AGILE: GIScience Series, 5, 5, 2024. https://doi.org/10.5194/agile-giss-5-5-2024
Proceedings of the 27th AGILE Conference on Geographic Information Science, 4–7 Sept. 2024.
Editors: Alison Heppenstall, Mingshu Wang, Urska Demsar, Rob Lemmens, and Jing Yao.
This contribution underwent peer review based on a full paper submission.

AGILE Association of Geographic Information Laboratories in Europe

Process Analysis in Humanitarian Voluntary Geographic Information: the case of the HOT Tasking Manager

Dagoberto José Herrera-Murillo [1], Héctor Ochoa-Ortiz [2], Umair Ahmed [2], Francisco Javier López-Pellicer [1], Barbara Re [2], Andrea Polini [2], and Javier Nogueras-Iso [1]

[1] Aragon Institute of Engineering Research (I3A), Universidad de Zaragoza, 50018 Zaragoza, Spain
[2] PROS - PROcesses & Services lab, Computer Science Division, School of Science and Technology, University of Camerino, 62032 Camerino (MC), Italy

Abstract.

Geographic information is vital for organising humanitarian campaigns and helping those in need. The leading Humanitarian OpenStreetMap Team (HOT) organises projects to create the necessary geographical information and connect to the organisations that need to make decisions on the ground. This work provides insights into project management dynamics and volunteers' interaction with user interfaces in Volunteered Geographic Information (VGI) in a humanitarian context. We do so by conducting a process analysis of 746 completed, fully validated, and archived projects in the HOT Tasking Manager (HOT-TM) over the past two years. The analysis encompasses a process discovery stage from the perspectives of control flow, time, organisation, and outcome of the mapping tasks that comprise a project. The findings offer valuable implications for future project planning and execution in similar contexts. Our process mining exploration of the task states found a clear path that involves mapping and validation operations with minor deviations. However, we did find a major bottleneck from the mapping to the validation phase, which could reflect that validation capabilities are a scarce resource. Proactive notification for validators, artificial intelligence adoption for task planning, user interface redesign, and strategies for better harnessing the collective intelligence of volunteers could improve the process.

Keywords. OpenStreetMap, Volunteered Geographic Information, Humanitarian OpenStreetMap Team, Volunteer Engagement, Humanitarian Campaigns, User Interface, Process Mining

1 Introduction

The emergence of digital mapping and crowd-sourced geographic information has significantly reshaped the landscape of humanitarian aid and disaster response. Volunteered Geographic Information (VGI) is a method that leverages the collaborative contributions of volunteers to collect, analyse, and distribute geographic data. This approach is particularly useful in scenarios where conventional data sources are neither accessible nor longer up-to-date (Goodchild, 2007), and has been proven to be a game-changer in managing humanitarian crises, where timely and accurate geographic information is crucial (Neis et al., 2010; Saganeiti et al., 2017).

The Humanitarian OpenStreetMap Team (HOT) is a leading initiative in the field of Humanitarian VGI. HOT uses the OpenStreetMap (OSM) platform to facilitate the creation of freely editable global maps, with a particular focus on areas in urgent need of humanitarian aid[1]. The HOT Tasking Manager (HOT-TM), which organises the global volunteer network to facilitate effective and concentrated mapping activities, is an important tool in this effort.

To leverage collaborative mapping projects effectively, HOT must understand those dynamics that facilitate or block mapping and validation processes. Previous efforts have focused on understanding some of the critical success factors of humanitarian mapping projects, such as levels of volunteer experience (Urrea and Yoo, 2023). However, the available literature has not delved into the fine grained details of how the mapping process works. Nonetheless, detailed guidelines into how HOT should design tasks or assign resources to maximise the performance of the process have yet to be developed.

For these reasons, the objective of this paper is to propose a methodology for making process analysis on HOT projects. Process analysis aims at providing decision-makers with the knowledge they need to streamline operations efficiently. In particular, we investigate how effectively the mapping process occurs using a sample of 746 completed and archived projects within the HOT-TM over the last two years. We focus on discovering this process from four perspectives: (1) control flow, i.e. the ordering

[1] https://www.hotosm.org/

of events, (2) time, i.e. the duration of processing and idleness in the execution of activities, (3) organisation, i.e. the people who execute or initiate the events, (4) outcome, i.e. the geographical product that is finally embodied in OSM. The final purpose is to speed up the completion of tasks, reduce waiting times and ultimately increase overall efficiency. These improvements are of great importance, especially in the context of humanitarian emergencies.

The remainder of this paper is structured as follows. Section 2 presents the related work. The proposed methodology is explained in section 3. Section 4 describes the results obtained after applying the methodology for the analysis of the OSM projects. Finally, section 5 provides some conclusions and outlines for future research work.

2 Related Works

Volunteered Geographic Information (VGI) use in disaster management and humanitarian aid has emerged as a critical area of research, with several studies contributing to its understanding and application.

Rey (2022) focus on the application of spatial information in disaster risk management. Their study emphasises how important spatial data is for improving the efficacy of disaster risk reduction tactics. Xin (2022) explored anomaly detection for VGI, using Safecast data as a case study. This research demonstrates the potential of VGI in monitoring and responding to environmental hazards.

Ke et al. (2023) examine intelligent management strategies for emergency shelters and resilient communities during disaster scenarios. Tzavella et al. (2022) conducted a comprehensive review of the application of VGI in crisis, emergency, and catastrophe management, highlighting the importance of VGI in facilitating emergency decision-making. This study highlights the evolving role of VGI in addressing the dynamic challenges in disaster situations.

Zhao et al. (2022) investigated the use of shared geospatial data in disaster management. Their research contributes to the understanding of data-sharing mechanisms and their role in improving the efficiency of disaster response operations. Safariallahkheili and Malek (2022) introduce a novel approach for assessing the reliability of volunteered geographic information during flood emergencies. The authors highlight the importance of reliable VGI in facilitating precise and quick decision-making in flood-related catastrophes. The study proposes a systematic approach to validate the information provided by volunteers.

Hosseinali and Farhadpour (2020) develop a spatial solution integrating VGI and genetic algorithms to improve earthquake crisis management in Tehran, Iran. It demonstrates how VGI can be effectively used in conjunction with computational methods to optimise resource allocation and response strategies following an earthquake. Arapostathis (2020) provide a foundational understanding of VGI in disaster management, particularly in the context of floods. This study offers insights into the use of VGI for flood risk assessment and response planning.

