

# Mass Participation During Emergency Response: Event-centric Crowdsourcing in Humanitarian Mapping

**Martin Dittus**  
ICRI Cities, UCL  
London, UK  
martin@dekstop.de

**Giovanni Quattrone**  
Dept. of Computer Science,  
UCL, London, UK  
g.quattrone@cs.ucl.ac.uk

**Licia Capra**  
Dept. of Computer Science  
UCL, London, UK  
l.capra@ucl.ac.uk

## ABSTRACT

Crowdsourcing platforms have become important information providers after disaster events. While they can build on some prior experiences, it is not yet well understood how contributor capacity for such activities is constituted. To what extent are initiatives building a dormant task force that springs to action when it is needed? Alternatively, do they mainly rely on the recruitment of new contributors during disaster events, possibly at the expense of contribution quality? We seek to develop a better understanding of these relationships, using the example of the Humanitarian OpenStreetMap Team. In a large-scale quantitative study, we assess the outcomes of 26 campaigns with almost 20,000 participants. We find that event-centric campaigns can be significant recruiting and reactivation events, however that this is not guaranteed. Our analytical methods provide a means of interpreting key differences in outcomes. We close with recommendations relating to the promotion and coordination of event-centric campaigns in HOT and related platforms.

## Author Keywords

Crowdsourcing; Event-centric Crowdsourcing; Emergency Response; Engagement; Contribution Quality; Retention

## ACM Classification Keywords

K.4.3. Organizational Impacts: Computer-supported collaborative work

## INTRODUCTION

When a large earthquake hit Nepal in early 2015, thousands of people found themselves without shelter and food. Humanitarian aid teams sought to provide urgent support, however they required updated maps in order to coordinate their efforts. In response, online volunteers of the Humanitarian OpenStreetMap Team (HOT) sprung to action. Based on satellite imagery they created a highly detailed map, highlighting the location of residential areas, and producing early damage assessments. In the words of HOT organiser Dale Kunce, “There are literally thousands of new mappers who I’ve never met before, and they’re by and large doing a great job” [21].

This was not the first time HOT had played an important role in relief efforts. Similar stories have been reported in relation to earlier events, such as typhoon Haiyan in the Philippines [20], and the Ebola epidemic in West Africa [22]. However the community is not only activated when a disaster has already taken place. Characteristic for HOT is the coexistence of two broad types of campaigns. *Event-centric campaigns* are initiatives that seek completion within days or even hours, typically in the context of urgent emergency response. This is a synchronised kind of activity, in that contributors participate in the specific moment when a particular urgent need arises. *Mission-centric campaigns* on the other hand are mapping initiatives without a particular deadline. They may seek to proactively map certain unmapped areas, or update existing maps. Some of these may be long-running initiatives, covering vast geographic areas. Mission-centric activity can be characterised as asynchronous, in that individuals largely contribute at their own leisure.

Does the coexistence of both event-centric and mission-centric activity have implications for contributor engagement? It is currently not clear how contributor capacity for these different activities is constituted. On one hand, event-centric campaigns may be important growth events: media coverage during crisis events may attract many new contributors. On the other hand, event-centric campaigns might also invite a kind of contributor engagement that is characterised by dormancy-reactivation cycles: an experienced yet passive membership that only reawakens when it is needed. However there are no existing studies of such a general effect.

An important organiser concern is the potential tradeoff between community growth and data quality. A quick response often matters when a disaster strikes. According to FEMA surveys, the value of updated ground surveys to coordinating aid teams decreases with every additional day [19]. Modern coordination technologies make it possible to satisfy such timelines with the help of a large global volunteer force. However for the data to be useful it also needs to be accurate, which makes it a concern who participates in these mapping efforts.

## Proposed Contributions

We present the first large-scale study of a crowdsourcing system to investigate the relationship between event-centric coordination practices, contributor engagement, and campaign outcomes. The study takes the form of a large-scale quantitative study of past campaign outcomes. We compare HOT community activity during emergency-centric and mission-centric campaigns in two complementary respects.

*Reactivation.* Is HOT building a volunteer task force that springs to action in response to specific external events? To what extent are dormant contributors reactivated by urgent emergency responses? What is the impact of such dormant capacity on campaign outcomes?

*Recruiting.* Are disaster events important growth moments because they attract more newcomers, compared to other campaigns? What is the impact of these newcomers on campaign outcomes, compared to more experienced volunteers? Do they join future campaigns?

To address these concerns, we observe participation outcomes by almost 20,000 HOT volunteers across 26 HOT campaigns, including both event-centric and mission-centric campaigns. We propose a measure of campaign burstiness to quantitatively distinguish these two kinds of campaigns. We further develop two intrinsic measures of contribution quality which are used to assess campaign outcomes: the share of untagged new objects, and the share of new objects that are subsequently deleted by other contributors.

On the following pages we first present an introduction to HOT and its practices, and provide an overview of related work on event-centric crowdsourcing and HOT contributor engagement. We then outline our research questions, and describe our methodology. Finally we address our research questions with a set of analyses of campaign participation and outcomes, and close with a discussion of our findings, and a brief outline of future work.

## HOT REMOTE MAPPING

### Emergency Response and Proactive Mapping

HOT emerged out of informal disaster mapping activities on the online mapping platform OpenStreetMap (OSM). The outcomes of their work support the work of humanitarian aid organisations. Field experts in these organisations require maps to coordinate their relief work, yet updated maps may be hard to come by, and many regions of the inhabited world are still unmapped. Early experiences demonstrated the utility and potential for volunteer-created mapping data, and resulted in the creation of HOT as a coordinating body of such volunteer activities [48]. HOT has since provided emergency response after typhoon Haiyan in the Philippines, the Nepal earthquake, the tropical cyclone Pam, recent earthquakes in Afghanistan and Ecuador, and many other disaster events.

Today, HOT engages in a wide range of activities in addition to its initial focus on such event-centric campaigns [52]. A growing number of initiatives seek to produce maps before they are needed, for example in preparation for humanitarian field work. These are typically more sustained mission-centric campaigns, covering vast regions over multiple weeks or months per campaign. These maps are often the first ever of their kind. Some of these mission-centric campaigns are focused on specific geographic regions such as the Congo, Central Africa, Lesotho, or South Sudan [32]. In other cases they are larger umbrella initiatives orchestrated by organisations who seek to support the creation of new maps, including the Peace Corps, MapGive, and Missing Maps [43, 22]. In the case of the Ebola response, an initially localised emergency response campaign

turned into a multi-regional and multi-month effort to produce maps for a growing number of affected regions [5].

Conceptually, we can distinguish three entry points by which participants discover projects, be they first-time contributors or experienced mappers. In some cases, projects are highlighted in public discussion in media, social media, and other forms of external promotion [20, 21, 22, 5]. High-profile disaster events may benefit from this in particular. In other cases, organisers may recruit participants through non-public channels, for example as part of partnerships with other organisations, or through the organisation of mapping events [47]. Such recruiting channels can play an important role for campaigns focused on community-building, where organisers have a direct relationship with prospective contributors [8, 47, 13]. Finally, contributors may start by browsing the published list of available projects, with or without a particular intention.

