "support site", accessed through a 'help' button added to the GB1900 homepage (Aucott et al. 2016). This initially included a more detailed tutorial and a set of Frequently Asked Questions. Zoho is basically a help desk, so volunteers could submit questions to the project team and access an archive of answers. As GB1900 progressed, sections for known issues and progress tracking were also added.

[Figure 2 near here]

Thirdly, the GB1900 system displayed the pins marking transcriptions only when tightly zoomed in to particular areas, as seen in Figure 2, so it provided no overview of progress; and neither could it display the transcribed names. Initially progress was visualized via traditional GIS software and released as map images on the support site, but the amount of data involved quickly became unmanageable. The National Library of Scotland (NLS) then created a quite separate "progress map" system within their maps.nls.uk site, using MapServer software to pre-render tiles at different zoom levels. This was based on monthly data dumps from GB1900, and towards the end the maps of "pins needing confirming" were more useful than "all pins". Individual pins were clickable, revealing the text already transcribed for that

14

pin, and NLS also added basic name searching, so it started to be usable as a gazetteer (Fleet 2017).

Fourthly, we were strongly criticized at launch for promising that data would be available under a Creative Commons license, but having no data actually available. In March 2017, we therefore made an inevitably incomplete data dump available via our Vision of Britain system, which already included a facility enabling us to track downloads, and ask downloaders a few questions (Great Britain Historical GIS Project 2017). By 23 February 2018 this file had been downloaded 188 times, although the CC0 license means it can be freely passed on or included in other download systems so we cannot measure total usage. In July 2018 this was replaced by the final versions of the data.

Overall, we utilized three additional web sites to overcome the limitations of GB1900: the support site on Zoho, the progress maps and the download system. We also used e-mail both for the newsletter via MailChimp and through enabling individual volunteers to contact us by publicizing the gb1900@port.ac.uk address. Conversely, our social media presence was limited. We chose not to create a Facebook page as, to be effective, it would have needed more regular monitoring than our resources permitted. While we did not have a dedicated Twitter account for this project, we did publicize progress via two existing Twitter accounts,

@gbhgis and @natlibscotmaps and through the personal accounts of some team members.

Volunteer behavior on GB1900

Understanding volunteer *motivations* required the depth interviews and questionnaire findings presented below, but the web site itself provided much information about volunteer *behavior*. In particular, to ensure that confirmatory transcriptions were independent of the initial transcriptions, volunteers were required to register on the system, enabling each transcription record to include a user ID as well as the exact date and time it was created; and of course each transcription also had a precise location on a map of Britain. When volunteers registered they accepted a "User Agreement", modelled on one from Galaxy Zoo, which assigned copyright in their work to us and also granted us permission to track individual activity. They were also asked three simple questions about their age group, gender and where they heard about the project, but there was a very high rate of non-response, so for example 16.6% of volunteers were female, 15.3 % were male and the rest unknown.

The actual user IDs within GB1900 are arbitrary 30 character alphanumeric strings, and, as they do not identify people, they now form part of the final dataset. During registration volunteers supplied an email address. In thirty cases, where an individual registered more than once using the same email address, we have

grouped together their contributions under a single user ID for each person. However, we did not merge together data for the three cases where we happen to know the same person used more than one email address, as there may well be others. The GB1900 data includes all data gathered by the earlier Cymru1900 project, including the volunteer registration details. The system inherited from Cymru1900 also enabled volunteers to log-in using Facebook accounts, but Facebook discontinued the relevant programming interface soon after we launched, and work by those volunteers is excluded from this analysis.

Although volunteers had to log-in, the web site tracked them via cookies and never required them to log-out, so we could not directly measure time spent on the system. We therefore devised the following methodology to approximately measure the time periods volunteers actually worked. Firstly, we group together individual transcriptions into "blocks", within which the gap between adjacent transcriptions cannot be more than five minutes. We then group blocks into "sessions", within which the maximum gap is an hour; so total time worked is the sum of block lengths, but sessions can include tea breaks. We also group volunteers by when they first registered: joining for Cymru1900; from launch to the end of 2016; and then the four quarters in 2017, including anyone joining in January 2018 in the final one.

There were 1,199 identifiable individual volunteers, two thirds (66.2%) registering in GB1900's initial four months. 310 volunteers registered but made no

transcriptions, representing just over a quarter (25.9%) of all registrations, but over half of those without transcriptions (166) signed up for Cymru1900 rather than GB1900, representing 41.1% of the total Cymru1900 registrations. The percentage of non-transcribers was even higher (53.8%) in the final period from October 2017, when only the dedicated could find new names to transcribe. Table 1 shows the number of "sessions", as defined above, broken down by period as a percentage of all those registered during each period. This shows the number of sessions did not vary significantly between volunteers who registered during different periods, but that those joining later were less likely to have multiple session counts, presumably because of the difficulty in finding targets to transcribe.

[Table 1 near here]

Overall 2,666,341 pins were created, of which 106,339 have been identified as errors. 22,733 of the remainder (0.89%) required further offline checking because they had conflicting transcriptions, and 4,673 (0.18%) because they lacked confirmatory transcriptions. A common rule of thumb for crowd-sourcing is that 20% of contributors will be responsible for 80% of the work, but GB1900 was more extreme: the top 52 contributors (5.3%) did 80% of the work, and Figure 3 shows that almost all transcriptions (97.4%) were done by the top 200. Over half were done by the top 14 volunteers, between them entering 2,818,265 text strings, and the single most active volunteer did 393,701 transcriptions, 7.2% of the total. For

the first year the project was live the average active number of volunteers per day was 33, creating on average 6,441 new pins and confirming an average of 6,506 pins per day between them. As shown in Figure 1, in the final four months most transcriptions were confirmations and these became progressively harder to find. Another way of looking at this is by computing the average number of transcriptions per estimated hour worked: this was always between 269 and 307 from launch to June 2017, but dropped to 157 in October 2017 and 96 in January 2018. All of these rates may seem high, but bear in mind that many transcriptions were just of initials, such as "W." for well and "F.P." for footpath.

[Figure 3 near here]

Excluding those volunteers who worked only on Cymru1900, 415 volunteers only worked during their first month after registration, and a further 100 worked only for three months or less. Another 234 worked on the project over the remaining time, 53 between 3-6 months, 48 volunteers each for both 6-9 and 9-12 months and a further 85 for twelve months or longer. The strongest commitment was shown by the volunteer responsible for over half the transcriptions in Cymru1900 who remained active in GB1900 until June 2017.

[Figure 4 near here]

Figure 4 shows when volunteers did their work: Sunday was the most popular and Saturday least, but there was a fairly regular spread over the week. Time of day was

19

unsurprisingly more skewed, with more in afternoons than mornings, and an evening peak. Some activity may be by volunteers in other time zones.

Who volunteered?

Towards the end of the live project the project team decided to contact specific volunteers to discuss their personal interaction with the GB1900 project. Potential candidates were selected by their position in the top fifteen on the leader board, all having transcribed at least 60,000 text strings each. Eleven volunteers were invited to participate and six agreed to engage in these in-depth telephone interviews, representing 54% of those invited. All interviews were audio-recorded, lasted between thirty and sixty minutes each and the lead author was involved in all six interviews.

Before each interview, we sent the interviewee a personal "profile" extracted from the transcription data, showing the distribution of their transcriptions by time of day, day of the week and month of the project, and also a map showing the locations of their initial transcriptions; two of these maps are included below. The discussion topics included how they began work on GB1900; their personal contribution using the profile and map; what they felt they got out of participation and what had annoyed them about it. We also found out more about their personal history, and explored their experience of other online projects.

Following on from the interviews we also ran a questionnaire survey of volunteers for 13 days in February 2018 after the end of the project. Invitations were sent to all 1,057 volunteers who had registered for GB1900 directly and who had not previously unsubscribed from the project mailing list. Email invitations were sent out via SurveyMonkey enabling individual responses to be tracked, and linked via email addresses to volunteers' transcription histories; so as with the interviews we did not need to ask about their contribution.

The questions were based around the same themes as the in-depth interviews and the multiple-choice categories reflected the results of those interviews and our observations of comments on social media. There were twenty-nine questions with selectable multiple-choice answers, including twenty Likert scales. Every question provided a non-response option, either 'decline to say' or offer 'no opinion'. The final question was a free text box for any additional comments the volunteers wished to make about the project.

A total of 162 or 15.3% of all volunteers invited to participate actually started the survey, and only four of these participants (2.5%) did not complete it. Participants included three of the interviewees. Responses were received from volunteers who joined at all stages of the project, with 24 from Cymru1900, 72 during the first activity period up to the end of December 2016, 44 in the first three months of 2017

and then small numbers from the final three periods: nine, eight and five respectively.

Of the total 162 who responded to the survey, 12 had registered but not completed a single transcription, 15 had completed only a single session, 11 completed two sessions and 12 completed three sessions. The remaining 112 had completed four or more sessions, with the most being 1300 sessions. Of the 158 who answered the gender question there was an almost equal split between male (51.3%) and female (46.2%) respondents, with only 2.5% choosing not to say. This result is similar to the partial data on gender from our initial registrations, and contrasts with findings from other projects, reviewed above, where one gender usually predominates.

Participation of different age groups varied by gender. 42.5% of the female respondents were aged 50-64 and represented the largest grouping of females while 38.4% were 65 and over. In comparison only 29.6% of males were aged 50-64 and by far the largest group 60.5% were aged 65 and above.

Among the 158 respondents the largest group by employment status grouping were the retired, at 56.8%. Surprisingly those working full or part-time came second (27.2%), and the semi-retired a distant third (6.8%), perhaps suggesting they might consider themselves to have less time available. Each employment status grouping had participants in each of the age groupings over 24, except the two students who were both in the oldest grouping, and the semi-retired who started at 50. Of those

who were retired, 20 were in the 50-64 category, but only 22 retirees did three or less sessions. In contrast all of the interviewees who accepted our invitation were male and either retired, or self-employed and approaching retirement. Interestingly two of them, Interviewees B and C, had worked as part of a group focused on Scotland and had existing links to the National Library of Scotland.

Only 21 of the total 162 survey participants do not engage in any kind of historical research, they were almost equally split by gender and of these only five had one or less sessions. Only one survey respondent declined to answer this question. This question allowed for multiple selections and of those that did participate, the highest number of survey respondents (105) participated in family history and nearly half did local history (75). Interestingly, family history has a higher number of identified female participants (52) over male (46), in contrast to local history where almost double the number of males (47) are involved compared to females (23). The "other" category was next (33) but the survey did not allow respondents to explain what kind of history "other" was. 17 were involved in "one-place studies" (https://www.one-place-studies.org/) and equal numbers (12) were involved in taught and industrial history; notably all but one involved in industrial history were male. This tallies with the interviewees who between them noted family, local and industrial history as part of their interests.

Distribution in terms of age was fairly uniform across each of the types of history plus the none category, with the number of survey participants increasing with age and the 65 or older group participating the most in historical research. The one exception was one-place studies, meaning those interested in the history of one particular place, where more than twice as many people aged 50-64 (10) undertook this kind of research than in any other age bands, albeit these are small numbers compared to some of the other categories. Of those that participated in family history 17.1% had one session or less on GB1900, and 35% had three or less sessions, while for local history these proportions were lower, 12.4% having one or less session and 21% having three or less sessions.

Excluding those who could not remember, the results for discovery of the project from the 162 survey respondents are quite different from the registration results. The largest number, 16.7%, found out about it from a genealogy magazine or website, the second most popular source being somebody they knew and thirdly Transcription Tuesday, although local libraries and archives continued to be the least likely. Of those interviewed three learnt of the project from people they knew, two via a genealogy website or magazine and one from a blog.

Volunteer motivations

Interaction and collaboration were key themes in the interviews. The two who worked as part of a team clearly felt they benefited from the social interaction and

teamwork element of working together. The others interviewed all expressed an interest in collaboration, even though they did not personally know anyone else who was participating in GB1900. The newsletters were considered a helpful part of the interaction, but some interviewees felt they needed to be sent more frequently. The significant time input by the interviewed volunteers obviously led them to expect a greater level of commitment from the project team in return. All interviewees had initiated some direct contact with the project team, often making suggestions and offering feedback. They clearly felt they had a relationship with us and they liked that interaction as it helped them feel like part of the team.

Interviewees reported that the project offered them a productive way to fill their time, an alternative to TV, another hobby or when ill health curtailed other activities. They all liked feeling that they were able to contribute to something useful in the long run, although it did not necessarily matter if it was an academic project or not. They liked the fact they could participate from home, and whenever they felt like it. There were no pre-requisites to involvement, and no push to do a minimum amount, so they liked having the freedom to participate as much or as little as they chose:

> [I] retired five years ago, and then, I suppose, left at a bit of a loose end, wondering what to do … I rather liked the idea of doing some voluntary work, which I had never previously done (Interviewee B).

[Table 2 near here]

Of the five Likert scale responses visible in Table 2, "I liked looking at the old maps" provoked the strongest positive correlation with 70.9% of the 158 respondents choosing "strongly agree" and none chose to "disagree" or "strongly disagree". Similarly all of the interviewees listed a love of maps as a key motivator for both their initial participation and their continued interest. In talking about his motivations in taking part to the project, Interviewee E, for instance, said: "I thought, ooh, that covers two bases: I'm interested in researching my own family history, and I have a love of maps anyway, so let's go for it".

For some the project sparked a new interest in history, while for others it continued an existing hobby;

> I just spotted GB1900 and the maps have been a minor obsession of
> mine for quite a long time as well and the historical and also the
> geographical and also the language implications, all of those things said
> ah yes, that's just the sort of thing that I'd like to be involved in
> (Interviewee A).

Over half of the 162 survey respondents (51.2%) identified where they live as the starting point for their GB1900 experience. A further 19.8% chose where they were brought up, another 14.8% where family members or friends live and 5.6% from holiday memories. Totaling 91.4%, this underlies the importance of the personal

26

connection to the maps for the participants. Of the three survey respondents who used the random place selection offered by GB1900, two joined in period 2 at the start of 2017 and had a significant number of sessions each and the other joined in period 4, and had six sessions.

Most interviewees also started with an area they knew well. They had a personal and local connection to the places on the map, through familiarity with their own local area or that of a family member now, or in the past. Figures 5 and 6 show the distribution of transcriptions over the first two months after registration for two of our interviewees, with the darkest dots representing the earliest data. Interviewee A started with places with strong family connections; where his grandfather grew up near Darwen, then his own home town of Wigan and current home in Liverpool and he then moved on to holiday locations and areas connected to his wife's family. Similarly Interviewee E started where he lived in Worcestershire, then jumped to Norwich, in eastern England, where he was born and extended across Norfolk where his family were from.

[Figure 5 near here]

[Figure 6 near here]

They all enjoyed looking at the local changes, from how the place was back in 1900 to visualizing how the place had changed through their own knowledge or using modern technology. Interviewee D notes "what I tend to do sometimes is actually

having a second page on my browser open, with … the modern mapping, to see … the difference between … 1900 to what … in the modern day".

Once these familiar places were completed the interviewees moved on to random areas, using them as an opportunity to learn about somewhere new. Indeed, the discovery element of learning something new acted as a motivator for most. Interviewee F, for instance, commented "You don't know what you're going to come across next. Some interesting place or some historic … I mean, I was constantly amazed".

The interviewees were all motivated by an enjoyment of the experience and in some cases they felt it became more like an obsession or addiction. Some asserted that they found working on GB1900 'very addictive' (Interviewer D). They appreciated the visual nature of the system, the instant gratification of the color-coded pin system. This allowed them to build a sense of satisfaction, as they completed an area they felt both pride and achievement. There was also a definite interest in accuracy among them, perhaps stemming from their career histories, where many were involved in jobs focused on detail and accuracy. This manifested as both an initial hesitancy to participate due to concerns about doing the task correctly and later a frustration with other volunteers not following the instructions correctly. Interviewee C expressed this frustration when he said "we are trying to get a good job done; I mean, that's what this is all about".

Interviewees were divided about the competitive element. Some were very clearly motivated by the leader board, aspiring to overtake the person ahead of them on the leader board or not let the person immediately behind overtake them

> in the early days, there was a couple of other names which were sort of level-pegging with me all the time, and I was going up and down (Interviewee D).

Others experienced competition as more of a personal challenge, the number of transcriptions they could do daily, finishing a particular area;

> There is a competitive element to this; can I put more new entries or confirmations into the project than other people (Interviewee E).

However, towards the end of the project this competitive element also acted as a de-motivator, when it was no longer possible to change position on the leader board due to the high numbers involved;

> I'm 3000 ahead of the next guy behind me, and I've got no chance of catching the guy in front of me, so I don't look at it so much now (Interviewee F).

The revised Doubloons rankings served to help rekindle motivation for some interviewees, but one interviewee felt by the time he knew about it he was too far behind to realistically attempt to successfully compete and another felt the rewards

had no meaning for him. This tension between competition as both a motivator and de-motivator is an interesting one and not obvious in the survey responses where the vast majority of the 158 respondents said they were not motivated by the competitive element, with 62.6% not inspired by the leader board and 79.1% uninfluenced by the Doubloon rankings.

The main demotivating factor for interviewees were technical issues with the system. Speed of the interface was an issue at times and meant volunteers abandoned their session when it was not working properly. There was also frustration at the inability of the system to enable volunteers to flag or correct errors of any kind. Similarly as seen in Table 3, the most frustrating aspect for the survey respondents was the inability to edit pins when they were wrong, closely followed by not being able to identify errors. Interestingly the fact the place name searching was so limited was not regarded as a major problem, perhaps reflecting the likelihood of volunteers to go to places they know via the map, rather than by typing names into the search box.

[Table 3 near here]

All but one of the interviewees had been involved in one or more citizen-science projects previously, although many did not recognize the term "citizen science" itself, while just under half (49.4%) of the 158 respondents had not been involved in any citizen science projects before. The most frequently cited projects were

Wikipedia (28 respondents), *Geograph* (12) and *Cynefin* (10); *Cynefin* focuses on historical tithe maps and is another project of the National Library of Wales, so it is unsurprising we share contributors. The commonest write-in answers were genealogy sites: *Familysearch* had the most participants with seven, followed by *FreeBMD* and *FreeCen,* projects of Free UK Genealogy, which both had four. Overall 48 different citizen-science or crowd-sourcing projects were identified by volunteers, 17 people had participated in multiple projects with ten identifying as working on another project weekly, which included three on *Cynefin* and two on *Wikipedia*. Of the sixteen who worked daily on other projects, three worked on *Cynefin*, one each on *Geograph* and *Old Weather* and each of the others named different projects.

Two interviewees had participated extensively in *Geograph*, a crowd-sourcing project which aimed to assemble at least one photograph taken in each square kilometer of Great Britain, using the Ordnance Survey's National Grid framework; in other words another project exhaustively covering Britain's geography through crowd-sourcing. Twelve questionnaire respondents had also volunteered on *Geograph*, of whom eight identified it as their top project, but only two volunteers had contributed more than a few times: one daily Geograph contributor who did 111 total transcriptions in GB1900 and one weekly Geograph contributor who did 42,840 total transcriptions.

Conclusions

What follows obviously draws on our research into volunteer motivation but is presented as guidance to other similar historical projects.

Firstly, the tasks asked of volunteers cannot always be fun, but they should always be satisfying even in short initial sessions. GB1900 succeeded in large part because volunteers found both tasks enjoyable: creating initial transcriptions by covering the map with "pins", and adding confirmations to make pins change color. That second feature was added following problems with Cymru1900, so having even very limited software development resources at, effectively, a project mid-point was crucial, and something many projects lack because they rely on an external contractor whose role ends with the web site launch. However, it is hard to see how any software "tweaks" could make the transcription of large lumps of unstructured text fun, "addictive" or immediately satisfying, and this is reflected in the experience of the Transcribe Bentham project discussed earlier: crowd-sourcing is not a panacea.

Secondly, "gamification" works, but should be applied with care. With GB1900 this was complicated by the two tasks which needed to be not quite balanced: the Cymru1900 leader board ignored balance between new transcriptions and confirmations; the GB1900 board promoted perfect balance, and was predictably problematic in the final months; so we created an additional "finishers' league

table" based on a more complex formula, which we tried to sweeten with a silly name which only some volunteers liked. Once past their initial month, both rankings incentivized only our most committed volunteers, everyone else being left intimidatingly far behind. Additional regional rankings might have helped but would have complicated both software development and the user interface; and anyway, most of the work was done by those "most committed volunteers".

Thirdly, GB1900 differed from physical science projects, such as Galaxy Zoo, in that volunteers were not significantly motivated by a sense of "contributing to science", or to "human knowledge", and instead benefited more directly through an engagement with particular places which had meaning for them. One fundamental and probably unavoidable problem with our surveys is that they were carried out at the end of the project, so volunteers who joined in the initial months, worked on their local area but nowhere beyond, and so would have become inactive more than a year before the survey ran. However, most of our interviewees had started on areas which had particular meaning for them, then became "addicted".

This is partly simply that people like working with maps, but all the mapping within GB1900 was already freely accessible within the National Library of Scotland site: our interface was more interesting because it enabled volunteers to more actively engage with the maps. Further, this was not just about looking back at the past, as several interviewees spoke of developing a growing understanding of how

particular places have changed from then to now. In this sense the system was truly educational.

Fourthly, the majority of our most active contributors were retired, and made very substantial time commitments to the project: our top ten contributors were typically spending 15 to 20 hours per week on GB1900 over the first year. This made communication with them very important, and our initial channels were clearly inadequate. We progressively added the email newsletters, the support site and the progress mapping, but it would have been much better if these had been included in the main site.

Finally, we found it important to, not only develop additional communication channels, but also develop continuing relationships with our contributors. This requires real commitments on both sides. For example, had we simply changed the leader board formula, rather than creating the "finishers' league table" as a parallel ranking, we would have destroyed the large investment some contributors had made in their ranking. To some extent we had anticipated this, as we changed the GB1900 registration process to include assent to a "User Agreement", explicitly presented as a contract, which Cymru1900 had lacked. Even so, the depth of commitment showed by the volunteers we interviewed made us realize GB1900 had become as much their project as ours.

References

Aucott, P., P. McCann, H. Southall, and C. Fleet. 2016. GB1900.org Support Site.
Zoho Desk. 2016. https://support.gb1900.org/portal/kb

Causer, T., and V. Wallace. 2012. Building A Volunteer Community: Results and
Findings from Transcribe Bentham. *Digital Humanities Quarterly* 6 (2).
http://www.digitalhumanities.org/dhq/vol/6/2/000125/000125.html

Crowston, K., and I. Fagnot. 2008. The Motivational Arc of Massive Virtual
Collaboration. In *Proceedings of the IFIP WG 9.5 Working Conference on
Virtuality and Society: Massive Virtual Communities, Lüneberg, Germany.*

Dodge, M. 2017. Cartography I: Mapping Deeply, mapping the Past. In *Progress
in Human Geography,* 41 (1): 89–98. doi.org/10.1177/0309132516656431

Ell, P., L. Hughes, and H. Southall. 2016. Digitally Exposing the Place Names of
England and Wales. In *Placing Names: Enriching and Integrating
Gazetteers*, eds. M. L. Berman, R. Mostern, and H. R. Southall, 146–62.
Bloomington: Indiana University Press.

Eveleigh, A., C. Jennett, A. Blandford, A. Cox, and P. Brohan. 2014. Designing
for Dabblers and Deterring Drop-Outs in Citizen Science. In *Proceedings of
the 32nd Annual ACM Conference on Human Factors in Computing Systems*,
2985–94. New York: ACM.

Fleet, C. 2017. GB1900 Dot Distribution Maps Site. National Library of Scotland
Website. http://geo.nls.uk/maps/gb1900/

Great Britain Historical GIS Project. 2017. A Vision of Britain through Time
Website. Great Britain Historical GIS. University of Portsmouth.
http://www.visionofbritain.org.uk/

Haythornthwaite, C. 2009. Crowds and Communities: Light and Heavyweight
Models of Peer Production. In *Proceedings of the 42nd Hawaii International
Conference on System Sciences, HICSS,* 1–10.
doi.org/10.1109/HICSS.2009.137

Holley, R. 2009. *Many Hands Make Light Work : Public Collaborative OCR Text
Correction in Australian Historic Newspapers.* National Library of Australia.
https://www.nla.gov.au/content/many-hands-make-light-work-public-
collaborative-ocr-text-correction-in-australian-historic

Nov, O., O. Arazy, and D. Anderson. 2011a. Dusting for Science: Motivation and
Participation of Digital Citizen Science Volunteers. In *Proceedings of the
2011 iConference*, 68–74. New York: ACM.
doi.org/10.1145/1940761.1940771

———. 2011b. Technology-Mediated Citizen Science Participation: A
Motivational Model. In *Proceedings of the AAAI International Conference*

on Weblogs and Social Media (ICWSM 2011). Barcelona, Spain, 249–56.

Nov, O., O. Arazy, K. Lotts, and T. Naberhaus. 2013. Motivation-Targeted Personalized UI Design: A Novel Approach to Enhancing Citizen Science Participation. In *ECSCW 2013: Proceedings of the 13th European Conference on Computer Supported Cooperative Work, Paphos, Cyprus,* eds. O. Bertelsen, L. Ciolfi, M. Grasso, and G. Papadopoulos, 287–97. London, Springer. doi.org/10.1007/978-1-4471-5346-7_15

Raddick, M. J., G. Bracey, P. L. Gay, C. J. Lintott, C. Cardamone, P. Murray, K. Schawinski, A. S. Szalay, and J. Vandenberg. 2013. Galaxy Zoo: Motivations of Citizen Scientists. *Astronomy Education Review* 12 (1). doi.org/10.3847/AER2011021

Raddick, M. J., G. Bracey, P. L. Gay, C. J. Lintott, P. Murray, K. Schawinski, A. S. Szalay, and J. Vandenberg. 2010. Galaxy Zoo: Exploring the Motivations of Citizen Science Volunteers. *Astronomy Education Review* 9 (1). doi.org/10.3847/AER2009036

Reed, J., M. J. Raddick, A. Lardner, and K. Carney. 2013. An Exploratory Factor Analysis of Motivations for Participating in Zooniverse, a Collection of Virtual Citizen Science Projects. In *Proceedings of the 46th Hawaii International Conference on System Sciences*, 610–19. doi.org/10.1109/HICSS.2013.85

Romeo, F., and L. Blaser. 2011. Bringing Citizen Scientists and Historians Together. In *Museums and the Web 2011: Proceedings*, eds. J. Trant and D. Bearman. Toronto: Archives & Museum Informatics. https://www.museumsandtheweb.com/mw2011/papers/bringing_citizen_scientists_and_historians_tog

Southall, H., P. Aucott, C. Fleet, T. Pert, and M. Stoner. 2017. GB1900: Engaging the Public in Very Large Scale Gazetteer Construction from the Ordnance Survey 'County Series' 1:10,560 Mapping of Great Britain. *Journal of Map & Geography Libraries* 13 (1): 7–28. doi.org/10.1080/15420353.2017.1307305

Registration Period	Percentage of volunteers registered				
	0 sessions	1 session	2 sessions	3 sessions	4 or more sessions
pre 22nd Sept 2016	41.1	26.2	7.2	5.2	20.3
22 Sept - 31 Dec 2016	17.7	30.2	11.2	4.5	36.5
01 Jan - 31 Mar 2017	12.9	28.4	11.9	7.2	39.7
01 Apr - 30 Jun 2017	19.2	39.7	11.0	6.8	23.3
01 Jul - 30 Sept 2017	21.8	47.3	5.5	1.8	23.6
01 Oct 2017 - 31 Jan 2018	53.8	34.6	3.8	0	7.7

Table 1: Percentage of volunteer sessions

Response to GB1900	n = 158				
	Strongly agree (%)	Agree (%)	No opinion (%)	Disagree (%)	Strongly disagree (%)
I found it enjoyable	46.8	43.7	4.4	4.4	0.6
I found it addictive	31.0	36.7	15.2	14.6	2.5
It gave me a sense of satisfaction	41.8	46.2	7.6	3.2	1.3
I felt I was creating something I might use in future	46.2	36.7	8.9	7.6	0.6
I liked looking at the old maps	70.9	25.3	3.8	0.0	0.0

Table 2: Positive motivations identified by questionnaire respondents

Short-comings of GB1900	n = 158				
	Very frustrating (%)	Frustrating (%)	Mildly frustrating (%)	Made no difference to me (%)	Prefer not to say (%)
Cannot see names once entered	7.6	36.7	32.3	21.5	1.9
Cannot flag errors	22.2	42.4	23.4	10.1	1.9
Cannot edit existing pins	24.1	34.2	26.0	13.9	1.9
Cannot tell when an area is complete	19.0	31.0	26.6	21.5	1.9
Place name searching limited	5.7	20.9	27.2	44.3	1.9

Table 3: GB1900 limitations identified by questionnaire respondents

List of Figures

Figure 1: Transcriptions per week for GB1900 and Cymru1900
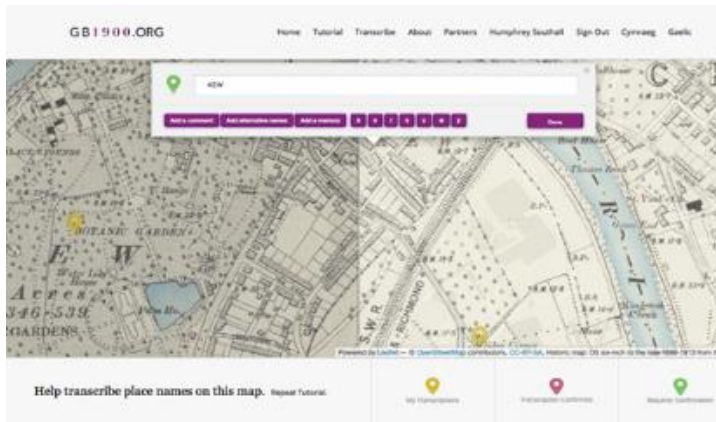
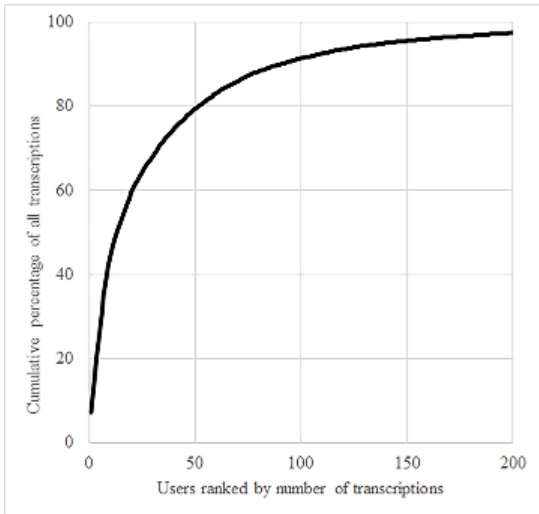Figure 2: Transcription screen for GB1900
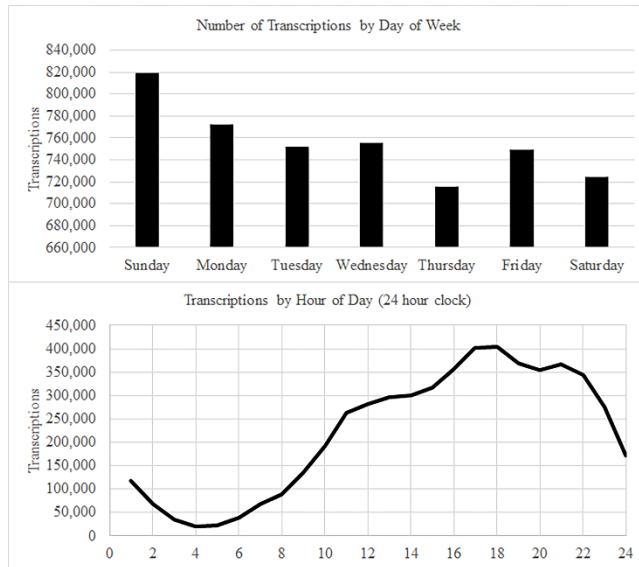
Figure 3: Cumulative transcriptions of top 200

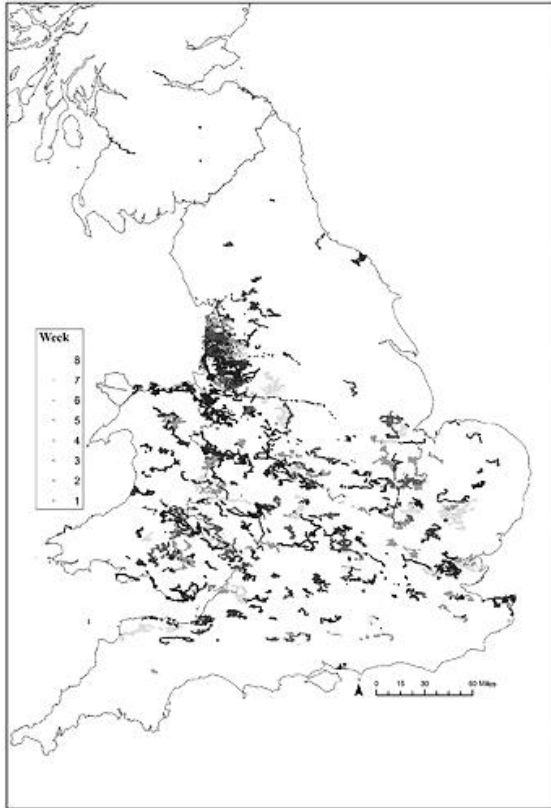Figure 4: Transcription Patterns

Figure 5: Geographical locations of transcriptions during first two months of GB1900 registration by Interviewee A
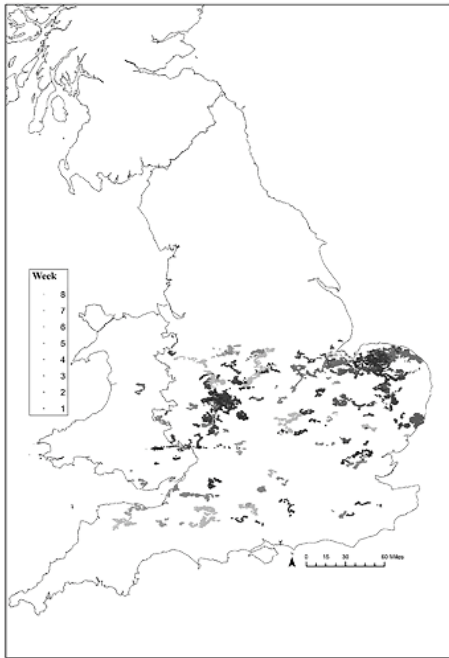
Figure 6: Geographical locations of transcriptions during first two months of GB1900 registration by Interviewee E