El Hatimi et al. (2020) focus on the quality assessment of VGI in the context of risk management applications. It addresses the challenges of verifying the accuracy and reliability of VGI, proposing methodologies to ensure the quality of data used in managing various risks, especially in crisis scenarios. Hardy (2020) delves into the use of VGI for tracking purposes in crisis management. The study discusses the potential of VGI to provide critical, up-to-date information for tracking developments in real-time during crises, highlighting its application in various emergency scenarios.

Zhao et al. (2019) discuss the extraction and classification of typhoon disaster information based on VGI sourced from Sina Weibo, the Chinese microblogging platform. This study underscores the utility of social media platforms as sources of VGI in disaster scenarios, providing real-time, location-specific information. Haworth et al. (2018) investigate the dual nature of VGI contributions to community disaster resilience. It explores how VGI can empower communities with real-time information during disasters, while also discussing the potential pitfalls and uncertainties associated with reliance on such data sources.

Urrea and Yoo (2023) examine the influence of volunteers' experience on achievement and retention rates in online platforms. Using data from HOT-TM, the study demonstrates that volunteer experience improves project completion rates. However, the study also finds that the impact of volunteer experience diminishes over time and varies depending on the urgency of the project. Additionally, it finds that experience-based incentives initially boost volunteer retention, but their influence wanes as volunteers gain more experience.

These studies collectively illuminate the multifaceted role of VGI in disaster management, ranging from data collection and analysis to its application in risk assessment, response coordination, and damage estimation. The insights gained from these works are instrumental in shaping future strategies and approaches in the field of humanitarian aid and disaster response.

3 Methodology

As depicted in Figure 1, the methodology proposed to analyse the mapping process involves three distinct phases: (1) understanding HOT-TM; (2) data collection and pre-processing; and (3) data analysis. The following subsections describe these phases.

3.1 Understanding HOT-TM

Before analysing raw data from mapping projects, it is essential to understand the context and operation of HOT-

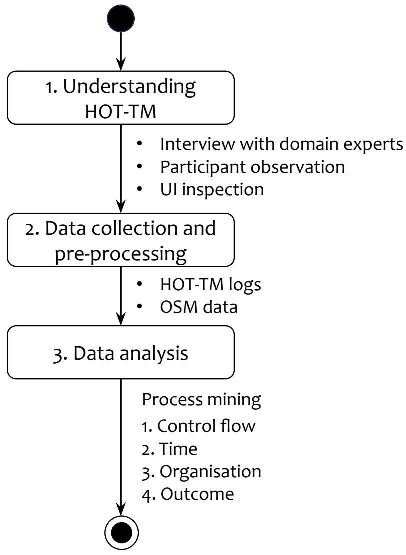

Figure 1. Methodology at a glance.

TM[2], the tool used by HOT to coordinate volunteer action in humanitarian projects.

To this end, three qualitative research techniques were used: in-depth interviews with domain experts, participant observation, and individual expert review of the user interface.

We interviewed a member of the HOT Open Mapping Hub for Latin America and the Caribbean (LAC Hub), and a member of the Médecins Sans Frontières (MSF) GIS Team. Given the exploratory nature of this phase, the interviews followed a semi-structured format with a guide of possible open-ended questions prepared by the research team.

The participatory observation involved the attendance of three research team members in an online mapathon organised by HOT, one of the researchers already had extensive experience in humanitarian mapping missions and OSM, and the other two were first-timers. The event was two hours long. The experienced mapper participated in the validation of tasks, reserved for expert contributors. In contrast, the novices first received a 30-minute explanation of HOT-TM and the iD Editor, which is an embedded editor for beginner users, and started to perform basic mapping operations afterward.

As part of the interface inspection, an expert scanned the interface to capture interface elements associated with specific actions in the mapping and validation phases.

In HOT-TM, each project is subdivided into tasks, in a grid pattern. Figure 3 shows a screenshot of a project page, with a map displaying the different tasks with their current state.

[2]https://github.com/hotosm/tasking-manager

Each task within the HOT-TM system can be categorised into specific states, reflecting its current state in the mapping process. They are described as follows.

- **LOCKED FOR MAPPING:** Tasks that are being actively worked on by mappers are locked to prevent overlapping work. This task has an associated locking time.
- **AUTO-UNLOCKED FOR MAPPING:** Locked for mapping tasks get unlocked due to timeout, usually after 2 hours. The associated LOCKED FOR MAPPING task does not appear in the log.
- **MAPPED:** This state represents tasks where the last mapper who locked for mapping considers the work to be completely mapped, but is yet to be validated.
- **LOCKED FOR VALIDATION:** Tasks enter this state when they are ready for validation, preventing further mapping edits. This task has an associated locking time.
- **AUTO-UNLOCKED FOR VALIDATION:** Locked for validation tasks get unlocked due to timeout, usually after 2 hours. The associated LOCKED FOR VALIDATION task does not appear in the log.
- **VALIDATED:** Tasks confirmed to meet the mapping standards are marked as validated.
- **INVALIDATED:** Tasks that do not meet the criteria during validation are marked as invalidated and may require remapping or further review.
- **SPLIT:** This state occurs when tasks are divided into smaller parts for easier management or to address complex mapping areas. When splitting a task, 4 new tasks are created. The original task disappears from the log and its states before the split are copied to the new tasks. The associated LOCKED FOR MAPPING task does not appear in the log.
- **BAD IMAGERY:** Tasks are given this state when the underlying imagery is insufficient for accurate mapping.

Figure 2 visually explains the workflow of the mapping process, with the different states for a task. The numbers given in the diagram can be associated with different interactions in HOT-TM, as seen in the interface screenshots (Figures 3, 4 and 5).