### The Remote Mapping Process

Remote mapping is HOT's main organisational output, it entails the tracing of satellite imagery by thousands of volunteers, using a shared set of tools and processes. Fundamental to this work is the HOT *tasking manager*, an online platform that is used to coordinate the volunteer community of remote mappers. It was introduced by HOT as a means of reducing edit conflicts, and to streamline the contribution process [39].

On the tasking manager, work is organised into projects, each aspiring to map specific features in a particular region of the world. Multiple projects are often part of larger mapping campaigns, for example covering different areas within a larger region. To facilitate volunteer coordination, project areas are geographically divided into smaller map segments, so-called tasks. Contributors register for a task when they begin their work, and mark it as 'complete' when their work is done.

A typical HOT project may ask volunteers to trace a simple basemap of roads and buildings. These map objects are drawn as line geometries, so-called ways, and annotated with basic metadata to distinguish them. The actual mapping work involves the use of OSM tools and platforms.

To improve map quality, HOT has an internal peer review process in the form of a *validation* stage. During validation, an experienced mapper reviews contributions to a task, makes changes as necessary, and either declares them 'valid' or 'invalid', for example because the work was incomplete. As part of this process, validators may also adjust geometries, refine annotations, and delete map objects which were of an insufficient quality. Invalidated work can then be picked up by other mappers, who may further refine any existing map data. The validation process is currently not standardised, however there are some general guidelines, and much community discourse about potential approaches [34, 33, 56, 37, 44].

Figure 1 shows an exemplary outcome of this work, a mapping campaign in rural Nepal after the 2015 earthquake. As the picture shows, in this case there was already some prior map data before the campaign started, however the work by HOT volunteers yielded a much more detailed picture.

<sup>1</sup> Screenshot taken from <http://osm-analytics.org>




Figure 1. Mapped buildings before (left) and after (right) a HOT campaign in rural Nepal.<sup>1</sup>

## RELATED WORK

### Event-centric Crowdsourcing

HOT emergency response campaigns can be considered instances of *event-centric crowdsourcing* [10, 11], or *synchronous work* in the Johansen time-space matrix [16]. They are efforts to recruit distributed participants for a particular event which has a start and an ending. On the other hand, mission-centric campaigns can be considered *asynchronous work*, in that activities are not typically focused on a timely response to an external event, and instead participants can contribute at their own leisure.

To our knowledge, there are no published studies on the impact of event-centric coordination on the *recruitment and reactivation* of HOT contributors over time. However work in related domains can introduce some initial expectations. In social media research, studies have documented the willingness of outsiders to participate remotely during crises, for example to help in information propagation [49, 54, 50]. It was further found that participation in such events may lead to an interest in becoming more deeply involved in future initiatives [51, 6]. In the context of Wikipedia, it was found that breaking news can lead to intense collective editing activity which surpasses that of most other Wikipedia articles [17]. Such media events can rally a diverse set of contributors: some may have contributed during a previous event, while others may be first-time contributors to Wikipedia who make some minor changes and never return. In comparison, other articles tend to have a more stable contributor network [17].

There is some early knowledge about how these different coordination practices may affect the HOT *contribution flow* over time. A recent MSc thesis compares OSM and Wikipedia edit patterns during major disaster events [12]. The study finds two patterns of behaviour across both platforms: contribution activity after earthquakes and hurricanes is characterised by large initial spikes followed by a long decay period, while a multi-month response to the West Africa Ebola epidemic was characterised by more sustained activity levels. It was further found that OSM activity was much reduced in instances where it emerged informally, rather than being coordinated by

HOT. In other words, there is evidence that HOT coordination can amplify and even foster increased community activity that otherwise would not have taken place. According to the study, organised mapping campaigns have a large impact on contribution patterns [12].

### HOT Engagement and Outcomes

To date, few studies have tried to assess how HOT contributor engagement is constituted over time. A growing body of published work seeks to assess the outcomes of HOT campaigns, with a focus on larger and more well-known event-centric campaigns: the Haiti earthquake [48, 39], typhoon Haiyan/Yolanda [39, 55, 8], the Nepal earthquake [9, 2, 41]. However such studies are typically limited to evaluations of these individual campaigns, and focused on questions of process and data quality.

An evaluation of the Nepal campaign finds that more contributions were made by experienced mappers, however that first-time mappers provided small but important contributions, such as the creation of notes for missing information. The authors observe that further research is needed to verify if and how prior contributor experience affects data quality [41].

A further study compares HOT contributor engagement across three large campaigns, typhoon Haiyan, the Ebola response, and Missing Maps. The authors find that newcomer retention is significantly lower during the former event-centric campaign than during the latter two mission-based campaigns. They suggest that this may be attributable to self-selection effects related to different recruiting practices, as well as differences in community-building practices between the campaigns [8].

In a detailed assessment of OSM contributions after typhoon Haiyan, it is further suggested that media coverage may influence mapping outcomes, however the specific relationship is not yet well understood. The authors of the study observe that a region which was more frequently covered by news media had been mapped differently than other nearby regions: the contributions were of a markedly lower quality [55].

### Research Contribution

We present the first study of a crowdsourcing system to investigate the relationship between event-centric coordination practices, contributor engagement, and campaign outcomes. To our knowledge this is also the first large-scale study of HOT contributor participation across a large number of different campaigns, and the first study of HOT to establish a link between prior contributor experience and contribution quality. We further introduce two intrinsic measures of map contribution quality which are used to assess the outcomes of the observed campaigns.

Our findings have direct operational relevance for HOT, and we outline opportunities to improve coordination mechanisms during large emergency responses. Furthermore, the findings have theoretical implications for other event-centric platforms where participants contribute towards a shared outcome. We discuss these in a concluding section.

## RESEARCH QUESTIONS

### RQ1: Contributor Recruitment and Reactivation

- Do event-centric campaigns attract a larger share of newcomers than mission-centric campaigns?
- Do they reactivate a larger share dormant contributors?
- What is the aggregate contribution impact by these groups, and how does this differ across campaign types?
- How long do newcomers remain active after they joined a particular kind of campaign?

We expect that event-centric campaigns are both recruiting and reactivation moments, however the relative proportions are not clear. We further expect that experienced mappers provide more contributions in aggregate than newcomers, regardless of campaign type. Finally, based on the reviewed literature we expect that newcomers who join during an event-centric campaign are less likely to be retained than those who join during mission-centric campaigns.

### RQ2: Contributor Performance

- Do the observed recruiting and reactivation patterns have an impact on the nature of incoming contributions?
- Do different contributor segments produce work at a different rate, or of a different quality? For example, how do the contributions by new recruits compare to those of more experienced contributors?
- How does this affect the outcomes of particular campaigns?

We expect that some contributor groups may be motivated to spend more time on event-centric campaigns. We also expect that on average, newcomers spend less time contributing, produce work at a lower pace, and of a lesser quality than experienced contributors.

## METHODOLOGY

### Data

All our analyses are based on two data sets:

1. Project information published on the HOT Tasking Manager.<sup>2</sup> This data was scraped for every project.
2. The OSM edit history of all map contributions, recording the creation and modification of map objects over time. This dataset is freely available for download.<sup>3</sup>


Using contributor lists from the Tasking Manager as a starting point, we extracted the OSM map contributions by all known HOT participants and cross-referenced them with HOT projects based on username, date, and location. For the purpose of this study, any creation or modification of a map object is considered an edit.

### Campaign Selection

We identified a list of HOT campaigns based on a review of tasking manager projects. We thematically grouped all HOT projects relating to a particular larger concern. We restricted our analysis to tasking manager projects with at least 50 participants. 45% of the projects could not be classified in this manner: they had less than 50 participants each, and were not

<sup>2</sup><http://tasks.hotosm.org>

<sup>3</sup><http://planet.osm.org/planet/full-history/>



**Figure 2. Campaign timeline.** Each line visualises the activity period per campaign, indicating how much time passed until 50% of the overall contributions were made.

linked to any of the larger campaigns. These projects represent 26% of all HOT contributions.

This process yielded 29 campaigns. The start of the first campaign in early 2012 marks the beginning of the observation period for our study. At the time of the study, the last available date in the contribution history was 4th July 2016. Our evaluation includes an analysis of newcomer retention, for which we chose an observation period of 90 days. This threshold was chosen to be significantly longer than 30 days: some regional HOT communities organise monthly mapping events [47], and attendees of such events who do not map at home should still be regarded active contributors.

This determines our study period:

- First observation date: February 2012
- Last inclusion date: 18th March 2016
- Last observation date: 16th June 2016

We only considered campaigns where at least 75% of contributions were made before the cutoff date of 18th March 2016. After this process, 26 campaigns remained part of our study. We determined all campaign participants whose first contributions were before the last inclusion date, amounting to a total of approximately 19,000 participants. They represent a majority share of 87% of the almost 22,000 recorded HOT participants before that date. Their contribution history is the basis for our study, it represents 76% of all HOT edits since early 2012, approximately 100 million edits.

Figure 2 shows a timeline of all campaigns we identified during this process, classified by type. The classification process is explained in the following section.

## Campaign Classification

On the tasking manager, projects may refer to specific events or long-term missions in their documentation, however event-centric and mission-centric are otherwise not explicitly labelled as such. For our study, we instead manually labelled all campaigns based on their primary cause, as stated in project titles and documentation. 11 campaigns relating to an external event such as a natural disaster (a flood, earthquake, landslide, typhoon, hurricane, or cyclone) were labelled as event-centric. The remaining 15 campaigns were labelled as mission-centric.

In addition to these manual labels, we sought to develop a quantitative classifier which discriminates between campaign types based on participation activity over time. Such a classifier could later be used to replicate our findings in larger studies where manual labelling is infeasible, and in studies of different systems. In particular, the classifier should be suitable for studies of systems which may have event-centric participation characteristics, but that do not make the same explicit distinction between urgent event-centric and more long-running mission-based campaigns. For such cases, we sought to derive the distinction from collective participation activity over time. Is a campaign an urgent response to an external event, or does it entail a sustained period of activity?

During a review of campaign contribution timelines we encountered two distinct temporal patterns. Event-centric campaigns were characterised by a single, short, and large burst of initial activity, followed by a longer decay period of minimal activity. In some cases this decay period lasted multiple weeks or even months. More long-term mission-centric campaigns were characterised by a sustained period of activity, typically over multiple weeks or months. Their contribution timelines may include multiple intermediate bursts of activity. This matches an observation by a recent study of HOT contribution flows across different campaigns [12].

After some iterations we chose the *campaign midpoint* as a discriminator between campaign types. The midpoint is the time that has passed until 50% of all contributions to a campaign have been made. We use this measure as an indicator of *campaign burstiness*. The median midpoint across all campaigns is 60 days: half the campaigns finished within or before this time. The remaining campaigns lasted much longer, some spanning many months of activity.

This median midpoint can be used as a threshold for binary classification. Almost all campaigns below this threshold were event-centric campaigns, with only two false positives. All campaigns above the threshold were mission-centric, with no false negatives. Overall this classifier has a false discovery rate of 7.7% of all campaigns (2 out of 26). This misclassification represents an effective error rate of 1.4% of the total edit volume, and 1.5% of all participants. We consider this an acceptable classification error for the purpose of this study.

The two false positives were both comparatively small-scale projects. A campaign to map the Tanzania road network [36] started with a single large automated import, which was succeeded by a longer period of mission-centric mapping. It represents less than 1% of total edits in the study, involving

less than 1% of all study participants. A mission-centric effort to digitise buildings in informal settlements in Gaza [26] was a relatively popular campaign with limited geographic scope, characterised by a bursty contribution flow and a short lifespan. It was similarly small in scale, at approximately 1% of total edits, involving 1% of study participants.

Since the classification error is so low, we decided to use the classifier in parts of our quantitative evaluation. Specifically, we used it for correlation analyses between campaign type (as measured in burstiness) and campaign outcomes. In our discussion of such analyses, we will refer to bursty campaigns with the term ‘event-centric campaigns’, unless the meaning of these terms diverges in a way that would affect the interpretation of our findings. Non-bursty campaigns will be called ‘mission-centric campaigns’.

## Contribution Profiles

For a comparative study of campaign outcomes we sought to determine who participated in particular campaigns, how much they contributed, and how well they contributed. We will first outline in a general manner how these aspects were derived, and then explain key aspects in more detail.

We first computed a session history for every study participant, using a method introduced for Wikipedia contributor analysis [14]. This yielded the number of edits per session, and an estimate of the time spent on these contributions, also called *labour hours*. The division of these yields a contributor’s *edit pace*, the rate at which they contributed during the session. We further computed each participant’s *campaign history*, starting from their first HOT campaign, and recording any subsequent campaign they joined.


Using these measures as a basis, we computed *contribution profiles* for every instance where a study participant contributed to a new campaign. These contribution profiles were later used to evaluate the outcomes of the different campaigns. They include measures across different areas of concern:

- The participant’s prior activity at the time of initial campaign participation. Are they a first-time mapper, were they already recently active in other campaigns, or have they been inactive for a longer period? We call this the *contributor segment*.
- The participant’s *contribution activity*: number of edits, labour hours, and edit pace.
- The participant’s *contribution quality*: the share of untagged new objects, and the share of objects which are later deleted by other users.
- For first-time contributors: their *retention* after the initial campaign contribution.

Retention was measured with a survival analysis over the duration of the observation period (90 days after the initial contribution). Contributor segment classification and measures of contribution quality are developed in more detail in the following sections.

## Contributor Segments

When can a contributor be considered dormant? We analysed the frequency and duration of contributor inactivity periods to



**Figure 3. Inactivity period between campaign engagements, for all contributors who participated in more than one campaign.**

inform our choice of an inactivity threshold: the time between the last edit of the previous campaign, and first edit of the following. Inactivity periods follow a long-tail distribution, as is shown in Figure 3. 3,500 contributors joined a second campaign (18%). Of these, the median time between engagements is 13 days. 40% engagements involved an intermittent dormancy period of 29 days or more, 30% of 62 days or more.

We sought an inactivity period long enough so that it can arguably be regarded as dormancy period, rather than merely an interruption. In particular, it should be significantly longer than 30 days: some regional HOT communities organise monthly mapping events [47], and attendees of such events who do not map at home should still be regarded active contributors. On the other hand our threshold should still be short enough to capture a significant number of occurrences.

After some early trials we chose 60 days of inactivity as a threshold for dormancy. This is a fairly high threshold above the 70th percentile, yet it still captures a large number of samples. More than 1,600 study participants (9%) at one point in their contribution history became inactive for at least 60 days, but then returned for a future campaign. They represent approximately half of the study participants who had contributed to more than one campaign.

Using this inactivity threshold, we classified all study participants at the time they joined a new campaign into one of three segments:

- If this is their first HOT contribution: *Newcomer*.
- If they contributed in the last 60 days: *Already active*.
- If they previously contributed, but were inactive in the last 60 days: *Previously dormant*.

### Assessing Contribution Quality

For our evaluation we further sought to compare the contribution quality of different contributors across different mapping activities. To this purpose we required simple indicators of a mapper’s contribution quality. These measures needed to be suitable for the HOT context, and manageable within the scale and scope of this study, encompassing many thousands of contributors, and millions of edits.

At this scale, ground truthing of contributions is infeasible. Instead we sought to develop *intrinsic* measures of contribution quality. In the OSM literature, there is a range of widely used intrinsic measures of map quality [4], however

many are not suitable in this context. As an example of this, the frequently-used indicator measure of Linus’ Law [15] assumes that regions are refined by multiple mappers over time, whereas the HOT contribution process assigns only a single mapper to each region.

Outcomes from the HOT validation process could in principle serve as an indicator of contribution quality, but they are unfortunately not made public. Map edits by validators are public, but not clearly attributable to validation, and indistinguishable from contributions by other mappers.

After a review of these options we decided to incorporate two complementary aspects. We computed the share of *untagged new objects* as an instance of easily identified mistakes during the mapping process, and the share of *objects that are eventually deleted* as example of modifications made to new objects some time after they have been added. Both measures can easily be derived from the edit history. They are discussed in more detail in the following sections.

### Untagged Map Objects

Prior research of the OSM contribution process suggests that beginners do not always annotate their map objects, which renders their contribution unusable [3, 23]. However there is some dispute in the literature about whether this is a regular occurrence, or even unique to newcomers: a study that assessed HOT contributions after Nepal finds that mappers with less experience do not necessarily produce less well-annotated map objects [2].

We included the share of untagged new objects as an evaluation criterion in our study to determine whether this reportedly low occurrence would still be sufficient for our large-scale study of contribution behaviour. However based on the prior evidence we expected a low rate of untagged new objects.

To compute the measure, we determined the number of map objects created by study participants which had no annotations at the end of each contribution session. We excluded annotations from this analysis which are automatically added by editing tools (*created\_by* and *source*). In total, only a low 0.6% of map objects created during the observation period were left untagged by the creator. 4,700 study participants (26%) had created at least one such untagged map object.

### Map Object Persistence

Studies on Wikipedia contribution quality introduced the notion of *contribution persistence*, also called transience. The concept describes the extent to which a contribution survives subsequent review by other contributors [1, 40, 42, 57]. It maps well to the HOT contribution process: validators may delete contributions they consider of a low quality.

A review of the HOT edit history showed that a significant number of new map objects were deleted by their creator, often within the same edit session. This indicates that deletions can also be a normal part of the contribution process, for example to fix mistakes as they occur. For our analysis we thus ignore instances where objects are deleted by the creator, and only observe deletions by different contributors. The median delay between object creation and deletion was 29 days, the 75th

% contributors	Newcomer	Prev. dormant	Already active
Event	38.4% 	16.1% 	45.2% 
Mission	61.2% 	7.7% 	33.3% 
Overall	50.2% 	10.1% 	38.3% 

Table 1. Median percentage of participants, by contributor segment and campaign type.

Type	Campaign	% newcomers	# newcomers
Event	Nepal earthquake	84.7% 	5,072 
Mission	Missing Maps	84.4% 	6,280 
Mission	Tanzania roads	80.2% 	134 
Event	Typhoon Haiyan	76.4% 	482 
Mission	Ebola response	76.2% 	2,184 
Mission	Indonesia	75.5% 	240 
Mission	MapLesotho	68.4% 	360 
Mission	MapUganda	66.5% 	177 
Mission	Peace Corps	65.5% 	898 
Mission	MapGive	61.2% 	112 

Table 2. The 10 campaigns with the largest newcomer share.

percentile 213 days. Based on this we chose an evaluation threshold of 90 days, matching our observation window.

During the study period, 3.1% of newly created map objects were deleted by a different contributor within 90 days. This is still a low proportion, but higher than the share of untagged new objects. Approximately 56% of study participants have created at least one object that was deleted by someone else within the next 90 days.

A correlation analysis across contributor records showed that the two contribution quality measures are not mutually correlated. Furthermore, neither of them is correlated with campaign burstiness or campaign start date. On the other hand, the account age of contributors at the time of object creation is negatively correlated with the rate of deleted objects (Spearman coefficient  $\rho_S = -0.16$ ,  $p < 0.001$ ), and weakly correlated with the rate of untagged new objects ( $\rho_S = -0.04$ ,  $p < 0.0001$ ), suggesting a relationship between contributor experience and the two contribution quality measures.

These factors indicate that the measures are useful for an evaluation of contributor engagement: they capture different user behaviours that are not obviously interrelated, and not obviously biased by campaign-specific processes. Instead they can serve as indicators of the contribution quality of individual contributors at particular moments in their contribution history.

## FINDINGS

### RQ1: Contributor Recruitment and Reactivation

Table 1 shows the distribution of participants across both event-centric and mission-centric campaigns. According to these numbers, emergency response does benefit from a clear reactivation effect. For event-centric campaigns, the share of previously dormant contributors doubles compared to mission-centric campaigns. A correlation analysis between campaign




Figure 4. Campaign participation per contributor segment. Every vertical bar represents a campaign.

burstiness and reactivation rate confirms this effect: mission-centric campaigns involve a lower share of reactivated dormant contributors ( $\rho_S = -0.74$ ,  $p < 0.0001$ ). However in comparison to the other contributor segments, overall participation by this group was relatively rare.

The inverse applies to newcomers: the share of first-time contributors during event-centric campaigns is markedly smaller than during mission-centric campaigns. Mission-centric campaigns tend to attract more newcomers ( $\rho_S = 0.40$ ,  $p < 0.05$ ). However manual inspection showed two exceptional outliers, both large event-centric campaigns, each with approximately 80% newcomers among their participants. These were emergency response campaigns to typhoon Haiyan in the Philippines in 2013, and the Nepal earthquake in 2015. In comparison, all other event-centric campaigns had only between 10 and 50% newcomer participants.

Table 2 lists the campaigns with the largest share of newcomers. The top ranks include Haiyan and Nepal, but otherwise only mission-centric campaigns, including Missing Maps, Indonesia, Ebola, Maplesotho, PeaceCorps and others, all with a newcomer share between 60-85%. In absolute terms, mission-centric campaigns recruited twice as many mappers as event-centric campaigns (11,600 vs 6,800).

The bar charts in Figure 4 further illustrate this relationship between campaign type, user segment, and participation: the newcomer share varies widely across campaign types, while participation by dormant reactivated contributors is low across all campaigns.

Table 3 shows how the contributions of these segments vary by campaign type. Contributors who had already been active were the most prolific, accounting for more than half of the total edits. The picture is more varied for newcomers. While they are the second-largest group in event-centric campaigns, they only contribute a low proportional share of edits: on average,

% edits	Newcomer	Prev. dormant	Already active
Event	11.0% █	9.9% █	67.4% █
Mission	40.3% █	5.5% █	43.6% █
Overall	29.9% █	8.9% █	53.9% █

**Table 3. Median percentage of contributed edits, by contributor segment and campaign type.**

they represent 40% contributors yet only provide 11% of edits. Here again, Nepal and Haiyan are an exception, their 80% newcomers produced around 80% of the overall campaign work. Correlation analysis confirmed that dormant provide a higher share of contributions during event-centric campaigns than during mission-centric campaigns ( $\rho_S = -0.46$ ,  $p < 0.02$ ), but was not significant for the other groups.


Overall, 80% of newcomers never joined a second campaign, this did not vary by campaign type. For a retention analysis we computed a survival analysis with Kaplan-Meier model, observing newcomer activity over 90 days after initial contribution. This revealed that newcomers who joined during an event-centric campaign have approximately a third of the retention rate: after 45 days, only 3.0% were still active, whereas for longer campaigns 9.6% were still active. For Nepal and Haiyan in isolation, retention rates were even lower (2.3% after 45 days), whereas newcomer retention in other event-centric campaigns was closer to mission-centric campaigns (6.2% after 45 days). Log-rank tests confirmed that the survival rates were statistically different between these groups ( $p < 0.02$ ).

In other words, retention rates were between 5-10% across most contributor segments and campaign types, with the exception of newcomers who joined during the Haiyan and Nepal campaigns. For these campaigns, average newcomer retention was only at 2.3%.


The share of already active contributors was not correlated with campaign type (as measured in burstiness), suggesting their choice to participate may be less related to campaign-specific considerations. However a review of their temporal contribution behaviour showed instances where these contributors joined large events as they occurred, and then resumed mission-based campaign work after the events were over. Figure 5 shows such contributor flows for two period of intense emergency activity in April-May and October-November 2015. In both cases, a large number of contributors who had been participating in mission-centric campaigns in earlier months joined these event-centric campaigns. As the emergency work was completed, many who had participated in the events returned to mission-centric work.

### RQ2: Contributor Performance

Overall, people spent more time and contribute more work during mission-centric campaigns, with a median of 5 hours compared to 3.3 hours during event-centric campaigns, and 2,500 edits compared to 1,600 edits. The average contribution rate is fairly stable across campaign types, at 630 edits per



**Figure 5. Alluvial flow visualisation of temporal contributor flows between event-centric (top) and mission-centric campaigns (bottom). Vertical bars indicate the monthly contributor count, edges the number of contributors who move between campaign types month over month.**



**Figure 6. Contribution volume. The number of edits per contributor, by contributor segment and campaign type.**

second. However on average, newcomers tended to work more slowly and be less prolific than the other contributor segments.

Average contributor performance according to our evaluation measures is further broken down in Table 4. This allows a more detailed comparison of contributor profiles across campaign types. The table shows that newcomer performance was lowest during event-centric campaigns: on average, they contributed less edits and spent less time contributing than any other group. On the other hand, already active contributors tended to increase their effort slightly during event-centric campaigns. Both effects are illustrated by the corresponding distributions of average contribution volumes in Figure 6.

Deletions happened rarely, as can be seen in Table 4. The median percentage of non-persistent objects is 0.0% across mission-centric campaigns, and around 1% for event-centric campaigns. However the distribution of deletion rates in Figure 7 illustrates that these outcomes vary significantly by contributor segment: newcomers who joined during event-centric campaigns created a larger share of non-persistent objects than any other group (2.5%), including newcomers who joined during mission-centric campaigns (0.33%). A Mann-Whitney U

	Newcomer		Prev. dormant		Already active	
	Event	Mission	Event	Mission	Event	Mission
Edits	1,252	2,508	3,178	2,741	2,340	2,996
Labour hours	2.8	5.4	5.2	4.4	4.1	5.2
Edits / hour	603	594	634	634	633	633
% untagged objects	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%
% eventually deleted objects	2.5%	0.33%	0.36%	0.00%	0.30%	0.00%

Table 4. Contribution profiles. Median contribution activity and contribution quality, by contributor segment and campaign type.




Figure 7. Contribution quality. The percentage of new objects per contributor which are deleted within 90 days.

test confirmed that this difference is significant ( $p < 0.01$ ). Further inspection of the data showed that this effect can mainly be attributed to the outlier campaigns of cyclone Winston and Nepal, where object deletion rates for newcomers were around 4%. All other campaigns had deletion rates of 1% or less.

Untagged objects were created even less frequently, and statistical tests comparing average rates per contributor were inconclusive. However we found that on average across the campaigns, a marginally higher share of newcomers created untagged objects (24%) compared to dormant contributors (20%). This difference is statistically significant (Mann-Whitney U test with  $p < 0.05$ ).

## DISCUSSION

Our findings provide a nuanced picture of several interrelated phenomena. They suggest that synchronous and asynchronous modes of campaign coordination are in fact strongly associated with different modes of engagement. However they also reveal the presence of further effects which are not fully captured by the data. We will first summarise the key observed effects and outline how they relate, and then discuss their implications.

*Most recruiting happened outside of event-centric campaigns*, against our initial expectation. On average, only 40% of participants in disaster campaigns were first-time contributors, compared to 60% in more long-running mission-centric campaigns. In absolute terms, mission-centric campaigns recruited twice as many mappers as event-centric campaigns. The top recruiting campaigns included Missing Maps and the Ebola response, which are specifically focused on sustained outreach

and community-building [8]. These two mission-centric campaigns had a significant impact on overall HOT recruitment.

*There are indicators for a reactivation pattern.* The share of reactivated dormant contributors was twice as high during event-centric than mission-centric campaigns, and on average these contributors were more prolific than any other group. This suggests that there is interest among a subset of the community to specifically help out during urgent emergencies, which matches observations from comparable participatory emergency response systems [51, 6]. However reactivation only happened at a small scale, on average only 16% of contributors to disaster campaigns exhibited this behaviour. As a result, the overall impact of these reactivations was low. The comparatively small effect may indicate that experienced but dormant mappers are not easily reactivated for synchronous on-demand participation, or that there is a need for better reactivation mechanisms.

*Contribution quality was lowest for newcomers during event-centric campaigns.* For this group, a median of 2.5% map objects are deleted within 90 days, compared to 0.3% for more experienced mappers during disaster response, and 0.3% or less for any contributor group during mission-centric campaigns. The effect was particularly pronounced during emergency responses for the Nepal earthquake and cyclone Winston: during these campaigns, 4% of contributions by newcomers were eventually deleted. This suggests that timeliness may have taken preference over data quality, although it is unclear whether this was a deliberate organiser choice.

*Campaign burstiness is a useful activity-based classifier for event-centric campaigns.* We demonstrated that event-centric emergency response campaigns can be identified with a simple heuristic measure, the number of days that passed until 50% of the work of a campaign was completed. A binary classifier based on this measure had a false discovery rate of only 7.7%, with 2 wrongly classified campaigns out of 26.

*Map object persistence is a useful intrinsic measure of HOT contribution quality.* Object persistence as measure of contribution quality was first introduced in research of Wikipedia contribution quality, and maps well to the HOT contribution process. It may be unsurprising that inexperienced participants can produce work of a lower quality, yet to our knowledge this is the first study to observe such an effect for HOT activity at this scale, and across this range of activities. However because deletions were rare overall, the measure could only be used to characterise a small percentage of contributors. The merit

of measuring untagged new objects is less clear. Statistical tests involving the average rate of untagged new objects were generally inconclusive.

### Differences in Event-centric Campaign Outcomes

The Nepal and Haiyan campaigns were noteworthy exceptions by almost all measures. These event-centric campaigns were significant recruiting events, each with 80% of first-time participants. Nepal in particular managed to recruit several thousand first-time mappers. These had a significant effect on campaign outcomes, accounting for the vast majority of all work in both campaigns. On the other hand, they were much less likely to be retained: only 2.3% were still active after 45 days, compared to 6.2% among newcomers who joined during any other event-centric campaign, or 9.6% during any mission-centric campaign. To our knowledge, this is the first published study to observe such a difference in outcomes among event-centric HOT campaigns.

The campaign characteristics observed so far do not offer a good means to reason about the difference. However, it is noteworthy that the two campaigns had the largest numbers of participants among event-centric campaigns, which may be indicative of other unobserved effects. For example, it is feasible that there were differences in promotion and recruiting strategy across campaigns, differences in media coverage, or differences in public interest.

To our knowledge, there is no public record of HOT promotion strategies for these campaigns. However, community members have collected references to international media coverage for some campaigns, and published them on the respective coordination pages. The pages for both Nepal and Haiyan each list more than 40 references per campaign, including coverage of HOT activities by the New York Times, the BBC, The Guardian, The Atlantic, and others [30, 35]. Where coordination pages exist for the remaining campaigns, they include few to no such references [28, 29, 31, 25, 27]. While such lists are not a good reflection of actual media coverage, they indicate a community interest to observe such coverage for Nepal and Haiyan. They may suggest that for these two campaigns, media promotion was perceived as important.

As a measure of public interest, we can compare the search volume relating to each campaign's disaster event as captured by Google Trends. A timeline of worldwide search activity is shown in Figure 8, and associated with each campaign. By this measure, three disaster events were most prominent: the Nepal earthquake, typhoon Haiyan, and hurricane Patricia which hit Central America in October 2015. Search volumes for the remaining disaster events were significantly lower.

These observations can serve as initial indication that promotion practices and media coverage may play a role in explaining campaign outcomes, in particular when they are accompanied by increased public interest. However, further research is needed to better understand the relationships between these factors. We offer some suggestions in a Future Work section at the end. In the meantime we advise caution when applying lessons from Nepal and Haiyan to other disaster campaigns.




Figure 8. Google Trends: relative search volume per disaster event. HOT campaign periods for each event are highlighted.

Experiences gained in these campaigns may not always be transferrable to other emergency response settings.

### Implications

*HOT as an on-demand task force for emergency response.* During event-centric campaigns, a large part of participants was already active in recent weeks. Only a comparatively small share of participants were reactivated dormant contributors. This suggests that in the current form, experienced HOT contributors are less characterised as a dormant task force that can be activated on demand, and more a stream of continuous activity which can be redirected. Most of the observed on-demand capacity during emergency response was instead provided by inexperienced newcomers who tended to produce contributions of a lower quality. In response, in the following paragraphs we provide some recommendations to organisers for capacity management during emergency campaigns.

*Promotion of urgent disaster campaigns, and newcomer training.* During some event-centric campaigns, quality of contributions may be just as important as expediency. In order to improve contribution quality for such cases, we recommend to promote upcoming disaster campaigns to the existing community first, and only promote more widely if a quick response is needed and a decrease in overall quality is acceptable. Furthermore, newcomers should be encouraged to first make early experiences in mission-based campaigns where they can receive feedback on their work, and where outcomes are not as time sensitive.

*A need for an explicit reactivation mechanism?* HOT community coordination takes place on high-traffic mailing lists, online chat, and wiki pages [48, 39, 46, 2]. Monitoring these on an ongoing basis can require significant time investment. There currently is no means of being informed about urgent emergency responses except by being actively engaged in the community, which is reflected in the low reactivation rates of dormant contributors observed during event-centric campaigns. A dedicated opt-in alert mechanism may allow experienced

mappers to join urgent initiatives without having to actively follow daily community discourse, and increase the participation of experienced mappers during emergency response.

*A need for a personalised project listing?* The urgent HOT responses after the Nepal earthquake and typhoon Haiyan were largely completed by inexperienced contributors. While an influx of newcomers contributed to the completion of the responses, it also resulted in a marked decrease in contribution quality. This is a systemic limitation, and organisers have limited ability to affect such outcomes: in cases where high map quality is an important concern, they cannot currently divert newcomer flows to less critical projects, instead all participants are presented the same list of projects. In such cases, organisers may benefit from an ability to personalise the project listing, for example to emphasise different kinds of projects depending on a contributor's prior experience.

In summary, we found that event-centric campaigns can be significant recruiting and reactivation events, however that this is not guaranteed. There is evidence that dormant contributors can be reactivated by event-centric campaigns, yet this was rare, possibly due to the lack of a coordinated reactivation process. In two cases, on-demand capacity was largely provided by first-time contributors, which had a negative effect on contribution quality. The causes for this difference in recruiting outcome are not clear, however we provide suggestions for further research at the end.

*Theoretical implications.* Our findings have theoretical implications for crowdsourcing platforms with event-centric contribution characteristics, including emergency response initiatives that involve public participants [18, 7, 38, 24], platforms which solicit contributions in response to breaking news [17, 53], and other event-centric platforms where participants contribute towards a shared outcome. Organisers of such platforms should consider emergent effects relating to recruiting, reactivation, and how they may affect the rate and quality of incoming contributions. This is a particular concern during events which draw large public audiences, and where contribution quality matters. On one hand, platform designers should consider means to manage increased flows of first-time contributors, for example by emphasising less quality-sensitive contribution opportunities. On the other hand, overall contribution quality may benefit from the provision of opt-in notification mechanisms that reactivate more experienced contributors during key events. Designers of opinion-gathering platforms further need to consider how public interest around large events may affect the recruitment of participants, and how this may in turn influence the distribution of responses. This includes online petition platforms [58, 45].

## CONCLUSION

We presented the first large-scale study of a crowdsourcing system to investigate the relationship between event-centric coordination practices, contributor engagement, and campaign outcomes. Using the example of the HOT volunteer community, we compared outcomes of 26 different humanitarian mapping campaigns, involving almost 20,000 participants. We distinguished two types of campaigns: event-centric campaigns which seek to provide outcomes quickly, and mission-

based campaigns which proactively map larger regions in the absence of urgent causes. A particular focus was placed on engagement effects relating to the reactivation of inactive contributions, and recruiting of newcomers. We further assessed the quality of contributions by these volunteer segments, making this the first large-scale HOT study to establish a link between prior contributor experience and contribution quality.

Overall, HOT appears to be a highly engaged community of significant proportions: every campaign studied here benefited from low hundreds to thousands of participants. However in most cases, HOT emergency response does not give the impression of an on-demand task force. Instead, most of the contributions were made by participants who were already active at the time. The two high-profile exceptions to this were instances where synchronous activity originated from outside of the community. During the Nepal and Haiyan campaigns, the vast majority of contributions were provided by first-time mappers. While this shows that HOT can successfully channel popular concern around humanitarian crises, it also raises data quality concerns. It may be preferable for newcomers to make early experiences in a less time-sensitive setting. Furthermore, while most mappers are recruited during more long-term campaigns involving proactive mapping efforts, few experienced mappers are then reactivated during emergency response, suggesting a potential for more explicit synchronous coordination.

## Future Work

Our findings suggest there are unexplained effects relating to the nature and outcomes of HOT emergency response. We close with a potential direction for further research, and a theoretical implication for how we may conceptualise community size and growth.

*A media effect of increased newcomer recruitment and reduced contribution quality?* Contribution quality was worst for newcomers during Haiyan and Nepal, the disaster campaigns with the highest newcomer recruiting rates, and the lowest newcomer retention. The combination of these factors suggests the effect may relate to the specific contexts and coordination practices of these campaigns, including their recruiting practices. To our knowledge, both campaigns received more mainstream media coverage than other emergency response campaigns under study, as is illustrated by long lists of media references on their respective coordination pages [30, 35]. It is feasible that audiences attracted by such coverage were curious about the phenomenon and interested in supporting a worthwhile cause, rather than specifically motivated to produce high-quality maps. Such a media effect was first suggested in an evaluation study of Haiyan contributions [55]. However while the research to date may be indicative of such a media effect, further study is needed to establish whether this was in fact the case.

*What is the size of the contributor community?* More broadly, our findings suggest that campaign-based organising is associated with distinct modes of contributor engagement. This affects how we might conceptualise community size and growth. In particular, the size of the currently active contributor community may not necessarily be indicative of the number of

participants who will join a new campaign. There is a need to also consider passive capacity that may be reactivated, as well as recruiting effects during highly promoted campaigns, and campaigns that focus on outreach. For this reason we see an opportunity to develop predictive models of *potential capacity* that may help in campaign planning.

*Why do contributors choose certain projects?* We have not yet studied the extent to which project selection is dependent on contributor circumstances, as opposed to project availability. This is an important limitation of our methodology: the HOT contribution history allows to observe when and where participants contributed, but not how they were informed of particular campaigns. This means we cannot currently determine whether participants found out about a particular campaign through public media channels, community discussion forums, by browsing the tasking manager listing, or in other ways. It is feasible that some contributors prefer contributing to event-centric or mission-centric causes. Furthermore, it is feasible that contributor choices are influenced by their social context, such as the frequent participation in mapping events, or membership of a particular organisation.

*What factors influence contribution quality?* Measures of object persistence and share of untagged new objects yielded plausible and interpretable results in our study. The measures make it possible to produce large-scale studies of HOT contribution quality without a need for manual labelling. Such studies may seek to identify contributing factors of low-quality contributions that are under organiser control, for example the impact of satellite imagery quality, software tools, editing workflows, and others. Further work may seek to compare potential reasons for map object deletion, for example through a manual review of contribution records.

## ACKNOWLEDGMENTS

This research was funded by the EPSRC and the Intel Collaborative Research Institute: Cities. We are grateful to the HOT community for their invaluable support. We would further like to thank the reviewers for their feedback and suggestions.

## REFERENCES

1. B Thomas Adler and Luca De Alfaro. 2007. A content-driven reputation system for the Wikipedia. In *Proc. WWW '07*. ACM, 261–270.
2. Johannes Anhorn, Benjamin Herfort, and João Porto de Albuquerque. 2016. Crowdsourced validation and updating of dynamic features in OpenStreetMap an analysis of shelter mapping after the 2015 Nepal earthquake. In *Proc. ISCRAM '16*.
3. Jamal Jokar Arsanjani, Christopher Barron, Mohammed Bakillah, and Marco Helbich. 2013. Assessing the quality of OpenStreetMap contributors together with their contributions. In *Proc. AGILE '13*.
4. Christopher Barron, Pascal Neis, and Alexander Zipf. 2014. A comprehensive framework for intrinsic OpenStreetMap quality analysis. *Transactions in GIS* 18, 6 (2014), 877–895.
5. Liat Clark. 2014. The race to contain West Africa's Ebola outbreak - Wired UK. (2014).  
<http://www.wired.co.uk/news/archive/2014-04/11/ebola-open-street-map>
6. Camille Cobb, Ted McCarthy, Annuska Perkins, Ankitha Bharadwaj, Jared Comis, Brian Do, and Kate Starbird. 2014. Designing for the deluge: Understanding & supporting the distributed, collaborative work of crisis volunteers. In *Proc. CSCW '14*. ACM, 888–899.
7. Stephen Collins. 2011. Conflict and Disaster Management in a Hyper-Connected World: Cooperative, Collaborative, Real Time. *Journal of Urban and Regional Research* 25, 2 (2011), 227–52.
8. Martin Dittus, Giovanni Quattrone, and Licia Capra. 2016. Analysing volunteer engagement in humanitarian mapping: building contributor communities at large scale. In *Proc. CSCW '16*.
9. Melanie Eckle and João Porto de Albuquerque. 2015. Quality assessment of remote mapping in OpenStreetMap for disaster management purposes. In *Proc. ISCRAM '15*.
10. Thomas Erickson. 2010. Geocentric crowdsourcing and smarter cities: Enabling urban intelligence in cities and regions. In *1st Ubiquitous Crowdsourcing Workshop, proc. UbiComp '10*.
11. Thomas Erickson. 2011. Some thoughts on a framework for crowdsourcing. In *Workshop on Crowdsourcing and Human Computation, proc. SIGCHI '11*. 1–4.
12. Gary Esworthy. 2016. *Comparing OpenStreetMap and Wikipedia Edit Patterns During Major Events*. Master's thesis. George Mason University.
13. Jane Feinmann. 2014. How MSF is mapping the world's medical emergency zones. *BMJ* 349 (2014), g7540.
14. R Stuart Geiger and Aaron Halfaker. 2013. Using edit sessions to measure participation in Wikipedia. In *Proc. CSCW '13*.
15. Mordechai Haklay, Sofia Basiouka, Vyrion Antoniou, and Aamer Ather. 2010. How many volunteers does it take to map an area well? The validity of Linus' law to volunteered geographic information. *The Cartographic Journal* 47, 4 (2010), 315–322.
16. Robert Johansen. 1988. *Groupware: Computer support for business teams*. The Free Press.
17. Brian Keegan, Darren Gergle, and Noshir Contractor. 2013. Hot off the wiki: Structures and dynamics of Wikipedia's coverage of breaking news events. *American Behavioral Scientist* (2013), 0002764212469367.
18. Patrick Meier. 2011. New information technologies and their impact on the humanitarian sector. *International review of the Red Cross* 93, 884 (2011), 1239–1263.
19. Patrick Meier. 2016. The Value of Timely Information During Disasters (Measured in Hours). (2016).  
<https://irevolutions.org/2016/04/06/the-value-of-timely-information-during-disasters-measured-in-hours/>

20. Robinson Meyer. 2013. How online mapmakers are helping the Red Cross save lives in the Philippines. (2013). <http://www.theatlantic.com/technology/archive/2013/11/how-online-mapmakers-are-helping-the-red-cross-save-lives-in-the-philippines/281366/>
21. Robinson Meyer. 2016. The Internet Mapmakers Helping Nepal. (2016). <http://www.theatlantic.com/technology/archive/2015/05/the-mapmakers-helping-nepal/392228/>
22. Chris Michael. 2014. Missing Maps: nothing less than a human genome project for cities - The Guardian. (2014). <http://www.theguardian.com/cities/2014/oct/06/missing-maps-human-genome-project-unmapped-cities>
23. Pascal Neis and Alexander Zipf. 2012. Analyzing the contributor activity of a volunteered geographic information project—The case of OpenStreetMap. *ISPRS International Journal of Geo-Information* 1, 2 (2012), 146–165.
24. Ory Okolloh. 2009. Ushahidi, or 'testimony': Web 2.0 tools for crowdsourcing crisis information. *Participatory learning and action* 59, 1 (2009), 65–70.
25. OpenStreetMap. 2016a. 2010 Chile Earthquake - OpenStreetMap Wiki. (2016). [https://wiki.openstreetmap.org/wiki/2010\\_Chile\\_earthquake](https://wiki.openstreetmap.org/wiki/2010_Chile_earthquake)
26. OpenStreetMap. 2016b. 2014 Gaza Strip - OpenStreetMap Wiki. (2016). [https://wiki.openstreetmap.org/wiki/2014\\_Gaza\\_Strip](https://wiki.openstreetmap.org/wiki/2014_Gaza_Strip)
27. OpenStreetMap. 2016c. 2014 Southeast Europe floods - OpenStreetMap Wiki. (2016). [https://wiki.openstreetmap.org/wiki/2014\\_Southeast\\_Europe\\_floods](https://wiki.openstreetmap.org/wiki/2014_Southeast_Europe_floods)
28. OpenStreetMap. 2016d. 2015 Eastern Afghanistan Earthquake: Media coverage - OpenStreetMap Wiki. (2016). [http://wiki.openstreetmap.org/wiki/2015\\_Eastern\\_Afghanistan\\_Earthquake#Media\\_coverage](http://wiki.openstreetmap.org/wiki/2015_Eastern_Afghanistan_Earthquake#Media_coverage)
29. OpenStreetMap. 2016e. 2015 Hurricane Patricia: Media coverage - OpenStreetMap Wiki. (2016). [https://wiki.openstreetmap.org/wiki/2015\\_Hurricane\\_Patricia#Media\\_coverage](https://wiki.openstreetmap.org/wiki/2015_Hurricane_Patricia#Media_coverage)
30. OpenStreetMap. 2016f. 2015 Nepal earthquake: Write-ups and in the News - OpenStreetMap Wiki. (2016). [https://wiki.openstreetmap.org/wiki/2015\\_Nepal\\_earthquake#Write-ups\\_and\\_in\\_the\\_News](https://wiki.openstreetmap.org/wiki/2015_Nepal_earthquake#Write-ups_and_in_the_News)
31. OpenStreetMap. 2016g. 2015 Pam Cyclone Response: Media coverage - OpenStreetMap Wiki. (2016). [https://wiki.openstreetmap.org/wiki/2015\\_03\\_Vanuatu\\_Pam\\_Cyclone\\_Response#Media\\_coverage](https://wiki.openstreetmap.org/wiki/2015_03_Vanuatu_Pam_Cyclone_Response#Media_coverage)
32. OpenStreetMap. 2016h. Humanitarian OSM Team - OpenStreetMap Wiki. (2016). [http://wiki.openstreetmap.org/wiki/Humanitarian\\_OSM\\_Team](http://wiki.openstreetmap.org/wiki/Humanitarian_OSM_Team)
33. OpenStreetMap. 2016i. JOSM/Validator - OpenStreetMap Wiki. (2016). <https://wiki.openstreetmap.org/wiki/JOSM/Validator>
34. OpenStreetMap. 2016j. OSM Tasking Manager/Validating data - OpenStreetMap Wiki. (2016). [https://wiki.openstreetmap.org/wiki/OSM\\_Tasking\\_Manager/Validating\\_data](https://wiki.openstreetmap.org/wiki/OSM_Tasking_Manager/Validating_data)
35. OpenStreetMap. 2016k. Typhoon Haiyan: Media coverage - OpenStreetMap Wiki. (2016). [https://wiki.openstreetmap.org/wiki/Typhoon\\_Haiyan#Media\\_coverage](https://wiki.openstreetmap.org/wiki/Typhoon_Haiyan#Media_coverage)
36. OpenStreetMap. 2016l. WikiProject Tanzania - OpenStreetMap Wiki. (2016). [http://wiki.openstreetmap.org/wiki/WikiProject\\_Tanzania](http://wiki.openstreetmap.org/wiki/WikiProject_Tanzania)
37. Lisa Marie Owen. 2015. Global Validation Procedure. (2015). <https://www.openstreetmap.org/user/Lisa%20Marie%20Owen/diary/35522>
38. Leysia Palen, Kenneth M Anderson, Gloria Mark, James Martin, Douglas Sicker, Martha Palmer, and Dirk Grunwald. 2010. A vision for technology-mediated support for public participation & assistance in mass emergencies & disasters. In *Proc. ACM-BCS '10*. British Computer Society, 8.
39. Leysia Palen, Robert Soden, T Jennings Anderson, and Mario Barrenechea. 2015. Success & scale in a data-producing organization: the socio-technical evolution of OpenStreetMap in response to humanitarian events. In *Proc. SIGCHI '15*. 4113–4122.
40. Katherine Panciera, Aaron Halfaker, and Loren Terveen. 2009. Wikipedians are born, not made: a study of power editors on Wikipedia. In *Proc. GROUP '09*. ACM, 51–60.
41. Thiago Henrique Poiani, Roberto dos Santos Rocha, Livia Castro Degrossi, and João Porto de Albuquerque. 2016. Potential of Collaborative Mapping for Disaster Relief: A Case Study of OpenStreetMap in the Nepal Earthquake 2015. In *Proc. HICSS '16*.
42. Reid Priedhorsky, Jilin Chen, Shyong Tony K Lam, Katherine Panciera, Loren Terveen, and John Riedl. 2007. Creating, destroying, and restoring value in Wikipedia. In *Proc. GROUP '07*. ACM, 259–268.
43. Missing Maps Project. 2015. Missing Maps - putting the world's vulnerable people on the map. (2015). <http://www.missingmaps.org>
44. Suzan Reed. 2015. Missing Maps Training Video Suggestions. (2015). <https://lists.openstreetmap.org/pipermail/hot/2015-May/008865.html>
45. