



FOUNDATION PROJECTS

SUPPORTING DOCUMENTATION

User engagement analysis report

Engaging Crowds: citizen research and
cultural heritage data at scale

Bernard Odgen
The National Archives

Contents

Introduction	1
Overview of the data	1
Scope of the analysis	5
Summary of key findings	6
Related work	7
Patterns of contribution	7
Task design	7
Volunteer motivation	8
Summary of related work	8
Research questions	9
When do volunteers contribute?	9
What are the volunteering patterns across the lifetime of a project?	12
Discussion	16
Key findings	16
Patterns of contribution	16
Task design	16
Volunteer motivation	17
Implications	17
Further work	17
Conclusion	19
References	20
Appendix	21

Introduction

This study looks at classification (transcribed) data from the three citizen research projects in *Engaging Crowds*, which were set up on the [Zooniverse](#) platform: [HMS NHS: The Nautical Health Service](#), [Scarlets and Blues](#) and [The RBGE Herbarium: Exploring Gesneriaceae, the African violet family](#).

Each project presents volunteers with a series of tasks (known as ‘workflows’ in Zooniverse terminology). The projects and their workflows are summarised in Tables 1, 2 and 3 below. From these we can see that *HMS NHS* is the largest and longest-running of the three projects and so gives us our best sample of volunteer behaviour. We can also see that the workflows in *The RBGE Herbarium* were completed most quickly by volunteers, possibly aided by the project’s familiarity: it is perhaps most similar to the majority of Zooniverse projects, which deal with observational scientific data in simple workflows with short tasks. Volunteers tended to contribute to just one of the *Engaging Crowds* projects. Within the data analysed, only 13 volunteers contributed to all three projects.¹

Overview of the data

Here we explore classification data from these projects to see what it suggests about patterns of volunteer engagement over time. A classification is a single completed workflow by a single volunteer on a single record. The data consists of 190,526 classifications made from the date and time of the Zooniverse launch announcement for each project. Classifications during beta testing and soft launch periods are thus excluded. A further 8,925 classifications from volunteers who were not logged in were also discarded as it is harder to link classifications to individuals for these volunteers. Ten classifications with a negative duration were discarded as uninterpretable: these classifications are recorded with a start time later than their end time and so it is hard to understand exactly when they took place. Five users were members of the project team. Project team members made a total of 616 classifications, 583 of which were in the ‘How Disposed Of’ workflow of *HMS NHS*. Any effect upon data relating to other workflows should be negligible. The remaining data covers the first phase of *HMS NHS*, the full run of *Scarlets and Blues* and the first few weeks of *The RBGE Herbarium*.

The data was analysed with Python scripts, using [Pandas](#) for data processing and [Plotly](#) to generate graphs. These scripts are available on [GitHub](#).

Following feedback from accessibility experts, the charts printed in the report have been modified from those originally used to perform this analysis in order to improve their accessibility. They represent the same data, but changing the form of the representation could possibly affect interpretation of the data. This is a risk particularly when interpreting shifts in colour in heat maps, which are not as clearly objective a measure as a data point on a line graph. However, our analysis focused on clear effects less likely to be illusions of representation and we have high confidence that it remains valid and useful.²

¹ None of these 13 volunteers were members of the project team. No member of the project team contributed to more than one project within the analysed data set.

² Other than a few ad hoc calculations, data and charts used in the analysis were generated from tag report_original of https://github.com/nationalarchives/engaging_crowds_user_analysis. The charts printed in this report were generated from tag report_final.

	HMS NHS	Scarlets and Blues	The RBGE Herbarium
Launch	29 June	16 November	11 January
Last Classification	22 January ³	19 January	31 January ⁴
Duration	207 days	64 days	20 days
Workflows	13	2	2
Discarded classifications ⁵	6,916 (4.3%)	849 (8.7%)	1170 (6.0%)
Analysed classifications	154,178	8,935	18,478
Total classifications ⁶	161,094	9,784	19,648
Volunteers	1,505	394	357
Repeat volunteers ⁷	1,207 (80.2%)	271 (68.8%)	276 (77.3%)

Table 1: Headline information on the projects. Volunteer counts are based only on the analysed data.

Project and Record Type	Workflow name	Type of transcription	Median Time to complete one classification (in minutes) ⁸
HMS NHS	Admission Number	Columns of numbers (number-type)	14
	Age		23

³ Completion of part one. *HMS NHS* launched a second phase on 1 February 2022. Only classifications from part one were included in this analysis.

⁴ *The RBGE Herbarium* completed on 8 February 2022. Classifications after 31 January 2022 were excluded from analysis due to project schedules.

⁵ Anonymous classifications and the ten classifications with negative duration. Number in brackets is the percentage of *total* classifications on the project.

⁶ Total classifications from the in-scope duration of each project.

⁷ Volunteers who submitted more than one classification. Number in brackets is the percentage of total volunteers on the project.

⁸ Rounded to the nearest minute. Some classifications take much longer: this is likely to indicate that a participant paused their work and returned to the browser window after a number of hours or days. We have used the median to prevent these high classification times from skewing the average.

Hospital Admission Register entries in a standard tabular format, handwritten.	Years at Sea ⁹		16
	Days Victualled ¹⁰		17
	Date of Entry	Columns of dates (date-type).	18
	Date of Discharge		22
	Quality ¹¹	Columns of text selected via dropdown (dropdown-type).	20
	How Disposed Of ¹²		14
	Name	Columns of text entered as free text (noun-type).	27
	Place of Birth		20
	Port Sailed Out Of		20
	Last Services ¹³		27
	Nature of Complaint		21
Scarlets and Blues Pages from minute books, handwritten and mostly unstructured.	People	Rows containing names and titles, free text and page numbers.	35
	Attendance ¹⁴	Attendance lists and 'standard' minute entries, entered via a list of options or by free text.	28
	Minutes ¹⁵	Other handwritten minute entries, entered as free text.	34
The RBGE Herbarium Handwritten or typed specimen labels	Latitude/Longitude	Latitude and longitude. Free entry of numbers, with a dropbox for direction.	1
	Geography	Country/locality/altitude information. Free entry of text and numbers, and dropboxes for altitude units and qualifiers.	0.4

Table 2: Headline information on the workflows.

⁹ 'Years at Sea' was more complex than other number-type workflows, requiring that two numbers be presented in a particular format and that months be translated into fractions of a year.

¹⁰ How long the patient was in the hospital.

¹¹ The patient's rank.

¹² The circumstances of departure from the hospital.

¹³ The last ship that the patient served on.

¹⁴ 'Attendance' and 'Minutes' are two branches of the 'Meetings' workflow as presented to users. They involve different tasks and so are analysed as separate workflows.

¹⁵ See footnote 14 above.

Total Unique Volunteers on...	Volunteers	Percentage of total volunteers
All three projects	13	0.6%
Exactly two projects	128	6.1%
Exactly one project	1,961	93.3%
HMS NHS only	1,374	65.4%
Scarlets and Blues only	297	14.1%
The RBGE Herbarium only	290	13.8%
HMS NHS and Scarlets and Blues only	74	3.5%
HMS NHS and The RBGE Herbarium only	44	2.1%
Scarlets and Blues and The RBGE Herbarium only	10	0.5%
(Total unique volunteers)	2,102	100.0%

Table 3: Volunteers contributing to each project.

Scope of the analysis

This is a limited study to identify areas that might be fruitful for further investigation. Although the classification data can be interrogated in several different ways, we focused on using it to examine two research questions that gave us an insight into user engagement.

When do volunteers contribute?

What are the volunteering patterns across the lifetime of a project?

The data has not been subjected to statistical analysis and so highlighted differences may not be statistically significant. Charts are generated independently from one another and so do not share common scales. Thus we compare the patterns not the absolute values. Throughout, we mainly highlight large differences as these are more likely to be meaningful. That is, they are somewhat less likely to be misreadings due to being led astray by a visualisation, and somewhat more likely to be statistically significant (i.e. to be due to some real effect rather than to expected random variation). Larger differences are also less likely to be over-interpretations of the appearance of a particular chart or visualisation.

This analysis does not attempt to understand the impact of the newly developed indexing tool on the Zooniverse platform, which allows volunteers to choose their own pathway through a project. This was evaluated through the volunteer survey and the project workshop 'After the Crowds Disperse'. We did not design the projects as A/B tests (in phases with and without the indexing tool) on the impact of the tool visible through classification data alone. While there are quantitative ways to investigate the tool's impact, we chose to rely on self-reporting data through the volunteer survey. The analysis of the survey results can be found in the [Annex 7](#) of the final report.

Summary of key findings

Our key findings are summarised below and explored in more detail in the main body of the report.

1. When do volunteers contribute?
 - Volunteers tend to contribute during the daytime (8 a.m. - 8 p.m.).
 - Volunteers tend not to contribute early in the morning (earlier than 8 a.m.).
 - An outlier to this pattern are tasks involving direct transcription of numbers only.
2. What are the volunteering patterns across the lifetime of a project?
 - Some workflows are more popular than others.
 - There may be a tendency for volunteers to run through some workflows in the order they appear on the platform rather than consciously choosing where to focus their efforts.

Related work

Several studies have investigated patterns of contribution and questions around volunteer motivation for participants in Zooniverse projects. This overview of that work gives some context to this study of user engagement in the *Engaging Crowds* projects.

Patterns of contribution

Many studies, both of Zooniverse and other crowdsourcing projects, have found that often a large proportion of work is undertaken by a relatively small number of active volunteers. There also tends to be a very large number of volunteers who make small contributions to projects.

Ponciano and Brasileiro (2014a) have explored this distinction further. They have considered in depth the same kind of classification data as used here, and developed quantitative metrics to identify five “engagement profiles” relating to volunteers on *Galaxy Zoo* and *The Milky Way Project*.

These profiles were:

- Hard-working: “Volunteers who exhibit this type of engagement profile typically work hard and regularly when arriving at the project, but may leave the project quickly”.
- Spasmodic: “Volunteers who exhibit this engagement profile provide an intense contribution, at a short period of time and with irregular periodicity within this period”.
- Persistent: “Volunteers who exhibit a persistent engagement profile remain in the project for a long period of time but contribute only a few days within this time period”.
- Lasting: “Characterised by an engagement pattern similar to persistent engagement, with the difference that volunteers exhibit here a much shorter activity duration”.
- Moderate: “moderate volunteers have intermediate scores in all categories of engagement metric”.

Task design

Other studies have considered how task design affects user engagement in Zooniverse projects. Spiers *et al* (2019) looked at classification data from 63 Zooniverse projects and found that small tweaks to project design could have an effect on who contributes and how often they do so. Sprinks *et al* (2019) focused on the *Planet Four: Craters* project and found that volunteers preferred greater variety and autonomy in their tasks. Simpler tasks were completed more quickly by volunteers but not necessarily to the highest standards of accuracy. Eveleigh *et al* (2013) explored elements of gamification in the *Old Weather* project and found that competitive features appealed to some volunteers but discouraged others. These insights into user behaviour led these scholars to suggest ways that tasks on citizen research projects could be explicitly designed to meet the preferences of different sorts of users.

Volunteer motivation

As Ponciano and Brasileiro have pointed out, understanding the motivations of different volunteers can further our understanding of their responses to citizen research projects. Raddick *et al* (2013) found through a volunteer survey that users of *Galaxy Zoo* had multiple motivations which shifted over time. The team also found that most volunteers identified with the project goals and got intrinsic enjoyment from the tasks they completed while fewer volunteers were interested in building their skills and participating in an online community. Cox *et al* (2018) echoed these findings. They looked across five different Zooniverse projects and found that the most active volunteers were motivated by the topic and values of a project more than social interaction or career development.

Summary of related work

Taken in the round, these studies support the idea of a range of different constituencies of volunteers on projects, each with their own motivations and working patterns.

The frequency, intensity and nature of contribution can differ between volunteers but it is possible to discern some general patterns across a spectrum of occasional to highly active contributors. Volunteers seem to prefer simple tasks and appreciate options for independence and variety in their work. Connection with the project goals appears to be a prerequisite for volunteer contributions but there are multiple shifting factors motivating volunteers beyond this.

We will refer back to these findings as part of our analysis of patterns of volunteer contribution across the *Engaging Crowds* projects. It must be noted that many of the studies mentioned here refer to projects based on data from the natural or physical sciences, where findings could differ from the cultural heritage sphere.

Research questions

We are looking at classification data to consider the following two questions:

When do volunteers contribute?

What are the volunteering patterns across the lifetime of a project?

Figures 1 to 20 in the [Appendix](#) of this report show the different representations of the data which we explore below.

When do volunteers contribute?

Understanding when people classify can help us learn more about the preferences of volunteers working on the *Engaging Crowds* projects.

- Figures 1 to 13 are heat maps of classification start times for volunteers working on each of the three projects, within given periods of the day.
- Figures 1 to 3 show volunteer activity for each individual project.
- Figures 4 to 8 show volunteer activity for selected workflows and workflow types in *HMS NHS*.
- Figures 9 to 11 show volunteer activity on the ‘People’, ‘Attendance’ and ‘Minutes’ workflows for *Scarlets and Blues*.
- Figures 12 and 13 show volunteer activity on the different workflows for *The RBGE Herbarium*.

The heat maps are divided into six time periods. Each time period covers four hours, starting from midnight at the bottom of the chart. Time here is the volunteer’s local time, rather than UTC (i.e. international time differences have been accounted for).¹⁶ The periods are labelled as described in Table 4.

Period label	Duration
Small hours	12 a.m. – 4 a.m.
Early morning	4 a.m. – 8 a.m.
Morning	8 a.m. – 12 p.m.
Afternoon	12 p.m. – 4 p.m.
Evening	4 p.m. – 8 p.m.
Night	8 p.m. – 12 a.m.

Table 4: Periods of the day shown in the heat maps (Figures 1 to 13).

¹⁶ The start time is the started_at field in the classification metadata, corrected to local time by subtracting the utc_offset field.

The bars at the top of each figure show total classifications on each day of the week; the bars to the right show total classifications within each time period. The shades in the heat map show the percentage of total classifications in a given period on a given day.

Darker shades indicate greater activity, as shown in the scale on the right of each chart: black is most active (hottest), white is least active (coldest). The charts do not have a common scale so, for example, a short bar on *HMS NHS* will likely be a larger number of classifications than a long bar on *Scarlets and Blues* due to the relative sizes of the projects.

Findings

Activity patterns: projects

The heat map for *HMS NHS* (Figure 1) shows most activity in the mornings and afternoons from Monday to Friday. Activity is also quite high at weekends in the same time periods, and in the evening through the full week. The other two projects (Figures 2 and 3) seem roughly consistent with *HMS NHS* in terms of active periods during the day, with most activity occurring from morning to evening. The pattern of days is less consistent. The strong Tuesday to Thursday activity in *The RBGE Herbarium* may be an artefact of the project having launched on a Tuesday and if this is the case, the effect would be diluted over a longer period of time.

The 'small hours' period (midnight to 4 a.m.) seems particularly inactive through all three projects, and the 'early morning' period (4 a.m. to 8 a.m.) is similarly inactive in *Scarlets and Blues* and *The RBGE Herbarium*. The relatively high 'early morning' activity in *HMS NHS* may be thanks to its number-type workflows, as these show higher early morning activity than other workflow types (see below).

Volunteers tend to classify during the daytime and early evening (8 a.m. to 8 p.m.), with activity being particularly low before 8 a.m.

Activity Patterns: workflows

In Figures 4 to 13 we break down activity in the three projects according to their different individual workflows and workflow types (see Table 2 above). This may allow us to detect activity specific to different types of task. We concentrate mainly upon periods of the day rather than days of the week, as changes are clearer between time periods than over days.

HMS NHS

The multiple workflows in *HMS NHS* give us the most opportunities to observe differences in volunteer activity across workflows. Figures 4 to 8 show volunteer activity for some of the different workflows and workflow types in the project.

Most workflow types follow the pattern for the overall project, with volunteers most active from morning to evening. Number-type workflows (Figure 4) differ the most from the overall project pattern: while there is higher activity on these workflows from morning to evening, activity is more evenly spread through the day. There is greater activity in the 'early morning' period than for the other workflow types. We speculate that this is because numbers are quicker and easier to transcribe and so these tasks perhaps both fit into smaller periods of time and require less effort.

Figure 5 focuses on 'Years at Sea', one of these number-type workflows. The chart shows that volunteer activity on this workflow is more concentrated in the morning to evening periods than the number-type

workflows as a whole (Figure 4).¹⁷ The pattern of activity on 'Years at Sea' is thus more similar to the overall *HMS NHS* project pattern than it is to the number-type workflows.

We have already suggested that the shorter and simpler nature of number-type workflows may influence their different activity patterns. If this is the case then it seems relevant that 'Years at Sea' is the most complex of the number-type workflows. Unlike the others, it requires reading two numbers and converting months into decimal fractions of a year. These data suggest therefore that volunteers may be less inclined to transcribe more complex workflows in the early morning.

Figure 6 shows volunteer activity for all the date-type workflows. This seems to resemble the overall project tendency for activity from morning to evening.

Figures 7 and 8 focus respectively on the 'Date of Discharge' and 'Date of Entry' workflows, which are both date-type workflows. Periods of volunteer activity on the 'Date of Discharge' workflow (Figure 7) are similar to the aggregated activity for both date-type workflows (Figure 6), with activity mainly from morning to evening. Volunteer activity in the 'Date of Entry' workflow (Figure 8), however, is more spread out across different times of day.

Although the 'Date of Discharge' and 'Date of Entry' workflows are similar, the content of the latter contains less variation because the column proceeds chronologically. We might therefore suspect that 'Date of Entry' has something in common with the number-type workflows, such as being easy to fit into different times of the day. It seems reasonable that dates in chronological order would be relatively easy to transcribe.

There are no variations strong enough for comment among the noun- or dropdown-type workflows.

Scarlets and Blues

The 'Attendance' (Figure 10) and 'Minutes' (Figure 11) branches of the 'Meetings' workflow show similar activity, both to each other and to the overall project (Figure 2). Volunteers tend to work from morning to evening.

The chart for the 'People' workflow (Figure 9) is hard to compare as it is scaled quite differently, but seems to show a greater concentration of activity in the afternoon. Perhaps this chart is more reflective of a few volunteers' working patterns than it is of the nature of the workflow, given that this workflow was much less popular than the others (75 total volunteers with a median of one classification each).

The RBGE Herbarium

Figures 12 and 13 show volunteer activity on the different workflows for *The RBGE Herbarium*. Scaling differences make the two charts difficult to compare, both to one another and to the overall project chart (Figure 3). The bar charts seem to indicate that there was a fairly consistent activity pattern across the project as a whole and the two individual workflows, with volunteers working mostly from morning to evening. The 'Geography' workflow (Figure 12) shows volunteer activity to be a little more concentrated in the morning through evening periods, but perhaps not meaningfully so.

The 'Latitude/Longitude' workflow (Figure 13) has relatively few volunteers but more total classifications, with a median of 21.5 classifications per volunteer, perhaps because it was quicker and more straightforward than the 'Geography' workflow. Its classifications have a much higher median than for any other workflow across all of the projects: the next highest is six. This suggests a small group of highly active volunteers. In some respects this resembles the 'People' workflow in *Scarlets and Blues*, but the very low

¹⁷ Note that Figure 4 includes activity on 'Years at Sea', as well as the other workflows. If 'Years at Sea' were excluded from this figure then the difference would be still more noticeable.

volunteer count and median for 'People' suggests a very small group doing nearly all of the work. We note that even a single classification in 'People' can be a significant investment of time given the difficulty of the handwriting and lack of context.

Limitations

These charts give an indication of volunteering patterns in terms of the local time when volunteers are working, based on the time zone recorded in their devices. Time of year and latitude may also be relevant, as volunteering patterns might shift seasonally. The significance of days and times of day could differ according to local culture. It may also be that different divisions of the day would yield different information: for example, we might (or might not) find that most of the 'night' classifications happen before 10 p.m. Time allowed by the project, including its overall duration, means we have not attempted to consider these additional factors.

The lack of a common scale makes comparisons between charts difficult. This can be especially tricky when shaded temperatures are scaled differently. An apparently small change in shade on one chart may be the same size as an apparently large change in shade on another chart with a scale covering half of the range.

We should also be aware that different charts aggregate different numbers of underlying workflows. For example, the date-type chart (Figure 6) aggregates only two workflows and so is effectively the average of the 'Date of Discharge' (Figure 7) and the 'Date of Entry' (Figure 8) workflows.

What are the volunteering patterns across the lifetime of a project?

We sought to monitor the activity of volunteers across the lifetime of each of these three projects. We looked at the number of active volunteers. On any given date, this is the number of volunteers who have submitted their *first* classification on or before that date minus the number of volunteers who have submitted their *final* classification on or before that date.¹⁸

All dates in this section are in UTC: we are not considering the time of day that a classification was made and so local time is less relevant. We have focused on the data for individual workflows (see Table 2), which allows us a sense of which activity types are most popular over the projects' lifetimes:

- Figures 14 to 16 show the number of active volunteers on individual workflows for each of the three projects.
- Figures 17 to 20 show the number of active volunteers on individual workflows for each of the workflow types in *HMS NHS*.

As the line climbs higher on these charts, a greater number of volunteers are making their first classification than are making their final classification (roughly, more people are joining than leaving). As the line falls lower, more volunteers are making their final classification than are making their first classification (roughly, more people are leaving than joining).

In *HMS NHS* charts (Figures 14 and 17 to 20), the number prior to the workflow name in the chart's key indicates where it appeared in the list of workflows in the project.

¹⁸ Using the `started_at` and `finished_at` fields in the classification metadata.

Findings

In the next sections, we use Figures 14 to 20 to take a closer look at workflow activity for each of the projects.

HMS NHS

Figures 14a and 14b show the number of active volunteers on all of the *HMS NHS* workflows. These two figures show the same data, once as a line chart and once as a scatter plot. Two workflows stand out as attracting the most active volunteers: 'Name' is very active through the first couple of months, with 'Years at Sea' taking over the top spot later in the project. 'Nature of Complaint' and 'Date of Entry' also quickly climb to over 40 active volunteers, though 'Date of Entry' soon falls off. The only other workflow to exceed 40 active volunteers is 'Date of Discharge', which reaches this level in November 2021.

It may be that these workflows were particularly appealing for volunteers. In the 'Nature of Complaint' workflow volunteers transcribe medical ailments. 'Date of Entry' and 'Years at Sea' are both relatively quick and simple. However, there are other simple workflows in *HMS NHS* that are not as popular in the data.

Figure 17 shows the number of volunteers on the number-type workflows in *HMS NHS*. We can see that three of these workflows are completed early in the project, and in order of their appearance in the list of workflows on the platform. We can tentatively hypothesise that this is due to volunteers completing one number-type workflow and moving on to the next one in the list. 'Years at Sea', the third of the number-type workflows, rises to high numbers of volunteers in the latter part of the project. The slow start might be because 'Years at Sea' is more difficult than the other number-type workflows.¹⁹

Figure 18 shows the number of volunteers on the two date-type workflows. We see that 'Date of Entry' is popular early on, while 'Date of Discharge' seems to accumulate volunteers fairly steadily. As we have seen, both are among the few workflows to exceed 40 active volunteers and again this might be due to dates being relatively simple to transcribe.

Figure 19 shows the number of volunteers on the two dropdown-type workflows to be at a similar level until mid-September 2021, when the 'Quality' workflow has a surge in volunteers. With the noun-type workflows shown in Figure 20, we see again the early popularity of the 'Name' and 'Nature of Complaint' workflows. The 'Last Services' workflow seems less popular than the others. There is no clear explanation for this, but perhaps transcription of ship names requires (or is perceived to require) a more specialised kind of knowledge than transcription of person, place and complaint names.

Generally, from Figures 17 to 20, it seems that workflows appearing earlier in the interface are often completed earlier. We can see this in Table 5, which shows the dates of 'thank you' Talk posts made as each workflow is completed as a proxy for workflow completion. It may be that there is a tendency to work on a particular type of task, from the top down: more work would be needed to confirm that this was an actual volunteer working pattern.

¹⁹ It required volunteers to deal with pairs of numbers and to write months as fractions of a year.

Date of 'Thank you' post	Workflow name	Position	Workflow type
13 August 2021	Admission Number	1	Number
6 September 2021	Date of Entry	2	Date
13 September 2021	Age	5	Number
18 September 2021	Name	3	Text
8 October 2021	Place of Birth	6	Text
14 October 2021	Days Victualled	13	Number
27 October 2021	Quality	4	Dropdown
29 October 2021	Port Sailed Out Of	7	Text
7 November 2021	Nature of Complaint	10	Text
28 December 2021	Years at Sea	8	Number
4 January 2022	Last Services	9	Text
20 January 2022	Date of Discharge	11	Date
22 January 2022	How Disposed Of	12	Dropdown

Table 5: HMS NHS workflows in order of completion, showing position in the user interface and the type of the workflow. The three "simple" number workflows complete early on. Workflows of the same type usually complete in order of their position in the interface.

Scarlets and Blues

Figure 15 shows the number of volunteers on three workflows in *Scarlets and Blues*. We see that the 'Attendance' and 'Minutes' branches of the 'Meetings' workflow have similar numbers of active volunteers throughout the lifetime of the project.

Both of the 'Meetings' branches have many more volunteers than the 'People' workflow, which peaks at 12 active volunteers.²⁰ Volunteers demonstrated a clear preference for the 'Meetings' workflow. The launch email focused more on the kinds of content covered by the 'Meetings' workflow, which may have encouraged volunteers to favour it. It might have been that the more descriptive form of the meeting minutes was more interesting than lists of names, or that the challenging handwriting and relative lack of

²⁰ That is, 12 active volunteers at any one time. There were a total of 75 volunteers on this workflow over the lifetime of the project.

context made transcribing names less enjoyable. The most important factor may simply have been that the 'People' workflow was long by Zooniverse standards. Some feedback on Talk (Zooniverse's forum) indicated that this was a problem for volunteers.

The RBGE Herbarium

Figure 16 shows the number of active volunteers on the two workflows in *The RBGE Herbarium*. This shows the initial spike in volunteers followed by a decline as the workflows approach completion. The 'Geography' workflow seems to attract more volunteers but 'Latitude/Longitude' completes quickly. This may be because the 'Latitude/Longitude' workflow is more straightforward, while the 'Geography' workflow requires extraction of more data and often some online research to verify locations: thus it is more difficult but also perhaps more interesting. This also might suggest that 'Geography' volunteers are more motivated by their interest in the project than by the simplicity of the task.

Limitations

This is just one of the possible measures of the number of volunteers taking part. Due to time considerations, we were not able to attempt to track how active these volunteers were: this measure makes no distinction between frequent and occasional volunteers.²¹ To give the most extreme example, a volunteer who classifies once at the beginning and once at the end of the project will be considered 'active' throughout the project. The number of classifications or amount of time spent by each volunteer is not shown in these charts and would be a useful avenue of further study.²²

Volunteers who classify once only are not represented in these charts. When volunteers classify just once, the start and end date of their classification period is the same and so, in line with our definition of 'active', the volunteer is both added to and subtracted from the total volunteers for the same day. The exception would be if their single classification spanned two days (i.e. starting before midnight and finishing after midnight UTC).

²¹ Ponciano and Brasileiro (2014a) have a similar and better developed measure.

²² Ponciano and Brasileiro (2014a), p. 253, note that duration of time is better than classifications as a measure of engagement.

Discussion

This is a brief exploration of a limited view of the data and there is a risk of over-interpreting the appearance of a particular chart or visualisation. For example, a large shade-shift in a heat map or a large step in a graph might only represent a slight change at certain scales. Deeper analysis of the data, and particularly statistical analysis of the significance of effects and selection of metrics, would help to protect against this. However, our work allows us to draw tentative findings that could inform future work.

Key findings

1. When do volunteers contribute?

Volunteers tend to contribute during the daytime – that is, during our morning to evening periods, 8 a.m. to 8 p.m. Volunteers tend not to contribute very early in the morning (earlier than 8 a.m. and especially not earlier than 4 a.m.).

Some *HMS NHS* workflows, which are arguably the easier ones, had more spread out activity.

2. What are the volunteering patterns across the lifetime of a project?

Some workflows appear to be more popular than others. Some *HMS NHS* workflows attracted more volunteers than others, and the *Scarlets and Blues* volunteers had a strong preference for the ‘Meetings’ workflow. Popularity may manifest in different ways for *The RBGE Herbarium*, with ‘Geography’ attracting a larger number of volunteers but ‘Latitude/Longitude’ showing high median classifications.

The data also suggests that there might be some tendency to run through workflows of a particular type in the order that they are listed.

When putting these findings in the context of related studies, we should keep in mind that much of this literature is based upon natural and physical science projects rather than cultural heritage projects.

Patterns of contribution

As posited by Ponciano and Brasileiro (2014), patterns of user behaviour can be discerned in Zooniverse projects (such as intense activity, long-lasting engagement or *ad hoc* contributions). Our analysis sheds light on possible patterns relating to our projects, such as the tendency to avoid working in the ‘small hours’ or to complete workflows of a given type in the order they appear in the interface.

Task design

Research around task design can help us think further about why more volunteers worked on particular workflows. Sprinks *et al* (2017, 2019) found that simple tasks could be completed quickly and easily, which may fit the preferences of some volunteers. This work perhaps provides some indication as to why volunteers appeared to prefer particular workflows in *HMS NHS* and *Scarlets and Blues*.

Volunteer motivation

The literature indicates that individual volunteers can have multiple motivations of varying importance which contribute to different behaviours. Researchers looking at other Zooniverse projects such as Cox *et al* (2018), Ponciano and Brasileiro (2014a) and Raddick *et al* (2013) have found that volunteer interest in the subject matter, identification with the project goals and enjoyment of the task were key. This could help us to understand why volunteers appear to have been drawn to particular workflows.

Implications

By looking at volunteer behaviour, we can start to appreciate the best ways to engage our audiences in citizen research. The data on times of classification suggests that simple workflows can be fitted into more parts of the day. Simple workflows may therefore draw in a wider range of contributors and maintain a higher rate of classification by being easier to fit around life commitments. Conversely, projects with more complex workflows may be limiting their audiences to those who can commit to regular classification sessions, and whose engagement may arguably be deeper or more focused.

As Spiers *et al* (2019) have suggested, we could think about how project design could affect volunteer participation. Different workflows within individual projects could be designed to meet the needs of particular constituencies, be they highly active contributors who have a deep connection with the project or people who dip in and want to complete their work quickly. The design of citizen research projects has implications for both volunteer well-being and data accuracy. If we assume that classifications improve with practice, more-active volunteers are more likely to produce accurate data. At the same time, we need to balance this with respect for the time of volunteers, aiming to prevent them becoming overworked.²³

Further work

The research questions considered here could be explored further, with more in-depth analysis of each project and further consideration of findings from other citizen research projects. Comparison across these projects would permit insights into variations in engagement across different cultural heritage projects.

Future work could also explore this dataset in new ways to generate further insights around user engagement in citizen research. The question of the impact of the new Zooniverse indexing tool could be further explored by investigating the extent to which volunteers move through the records in order. If we find that volunteers are not simply progressing page by page then it may be that they are using the index to focus on material of interest. Looking at the paths taken through the records by individual volunteers could identify particular working patterns enabled by the indexing tool.

Looking at the total number of classifications completed, or the number of classifications completed in a single sitting by a volunteer, could give a different view of engagement with these projects.²⁴ Combining this

²³ Ponciano and Brasileiro note that, by monitoring volunteer activity, we may be able to intervene to stop volunteers from working too hard.

²⁴ Mao, Kamar, and Horvitz present a means of grouping classifications into sessions.

with information about what material is being worked on may help us to see whether particular tasks are associated with particular patterns of contribution.

We could also study the Talk pages to understand how volunteers participate in online discussion and how far this relates to the frequency or quality of their classification activity. For example, BrodeFrank *et al* followed Talk activity to assess engagement in their mission to incorporate citizen research activity into a museum exhibition.

Conclusion

This preliminary analysis of volunteer activity across the three citizen research projects was focused on areas of particular interest to project partners in the *Engaging Crowds* project. Our tentative conclusions set out here are consistent with other work on engagement with citizen research projects.

Our findings indicate that volunteers on these projects tended to contribute between 8 a.m. and 8 p.m. and to avoid working during the 'small hours', between midnight and 4 a.m. They suggest that certain workflows were more attractive to volunteers than others. It may be that there is a tendency for volunteers to run through workflows of a particular type in the order that they are listed. Lastly, it may be that some workflows are associated with particular activity patterns. There is the potential for further work to give more insight into volunteer behaviour.

References

- BrodeFrank, Jessica, and others. 2020. 'Crowdsourcing Knowledge for Representation: Interactive Learning and Engagement with Collections Using Zooniverse's Mapping Historic Skies', *Theory and Practice*, 3 <https://articles.themuseumscholar.org/2020/04/24/tp_vol3_brodefrank/>.
- Cox, Joe, and others. 2018 'Doing Good Online: The Changing Relationships between Motivations, Activity and Retention among Online Volunteers', *Nonprofit and Voluntary Sector Quarterly* (Vol. 47, Issue 5) <<https://doi.org/10.1177/0899764018783066>>.
- Eveleigh, Alexandra, Jennett, Charlene, Lynn, Stuart, and others. 2013. "'I Want to Be a Captain! I Want to Be a Captain!": Gamification in the Old Weather Citizen Science Project', in *Proceedings of the First International Conference on Gameful Design, Research, and Applications* (presented at the Gamification '13, New York: ACM), pp. 79–82 <[dx.doi.org/10.1145/2583008.2583019](https://doi.org/10.1145/2583008.2583019)>.
- Ponciano, Lesandro and Brasileiro, Francisco. 2014a. 'Finding Volunteers' Engagement Profiles in Human Computation for Citizen Science Projects', *Human Computation*, (Vol. 1, Issue 2) <<https://doi.org/10.15346/hc.v1i2.12>>. Andrew Mao, Ece Kamar, and Eric Horvitz. 2013. 'Why Stop Now? Predicting Worker Engagement in Online Crowdsourcing', in *Proceedings of the AAAI Conference on Human Computation and Crowdsourcing*, pp. 103–11 <https://econcs.seas.harvard.edu/files/econcs/files/mao_hcomp13b.pdf>.
- Ponciano, Lesandro, Brasileiro, Francisco, Simpson, Robert, Smith Arfon and others. 2014b. 'Volunteers' Engagement in Human Computation for Astronomy Projects', *IEEE Computing in Science & Engineering* (Vol. 16, Issue 6) <DOI:10.1109/MCSE.2014.4>.
- Raddick, M. Jordan, Bracey, Georgia, Gay, Pamela L., Lintott, Chris J., Murray, Phil and others. 2010. 'Galaxy Zoo: Exploring the Motivations of Citizen Science Volunteers', *Astronomy Education Review*, 9.1 (2010), 010103 <<https://doi.org/10.3847/AER2009036>>.
- Raddick, M. Jordan, Bracey, Georgia, Gay, Pamela L., Lintott, Chris J., Cardamone, Carie, Murray, Phil, Schawinski, Kevin, Szalay, Alexander S. and Vandenberg, Jan. 2013. 'Galaxy Zoo: Motivations of Citizen Scientists'. arXiv <arXiv:1303.6886>.
- Spiers, Helen, Swanson, Alexandra, Fortson, Lucy, Simmons, Brooke, Trouille, Laura, Blickhan, Samantha and Lintott Chris. 2019. 'Everyone counts? Design considerations in online citizen science', *Journal of Science Communication* (Vol. 18) <<https://doi.org/10.22323/2.18010204>>.
- Sprinks, James, Wardlaw, Jessica, Houghton, Robert, Bamford, Steven, Morley, Jeremy. 'Task Workflow Design and its impact on performance and volunteers' subjective preference in Virtual Citizen Science', *International Journal of Human-Computer Studies* (Volume 104, 2017), pp. 50-63 <<https://doi.org/10.1016/j.ijhcs.2017.03.003>>.
- Sprinks, James, Houghton, Robert, Bamford, Steven, Morley, Jeremy. 2019. 'Planet Four: Craters—Optimizing task workflow to improve volunteer engagement and crater counting performance', in *Meteoritics and Planetary Science* (Vol. 54, Issue 6), pp. 1325-1346. <<https://doi.org/10.1111/maps.13277>>.

Appendix

Classifications by all classifiers [b3f60111c8868f11fcd2f3397ba8ca5b2fc8f55 ## charts...origin/charts]
All workflows in project 'HMS NHS' per weekday and period, in local time (1505 volunteers) [b3f60111c8868f11fcd2f3397ba8ca5b2fc8f55 ## charts...origin/charts]
154178 classifications (1505 volunteers, 154178 classifications, median = 6.0, mean = 102.44 ($\sigma = 506.83$))




Figure 1: HMS NHS classifications in local time. Activity is greatest on weekday mornings and afternoons (8 a.m. to 4 p.m.).

Classifications by all classifiers [b3f60111c8868f11cdf2f3397ba8ca5b2fc8f55 ## charts...origin/charts]
 All workflows in project 'Scarlets & Blues' per weekday and period, in local time (394 volunteers) [b3f60111c8868f11cdf2f3397ba8ca5b2fc8f55 ## charts...origin/charts]
 8935 classifications (394 volunteers, 8935 classifications, median = 3.0, mean = 22.68 ($\sigma = 67.52$))




Figure 2: Scarlets and Blues classifications in local time. Activity is mainly from morning to evening, more or less consistently throughout the week.

Classifications by all classifiers [b3f60111c8868f11cdf2f3397ba8ca5b2fc8f55 ## charts...origin/charts]
 All workflows in project 'RBGE Herbarium' per weekday and period, in local time (357 volunteers) [b3f60111c8868f11cdf2f3397ba8ca5b2fc8f55 ## charts...origin/charts]
 18478 classifications (357 volunteers, 18478 classifications, median = 5.0, mean = 51.76 (σ = 231.92))




Figure 3: The RBGE Herbarium classifications in local time. Activity is again mainly from morning to evening. Tuesday to Thursday see the most activity.

Classifications by all classifiers [b3f60111c8868f11fcd2f3397ba8ca5b2fc8f55 ## charts...origin/charts]
 All workflows of type col-num per weekday and period, in local time (800 volunteers) [b3f60111c8868f11fcd2f3397ba8ca5b2fc8f55 ## charts...origin/charts]
 47372 classifications (800 volunteers, 47372 classifications, median = 5.0, mean = 59.22 (σ = 332.77))




Figure 4: HMS NHS number-type workflows. Activity is more evenly spread over time, compared to the project as a whole (in Figure 1).

Classifications by all classifiers [b3f60111c8868f11cdf2f3397ba8ca5b2fc8f55 ## charts...origin/charts]
 Workflow 08. Years At Sea per weekday and period, in local time (419 volunteers) [b3f60111c8868f11cdf2f3397ba8ca5b2fc8f55 ## charts...origin/charts]
 12258 classifications (419 volunteers, 12258 classifications, median = 4.0, mean = 29.26 (σ = 130.34))




Figure 5: The HMS NHS' "Years at Sea" workflow is more similar to the overall project pattern (Figure 1) than it is to the 'number-type' pattern (Figure 4).

Classifications by all classifiers [b3f60111c8868f11fcd2f3397ba8ca5b2fc8f55 ## charts...origin/charts]
 All workflows of type **col-date** per weekday and period, in local time (517 volunteers) [b3f60111c8868f11fcd2f3397ba8ca5b2fc8f55 ## charts...origin/charts]
 24480 classifications (517 volunteers, 24480 classifications, median = 3.0, mean = 47.35 (σ = 227.83))




Figure 6: HMS NHS 'date-type' workflows resemble the overall project pattern in Figure 1, but show most activity from morning to evening, and less of a preference for weekdays over weekends.

Classifications by all classifiers [b3f60111c8868f11fcd2f3397ba8ca5b2fc8f55 ## charts...origin/charts]
 Workflow 11. Date of Discharge per weekday and period, in local time (277 volunteers) [b3f60111c8868f11fcd2f3397ba8ca5b2fc8f55 ## charts...origin/charts]
 12458 classifications (277 volunteers, 12458 classifications, median = 3.0, mean = 44.97 (σ = 173.88))




Figure 7: The HMS NHS 'Date of Discharge' workflow is similar to the overall date-type chart (Figure 6).

Classifications by all classifiers [b3f60111c8868f11fcd2f3397ba8ca5b2fc8f55 ## charts...origin/charts]
 Workflow 02, Date of Entry per weekday and period, in local time (295 volunteers) [b3f60111c8868f11fcd2f3397ba8ca5b2fc8f55 ## charts...origin/charts]
 12022 classifications (295 volunteers, 12022 classifications, median = 3.0, mean = 40.75 (σ = 227.48))




Figure 8: The HMS NHS 'Date of Entry' workflow has a more 'spread out' pattern of activity, which seems more similar to number-type workflows (Figure 4) than to date-type workflows (Figure 6).

Classifications by all classifiers [b3f60111c8868f11fcd2f3397ba8ca5b2fc8f55 ## charts...origin/charts]
 Workflow People per weekday and period, in local time (75 volunteers) [b3f60111c8868f11fcd2f3397ba8ca5b2fc8f55 ## charts...origin/charts]
 441 classifications (75 volunteers, 441 classifications, median = 1.0, mean = 5.88 ($\sigma = 12.08$))




Figure 9: The Scarlets and Blues 'People' workflow is less active in the morning and more active in the afternoon. The 'Attendance' and 'Minutes' branches of the 'Meetings' workflow are both similar to the overall project chart (Figure 2), which can be used for comparison. However, direct visual comparison is difficult due to the very different heat scale.

Classifications by all classifiers [b3f60111c8868f11fcd2f3397ba8ca5b2fc8f55 ## charts...origin/charts]
 Workflow **Attendance** per weekday and period, in local time (263 volunteers) [b3f60111c8868f11fcd2f3397ba8ca5b2fc8f55 ## charts...origin/charts]
 2340 classifications (263 volunteers, 2340 classifications, median = 2.0, mean = 8.90 ($\sigma = 21.40$))




Figure 10: The Scarlets and Blues ‘Attendance’ branch of the ‘Meetings’ workflow has a similar activity pattern both to the overall project chart (Figure 2) and to the ‘Minutes’ branch (Figure 11). Activity is mainly from morning to evening, more or less throughout the week.

Classifications by all classifiers [b3f60111c8868f11fcd2f3397ba8ca5b2fc8f55 ## charts...origin/charts]
 Workflow Minutes per weekday and period, in local time (284 volunteers) [b3f60111c8868f11fcd2f3397ba8ca5b2fc8f55 ## charts...origin/charts]
 6154 classifications (284 volunteers, 6154 classifications, median = 3.5, mean = 21.67 ($\sigma = 56.39$))




Figure 11: The Scarlets and Blues ‘Minutes’ branch of the ‘Meetings’ workflow has a similar activity pattern both to the overall project chart (Figure 2) and to the ‘Attendance’ branch (Figure 10). Activity is mainly from morning to evening, more or less throughout the week.

Classifications by all classifiers [b3f60111c8868f11fcd2f3397ba8ca5b2fc8f55 ## charts...origin/charts]
 Workflow **Geography** per weekday and period, in local time (300 volunteers) [b3f60111c8868f11fcd2f3397ba8ca5b2fc8f55 ## charts...origin/charts]
 8360 classifications (300 volunteers, 8360 classifications, median = 4.0, mean = 27.87 ($\sigma = 131.91$))




Figure 12: The RBGE Herbarium 'Geography' workflow has a similar activity pattern to the overall project (Figure 3). Note the different heat scale.

Classifications by all classifiers [b3f60111c8868f11cdf2f3397ba8ca5b2fc8f55 ## charts...origin/charts]
 Workflow **Latitude/Longitude** per weekday and period, in local time (92 volunteers) [b3f60111c8868f11cdf2f3397ba8ca5b2fc8f55 ## charts...origin/charts]
 10118 classifications (92 volunteers, 10118 classifications, median = 21.5, mean = 109.98 ($\sigma = 347.70$))




Figure 13: The RBGE Herbarium 'Latitude/Longitude' workflow seems quite different from the overall project pattern, despite representing more than half of the total classifications. Note that the heat scale again is quite different from the other The RBGE Herbarium charts (Figures 3 and 12).

Active volunteers on HMS NHS [b3f60111c8868f11cdf2f3397ba8ca5b2fc8f55 ## charts...origin/charts]




Figure 14a: Active volunteers on HMS NHS across all workflows, displayed as a line graph for ease of distinguishing patterns in the data. 'Name' and 'Years at Sea' show the highest numbers of volunteers, while 'Date of Entry', 'Date of Discharge' and 'Nature of Complaint' all peak at over 40 active volunteers.




Figure 14b: Active volunteers on HMS NHS across all workflows, displayed as a scatter plot. ‘Name’ and ‘Years at Sea’ show the highest numbers of volunteers, while ‘Date of Entry’, ‘Date of Discharge’ and ‘Nature of Complaint’ all peak at over 40 active volunteers.




Figure 15: Active volunteers on Scarlets and Blues. The 'Attendance' and 'Meetings' branches of the 'Minutes' workflow look very similar, while the 'People' workflow attracts fewer volunteers and is completed early.

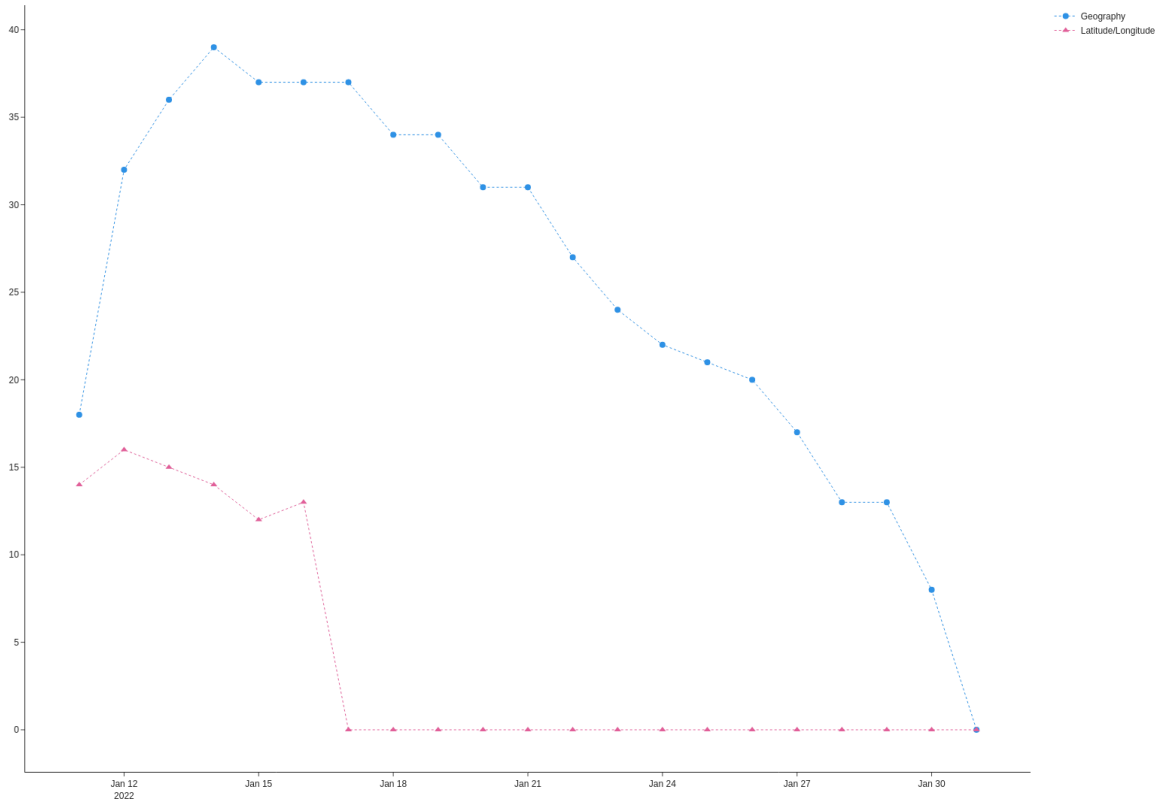


Figure 16: Active volunteers on The RBGE Herbarium. 'Latitude/Longitude' completes quickly. The dramatic drop in 'Geography' may be due to the workflow nearing completion.

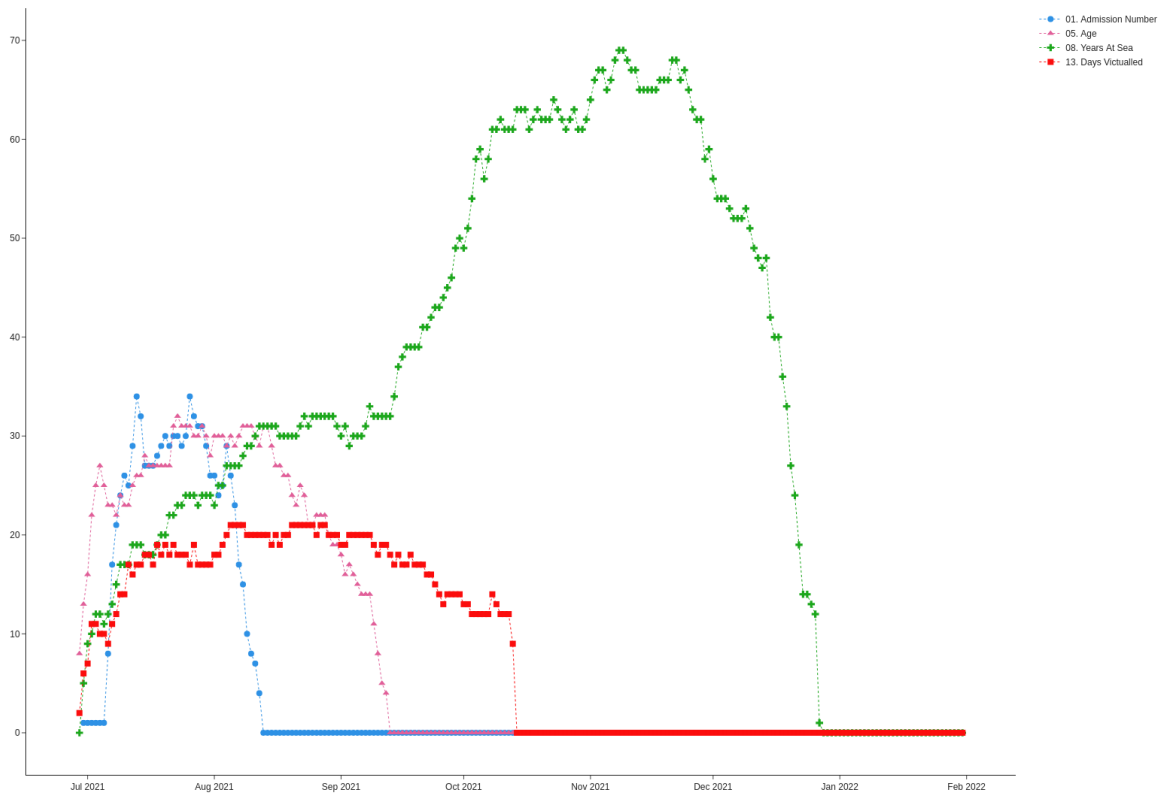


Figure 17: Active volunteers on HMS NHS number-type workflows. 'Years at Sea' climbs to a high peak as the other number-type workflows complete. With the exception of 'Years at Sea', the workflows complete in order of their appearance on the project page.




Figure 19: Active volunteers on HMS NHS dropdown-type workflows. Each attracts equal numbers of volunteers at first, but 'Quality' volunteers surge in September and October.

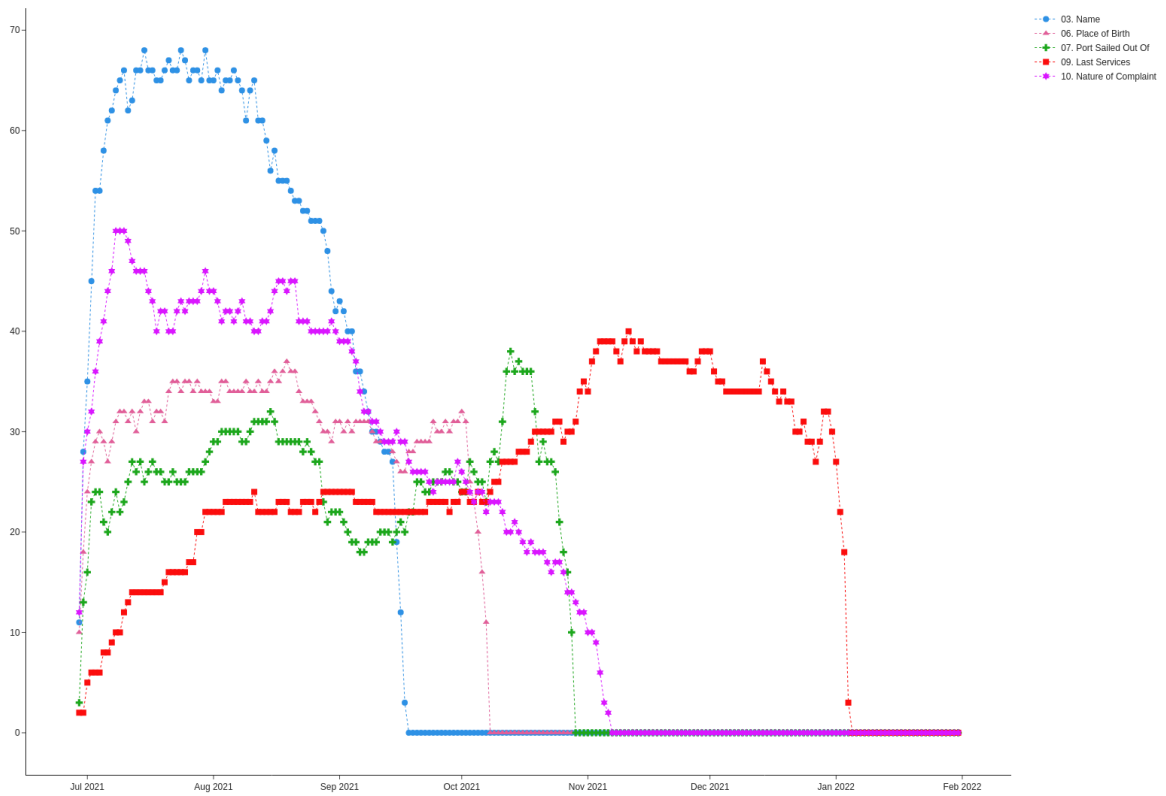


Figure 20: Active volunteers on HMS NHS noun-type workflows. ‘Name’ and ‘Nature of Complaint’ attract the most volunteers early on. ‘Last Services’ attracts the least volunteers early and takes the longest time to complete.