The mapping process can be divided into two phases: Mapping and Validation. In the mapping phase, users select a task that is ready to be mapped or get one assigned randomly. An embedded editor is open inside HOT-TM (there is also the possibility to open a desktop editor instead), where users have to draw the missing elements, according to the project description (a typical example would be to draw buildings or roads), without drawing outside

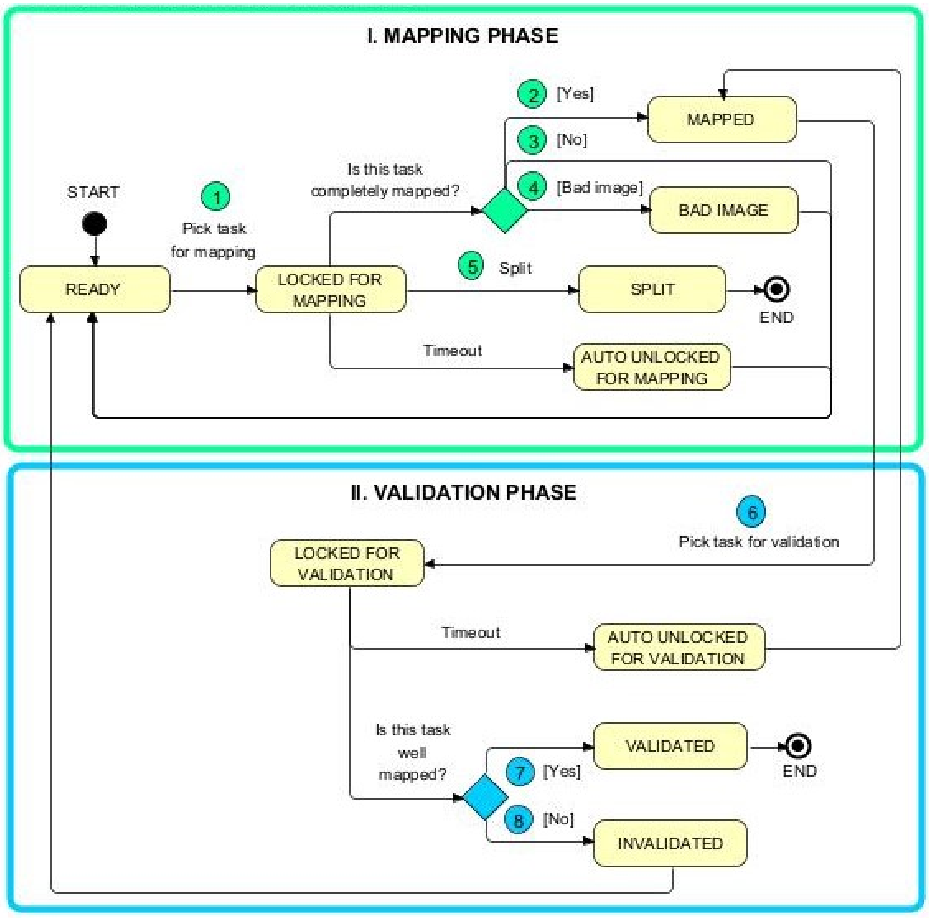

Figure 2. Task state diagram illustrating the mapping workflow. The numbers represent the different decisions made on the interface and correspond to the numbers on the screenshots.

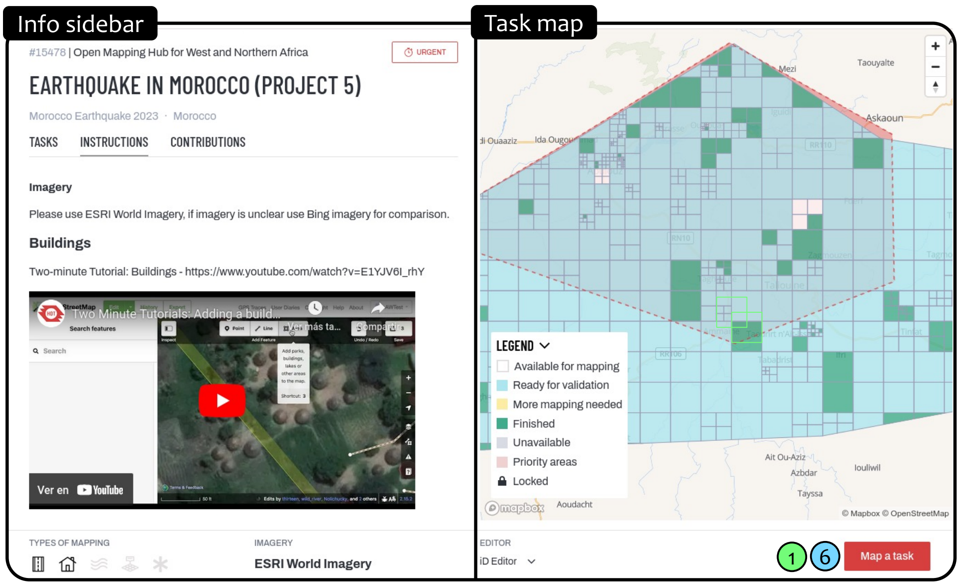

Figure 3. Screenshot of a project tasks screen in HOT-TM. On the left sidebar, information about the project appears, with task state on a list, instructions for completing each task, and statistics about the contributors. On the right, there is a map showing the task state in a more visual way. Users can either select a random task to map with the bottom right button, or manually select a task on the map, after which the button will then change to either map or validate the selected task. Source: https://tasks.hotosm.org/projects/15476/tasks/

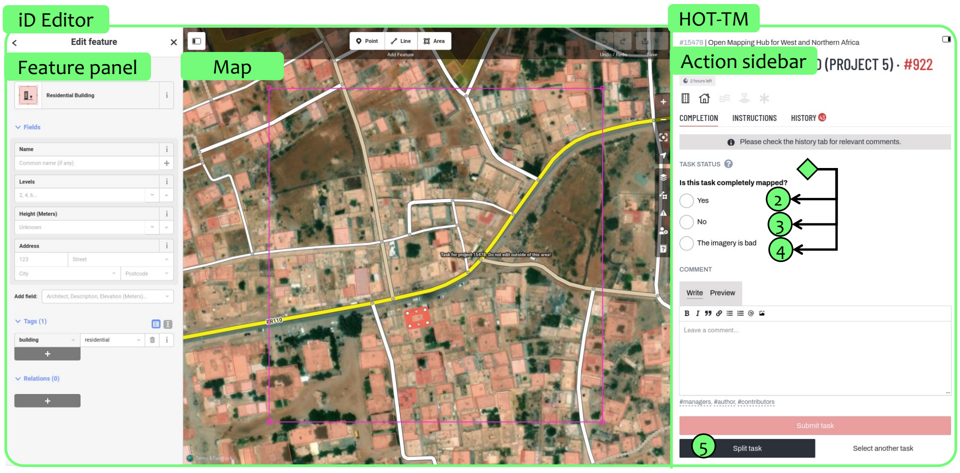

Figure 4. Screenshot of the mapping phase in HOT-TM, using the embedded iD editor. The iD editor changes will be reflected in the OSM database after hitting the Save button inside it. The user interaction with the action sidebar will trigger changes in the task state. HOT-TM and the iD editor are unrelated and do not communicate between them. The user is expected to edit and save the changes on the editor as many times as needed and then interact with the sidebar after they are done editing.

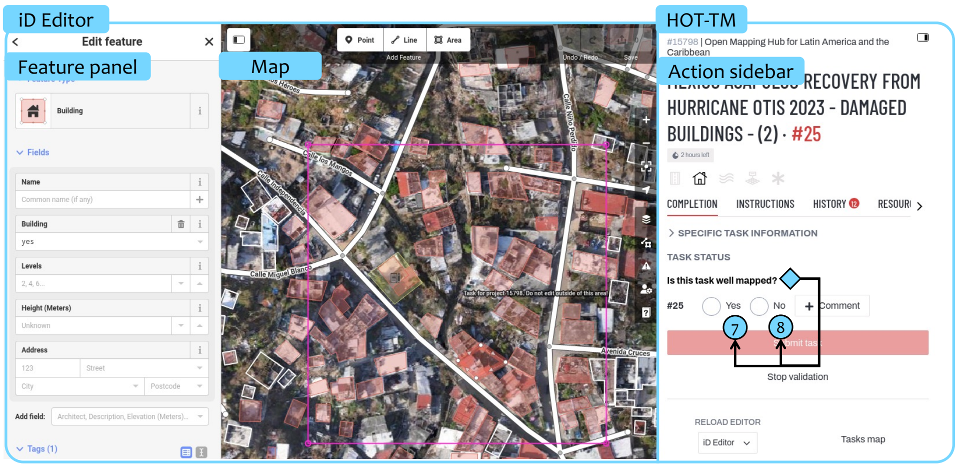

Figure 5. Screenshot of the validation phase in HOT-TM, using the embedded iD editor. It is equal to the mapping phase interface, but the action sidebar buttons change. The user interaction with the action sidebar will trigger changes in the task state. The user is expected to navigate through the data on the editor to check its quality and edit it if needed. Then interact with the action sidebar after they are done validating.

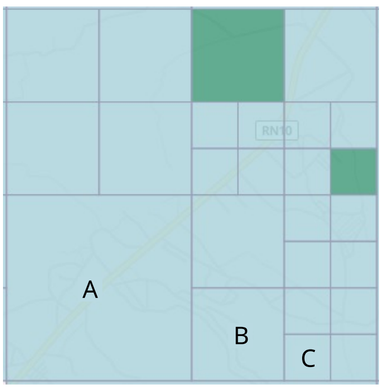

Figure 6. Screenshot of tasks in a project. Task A has not been split, task B has been split once, and task C has been split twice.

the marked area. The user then has the ability to select whether the task is, according to their criteria, completely mapped, not completely mapped, the imagery is bad or the task should be split (see Figure 4). Splitting a task creates 4 new tasks with a smaller area, as seen in Figure 6. This splitting can be repeated multiple times.

The validation phase presents validators with tasks marked as completely mapped. An editor is opened in a similar fashion to the mapping phase, but different buttons appear in the action sidebar (see Figure 5). The validator then decides if the task is completed (well mapped), moving the task to the validated or invalidated state. Invalidated tasks move back to the mapping phase for further mapping.

3.2 Data Collection and Pre-processing

We collected data on tasks from 746 completed and archived HOT-TM projects, created over the last two years. Table 1 shows the distribution of projects according to their level of difficulty and the regional centre where they take place. Two-thirds of these projects are classified as easy, one-third as moderate, and difficult projects are rare. Eastern and Southern Africa is the hub with the most associated projects, followed by Latin America and the Caribbean, West and Northern Africa, and Asia Pacific. Projects in countries outside the scope of the hubs are the least frequent. The size of the project is another factor to be taken into account. The final number of tasks per project reveals a wide variation ($mean = 418.6$, $sd = 487.3$, $min = 8.0$, $Q1 = 107.3$, $median = 286.0$, $Q3 = 564.0$, $max = 4483.0$).

Difficulty	n°	%
Easy	501	67.2
Moderate	239	32.0
Challenging	6	0.8
Hub	**n°**	**%**
Eastern & Southern Africa	309	41.4
Latin America & Caribbean	194	26.0
West & Northern Africa	129	17.3
Asia Pacific	82	11.0
Other	32	4.3

Table 1. Distribution of projects according to difficulty and Hub.

Data was collected from two different APIs: HOT-TM and Bunting Labs. The HOT-TM API provides detailed information about projects, the states occurring in the projects, and users associated with those events. Later, task grids served as a parameter for the Bunting Labs API to return the building datasets of each mapped task based on the data available in OSM. The task grid and building datasets were reprojected from geographical coordinates to UTM before calculating their areas. Buildings are overlaid onto the grid, to calculate the percentage of each cell covered by buildings.

The pre-processing of the data was aimed at constructing an event log (See table 2) that would subsequently allow process mining techniques to be applied. In a typical event log several key assumptions are met: a) a process consists of cases -in this analysis the cases are specified at the task level-; b) a case consists of events and each event refers exclusively to one case, c) events within a case are displayed in order, d) events can have attributes to model the process. The set of possible states that represent the events are shown in Figure 2, except that the READY state is not explicitly recorded in the log.

It is assumed that the tasks retain READY from the start of the mapping phase until the point at which they become mapped. The start and completion timestamps of the events are available. The resource dimension is represented by the user who performed the events and its mapping experience. Finally, attributes about the area of the task and its building coverage according to OSM are incorporated.

Data from the APIs is collected and saved into files, which then get accessed for further steps. The collection and preprocessing notebooks also store intermediate states in files whenever possible, to recover from closing or failing in the middle of execution, especially in the resource-intensive calculations.

It is also important to note that the two-year period selected for this study provides a snapshot of projects with relatively consistent user skill levels, as we were limited to the current user-level data. We chose this duration based on the assumption that OSM users typically advance from intermediate to higher skill levels within four years (Bégin et al., 2018).

3.3 Data analysis

Once the event log was extracted, we employed process mining methods to analyse the data. The process mining exercise concentrated on discovering the mapping process from four perspectives:

- **CONTROL FLOW:** This perspective considers the ordering of events. It aims to find a good characterisation of all possible pathways. The analysis included frequency and case coverage of task states, frequent process variants and a directly-follows graph to show transitions between states.

- **TIME:** This perspective deals with the timing and frequency of events. It allows the discovery of bottlenecks, measurement of service levels, and control of resource utilisation. The analysis included a directly-follows graph to show the median duration of states and transitions. The temporal analysis also distinguishes between processing time (the duration of all instances of active mapping or validation states) and idle time (the time when no instances of mapping or validation states are active).

- **ORGANISATION:** This perspective focuses on the actors present in the record (e.g. people, systems, and roles) and how they are related. Our analysis paid particular attention to the mapping level of the contributors. For each type of state in the mapping phase, we calculated the frequency of execution per mapping level. Validation events were excluded because by default they are not enabled for beginners.

- **OUTCOME:** This perspective takes advantage of the available data on the geographical product resulting from a task. We used the percentage of area covered by buildings as a proxy of the outcome from a HOT-TM project because the vast majority of projects in-

projectId	taskId	state	start	complete	actionBy	mappingLevel	areaSqkm	buildingCoverage
15796	70	SPLIT	2021-12-01 04:45:26	2021-12-01 04:46:29	Redacted	ADVANCED	0.24	2.55

Table 2. Example of how a row of the generated event log looks like, after our pre-processing.

clude building mapping. A zero-inflated beta regression model was used on task data to describe the percentage of area covered by buildings as a function of the number of times the task was locked for mapping, whether instances of splitting, invalidation, and bad imagery were observed (binary encoding), and controlling for the area of the task expressed in square metres (area_sqm), and the difficulty of the project to which the task belongs (EASY, MODERATE, CHALLENGING). Other states were excluded because their executions were not necessarily independent.

By examining the log it is possible to discover that when a SPLIT occurs the logging system creates four new tasks by replicating for each of them the events corresponding to the initial task before the split and adding the new events afterward. The original task is deleted. Duplication introduces noise that affects the event count and makes it difficult to interpret directly-follows graphs. For this reason, a cleaning routine was deployed to retain only the initial task logs in the first three perspectives that require these techniques.

Following common practice, the directly-follows graphs shown in the control flow and time perspectives were pruned with a trace frequency filter to facilitate the visualisation of the flow of states. Thus, using a cut-off threshold of 0.95, we will select at least 95% of the cases, starting with those that have the highest frequency. For the outcome perspective, final tasks (i.e., the set including the tasks resulting from the splits) are maintained.

Process discovery was performed using *bupaR*, a suite of open-source R packages for business process data analysis.

4 Results

To facilitate the understanding of the paper, we already explained in section 3.1 the conceptual dynamics of the mapping phases in HOT-TM. Therefore the first part of the results just focuses on reporting the feedback received by interviews and the second on data analysis with process mining.

4.1 Understanding HOT-TM

The pair of interviews with domain experts provided valuable insights into the preparation, field operation, and subsequent exploitation of the data resulting from a humanitarian mapping project.

According to the LAC Hub member, HOT-TM mapping process relies heavily on OSM mappers from the Global North contributing as a hobby, but they want to rely more on local mappers by investing in their training and engagement. HOT is now undergoing a transformation to empower contributors of the regions that need to be mapped, with the concept of regional Hubs. Local people better know areas in need, and can better interpret the local context when mapping. Also, the hubs aim to coordinate a local network to transform the gathered data into useful help, by connecting to relevant organisations on the terrain.

Regarding the HOT-TM tool, data should be validated before use, and that is their current bottleneck because of a lack of resources. They believe that data quality depends on the individual mapper, so some critical projects have a minimum experience level set for users to map on it. They also stated that usability should be improved on the portal, to make it easier for mappers, including the possibility of mapping from a mobile phone.

The MSF member explained that OSM and HOT-TM are relevant inputs for their operations. They contact the HOT community and set up projects based on MSF's necessities on the terrain. They trust the edits made by the users but believe that beginners may have more invalid edits, as they are new to it. Also, the quality of the underlying imagery used for edits matters and changes the final result. With their experience in in-person mapathons, some users do not read the instructions, do not map correctly, struggle to save their progress or forget to unlock the tasks. According to them, retention of beginner users and a balanced mix between beginner and advanced users is the key to a project success. Another thing that they mentioned, is that they have manual and automatic validation mechanisms set in place, to make sure the data is usable.

4.2 Data analysis

In this section, we present the result of the analysis of the event logs generated during the task mapping and validation from four perspectives: control flow, time, organisation, and outcome.

4.2.1 Control flow

The frequency and case coverage of the task states shown in Table 3 suggest that LOCKED FOR MAPPING is the most frequent states and is present in almost all cases. LOCKED

FOR VALIDATION, MAPPED, and VALIDATED are also frequent states with high case coverage. The occurrence of the remaining types of states is relatively infrequent. Only 2.8% of the initial tasks become split, but new tasks make up about one-tenth of the total number of tasks at the end of mapping.

States	nº	Task coverage (n=277,745)
LOCKED_FOR_MAPPING	406,324	98.8%
LOCKED_FOR_VALIDATION	308,244	97.7%
MAPPED	284,214	97.6%
VALIDATED	284,035	97.2%
AUTO_UNLOCKED_FOR_MAPPING	20,070	5.3%
INVALIDATED	9,429	3.0%
SPLIT	7,680	2.8%
BADIMAGERY	4,502	1.5%
AUTO_UNLOCKED_FOR_VALIDATION	2,956	1.0%
EXTENDED_FOR_MAPPING	2,801	0.3%
TOTAL	1,330,255	

Table 3. Frequency and case coverage of task states.

Figure 7 shows an absolute frequency map, where the nodes indicate the absolute number of state instance executions and the edges indicate the absolute number of times the source and target states were executed directly after each other. The thickness of the edges and the colour saturation of the vertices is proportional to a higher frequency. The flow reveals a main route for mapping tasks that involves users locking the task for mapping in one or more iterations until mapped state is reached, mapped tasks are then locked for validation and validated on the first attempt. This is also confirmed by the most common variants shown in Figure 8. Only a small proportion of tasks lead to invalidation (3%) and re-enter the mapping cycle until validation is achieved.

4.2.2 Time

The time perspective is shown in Figure 9, where the nodes and the edges indicate the median duration of states and waiting times respectively. LOCKED FOR MAPPING, LOCKED FOR VALIDATION, AUTO-UNLOCKED FOR MAPPING, and AUTO-UNLOCKED FOR VALIDATION are the only states with duration. With median execution times of 2.2 and 4 minutes for the first two states and 2 hours for the remaining two. The other states are immediate execution flags. Considering the transition frequency, the main bottleneck in the process is the time a mapped task has to wait for validation, which has a median of almost 14 days. Generally speaking, the duration of a task consists mainly of idle time. The median proportion of idle time of a task is 99.94% of its throughput time. The actual service time spent on assignment and validation is minuscule.

4.2.3 Organisation

Table 4 displays the profile of the mapping phase states according to the mapping level of the contributor performing them. The first row of the table contains the composition of the total number of contributors of the analysed projects according to their mapping level. The lower part of the table shows the breakdown of state execution frequency per mapping level. Note that the percentages correspond to the total of each row.

The first relevant finding is that the overwhelming majority of contributors to the mapping projects are beginner users in a ratio of 9:1 to more advanced users. However, when reviewing the output of each group in terms of operations performed during mapping, it is possible to appreciate that the activity of a standard advanced or intermediate user significantly exceeds that of a standard beginner. The result is that only half of the total volume of events is executed by beginner users. Approximately the same proportion as reported by LOCKED FOR MAPPING states. For the remaining state types the average advanced or intermediate user exceeds the execution frequency of novice users, but the propensity varies significantly by state type. Of the total volume of AUTO-UNLOCKED FOR MAPPING and BAD IMAGERY transactions, a clear majority are the product of novice user actions, while SPLIT transactions come from advanced users.

4.2.4 Outcome

The distribution of the percentage of area covered by buildings reveals that the vast majority of tasks have a small coverage, with some atypically high observations ($mean = 1.91\%$, $sd = 5.57$, $Q1 = 0\%$, $median = 0.06\%$, $Q3 = 0.92\%$, $max = 100\%$).

Table 5 shows the results of the zero-inflated beta regression model that was used on the task data to describe the percentage of area covered by buildings. By examining the magnitude and polarity of the coefficients, it is possible to identify that the number of mapping operations and the presence of splits, invalidations, or bad images are associated with higher coverage. The larger the area of the resulting task, the smaller the mapped area. Projects of easy difficulty are associated with higher coverage and those of moderate difficulty with lower coverage compared to the challenging difficulty that served as a reference category.

Among the state variables, the SPLIT flag has the strongest connection to higher building coverage. Figure 10 shows the distribution of the percentage of area covered by buildings as a function of whether the SPLIT state was observed during task execution or not.

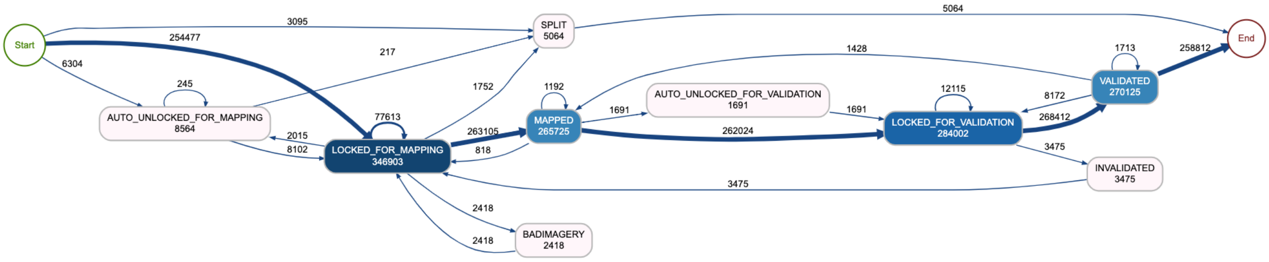

Figure 7. Frequency map of task states and transitions within the HOT-TM.

		Advanced		Intermediate		Beginner	
	N°	%	n°	%	n°	%	n°
Total contributors	32,991	6.7%	2,218	2.6%	852	90.7%	29,921
			Avg.		Avg.		Avg.
Total states	722,790	44.5%	145.6	6.4%	54.1	49.1%	11.9
Locked for Mapping	406,324	42.6%	78.3	6.7%	31.7	50.8%	6.9
Mapped	284,214	48.7%	62.6	6.1%	20.2	45.2%	4.3
Auto Unlocked for Mapping	20,070	19.2%	1.7	6.3%	1.5	74.4%	0.5
Split	7,680	66.9%	2.3	5.3%	0.5	27.8%	0.1
Bad Imagery	4,502	28.1%	0.6	3.8%	0.2	68.1%	0.1

Table 4. Breakdown of state execution according to the contributor mapping level.

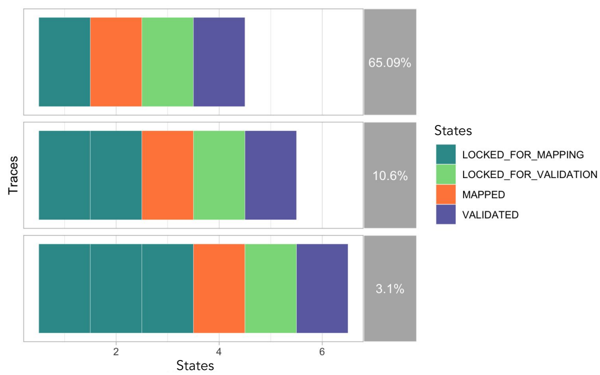

Figure 8. Top 3 process variants (% of tasks).

Mu Coefficients	Estimate	t value	
(Intercept)	-3.49	-509.31	***
split_yes	0.44	57.92	***
inval_yes	0.16	14.47	***
bad_imagery_yes	0.27	18.51	***
locked_for_mapping	0.02	24.13	***
area_sqm	-1.4e-08	-25.35	***
difficultyEasy	8.7e-03	2.08	*
difficultyModerate	-0.31	-44.67	***
Single term deletions	**Df**	**LRT**	
split_yes_no	1	3099.8	***
inval_yes_no	1	206.8	***
bad_imagery_yes_no	1	322.8	***
locked_for_mapping	1	538.8	***
area_sqm	1	786.2	***
difficulty	2	3566.7	***
Signif. codes: 0 '***' 0.001 '**' 0.01 '*' 0.05 '.'			

Table 5. Zero-inflated beta regression model for the percentage of area covered by buildings.

5 Discussion and Conclusion

This paper has presented a detailed process analysis of 746 completed and archived projects in HOT-TM over two years. The analysis of project details, task states, user contributions, and temporal aspects of participation has revealed significant patterns in how volunteers engage with tasks. These findings underscore the importance of understanding volunteer behaviour to optimise the efficiency and impact of humanitarian mapping initiatives. Our study contributes to the broader field of humanitarian aid and disaster response by highlighting the critical role of VGI and the effectiveness of platforms like HOT in coordinating these efforts.

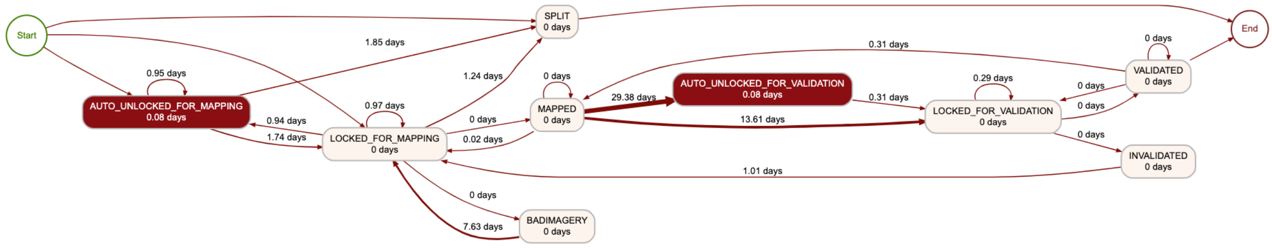

Figure 9. Time map of task states and transitions within the HOT-TM.

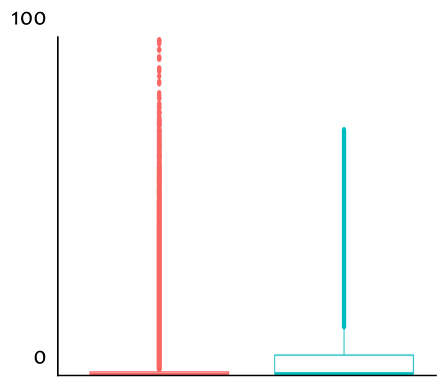

	split_no	split_yes
n°	269,775	42,033
Min.	0.00	0.00
Q1	0.00	0.00
Median	0.05	0.32
Mean	1.44	4.93
Q3	0.74	5.79
Max.	100.0	72.97

Figure 10. Distribution of the percentage of area covered by buildings of tasks (Box plot).

Our process exploration revealed a smooth path between mapping and validation with minor deviations. However, we have found that the time devoted to validation tends to be longer than the effective time devoted to mapping tasks. This could indicate either that the editing is quite straightforward or that validators are spending time fixing wrong or incomplete edits while validating. A way to check this would be by analysing the OSM changesets (a group of edits to the OSM database made by a user, uploaded with metadata such as a brief comment explaining the performed changes, hashtags, editor and imagery used, and timestamps) associated to each mapping and validation states. As seen in figures 4 and 5, the editor and HOT-TM are two unconnected pieces of software. Therefore the association between task states and OSM changesets is not always straightforward.

The time a mapped task has to wait to be validated -a median duration of almost 14 days- is a major bottleneck in the process. This finding confirms the insight received from the LAC Hub member who indicated that validation capabilities are a scarce resource. Possible validators may not necessarily be aware that they can validate in a given project as it is not shown clearly on the interface.

Advanced contributors are more active than beginners and are more prone to decisive tasks like MAPPED and SPLIT. While beginners are more associated with events like AUTO-UNLOCKED FOR MAPPING. The propensity for auto-unlocking could indicate that beginners do not fully understand the interface and unknowingly leave the task locked.

More instances of LOCKED FOR MAPPING and the presence of SPLIT, and INVALIDATED states are related to higher building coverage. That may be explained by areas having a greater editing complexity. Anticipating complexity could be key to having more efficient mapping and resource allocation.

Our study has several limitations due to the nature of the used data. The HOT-TM logs and OSM datasets, and the interfaces of HOT-TM and OSM editors are constantly changing and developing, making conclusions applicable to the point where the data download was made. In addition, the availability of user levels on the HOT-TM API was limited to the moment of download, not having a history of user levels, to associate with the moment the user contributed to each project. It could be, for example, that an advanced or intermediate user was a beginner at the time they participated in a given project. We tried to mitigate this effect by selecting data from archived projects created in the last two years.

As another limitation, the study relies on quantitative data from project details, task states, and user contributions. It may overlook the qualitative aspects of volunteer experience, motivation, and challenges.

The lessons learned in this case can inform the design of other collaborative mapping applications within the

OSM ecosystem, whether it is to improve the performance of existing applications (e.g. MapRoulette[3], MapSwipe[4], StreetComplete[5]) or to inspire the design of in-development (e.g. fAIr[6], Field Mapping Tasking Manager[7]) and future platforms. All of these applications must take into account the great impact that the experience of contributors has on the quantity and quality of their interactions with the system. Effective onboarding of new users is a major challenge for the design of such interfaces. This last observation presumably has major implications for the design of micro-tasks in other areas beyond VGI.

As per future work, we suggest several research lines. Firstly, we propose improvements to the validation process, to remove the observed bottleneck. This would see possible research in the direction of how to get higher retention mechanisms for achieving an increased pool of validators, how to improve the HOT-TM interface to make validators aware of projects that need their attention, and how to use the collective intelligence of the big pool of beginner users to validate.

Secondly, we suggest to improve the HOT-TM interface with user-centered design techniques, including testing. As seen in the mapathon we participated in, the first 30 minutes of it are devoted to explaining the project and the interface to beginners. A more intuitive interface can allow for less time in explanation, allowing more mapping to happen.

Thirdly, we propose to examine the effects of the HOT-TM and editor interfaces in the generated OSM data. HOT-TM separates the area of a project into a grid, which could be mapped by different users. This may be affecting the overall homogeneity of the data between the grid cells and at their borders. On top of that, another study could be made to check for differences between mappers who used iD, JOSM, or other editors in a humanitarian context, and the effects they could have on the generated geodata. In addition, and as said before, the HOT-TM and iD editor interfaces should be integrated more and communicate between them. This could mean storing in the HOT-TM logs the OSM changesets associated with each task state.

Lastly, the use of artificial intelligence could improve the efficiency of the humanitarian mapping process. For instance, preemptively splitting complex project areas using AI-generated building cover datasets from satellite imagery, to lower the invalidation rates in these areas. These datasets are already being used in some editors like Rapid to speed up the editing process (Rapid Editor, 2023), but have not been used in the task creation. AI can also play an instrumental role in coordinating the group of volunteers suggesting tasks to users based on their level of experience.

Furthermore, the insights gained from this research offer practical implications for future project planning and execution. By understanding the nuances of volunteer engagement and task management, organisations involved in humanitarian mapping can better strategise their efforts, ensuring that resources are utilised effectively and that the generated maps are of the highest quality and relevance. Further research on user experience design to drive volunteer motivation also offers great potential.

In conclusion, the patterns identified in this study not only enhance our understanding of volunteer-driven humanitarian mapping but also provide a foundation for future research in this area. As the field of VGI continues to evolve, studies like ours will be crucial in shaping strategies that harness the full potential of crowdsourced mapping in addressing global humanitarian challenges.

6 Data and software availability

Due to privacy concerns with the user data, we have decided not to publish the used data. However, a Jupyter notebook is provided to download the data from the HOT-TM and Bunting Labs APIs. HOT-TM data was downloaded between the 1st and 3rd December 2023, and OSM data from Bunting Labs API on the 4th December 2023. The projects downloaded were those created from the 1st December 2021 (two years before 1st December 2023). An R Notebook is provided for the analysis of the fetched data. The full code is available at https://github.com/Robot8A/beh-analysis-hotosm-tm.

Acknowledgements. This paper is part of the "ODECO: Towards a sustainable Open Data ECOsystem" project[8].

This project has received funding from the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No 955569.

The opinions expressed in this paper reflect only the authors' view and in no way reflect the European Commission's opinions. The European Commission is not responsible for any use that may be made of the information it contains.

References

Arapostathis, S. G.: Fundamentals of volunteered geographic information in disaster management related to floods, Flood Impact Mitigation and Resilience Enhancement, 2020.

Bégin, D., Devillers, R., and Roche, S.: The life cycle of contributors in collaborative online communities-the case of OpenStreetMap, International Journal of Geographical Information Science, 32, 1611–1630, 2018.

El Hatimi, B., Oulidi, H. J., and Fadil, A.: Quality assessment in volunteered geographic information for risk management applications, in: 2020 IEEE International conference of Moroccan Geomatics (Morgeo), pp. 1–4, IEEE, 2020.

[3] https://maproulette.org/
[4] https://mapswipe.org/
[5] https://streetcomplete.app/
[6] https://fair-dev.hotosm.org/
[7] https://github.com/hotosm/fmtm/
[8] https://odeco-research.eu/

Goodchild, M. F.: Citizens as sensors: the world of volunteered geography, GeoJournal, 69, 211–221, https://doi.org/10.1007/s10708-007-9111-y, 2007.

Hardy, A.: Tracking via Volunteered Geographic Information, 2020.

Haworth, B. T., Bruce, E., Whittaker, J., and Read, R.: The good, the bad, and the uncertain: Contributions of volunteered geographic information to community disaster resilience, Frontiers in Earth Science, 6, 183, 2018.

Hosseinali, F. and Farhadpour, S.: Developing a spatial solution for earthquake crisis management using volunteered geographic information and genetic algorithm: A case study of an earthquake, Tehran, Iran, Earth Observation and Geomatics Engineering, 4, 109–118, 2020.

Ke, K.-Y., Tsai, M.-H., Tan, Y.-C., Lai, M.-C., Jhou, G.-H., and Chang, T.-L.: Smart Management of Emergency Shelter and Resilient Community, Tech. rep., Copernicus Meetings, 2023.

Neis, P., Singler, P., and Zipf, A.: Collaborative mapping and Emergency Routing for Disaster Logistics - Case studies from the Haiti earthquake and the UN portal for Afrika, 2010.

Rapid Editor: Rapid Editor, https://rapideditor.org/, 2023.

Rey, M. C.: Spatial information applied in disaster risk management, Geographic Magazine of Chile Terra Australis, 58, 2022.

Safariallahkheili, Q. and Malek, M. R.: A method for assessing the credibility of volunteered geographic information in case of flood crisis, Procedia Computer Science, 207, 1611–1622, 2022.

Saganeiti, L., Amato, F., Murgante, B., and Nolè, G.: VGI and crisis mapping in an emergency situation. Comparison of four case studies: Haiti, Kibera, Kathmandu, Centre Italy, GEOmedia, 21, https://mediageo.it/ojs/index.php/GEOmedia/article/view/1445, number: 3, 2017.

Tzavella, K., Skopeliti, A., and Fekete, A.: Volunteered geographic information use in crisis, emergency and disaster management: a scoping review and a web atlas, Geo-spatial Information Science, pp. 1–32, 2022.

Urrea, G. and Yoo, E.: The role of volunteer experience on performance on online volunteering platforms, Production and Operations Management, 32, 416–433, 2023.

Xin, Y.: Anomaly detection for volunteered geographic information: a case study of Safecast data, International Journal of Geographical Information Science, 36, 1423–1442, 2022.

Zhao, Q., Chen, Z., Liu, C., and Luo, N.: Extracting and classifying typhoon disaster information based on volunteered geographic information from Chinese Sina microblog, Concurrency and Computation: Practice and Experience, 31, e4910, 2019.

Zhao, T., Liu, W., Di, X., Li, R., Wang, X., Zhai, X., Zhu, X., Zhang, Y., and Peng, Y.: Research on the way of Sharing Geographic Information Data in Disaster Management, The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences, 48, 103–109, 2022.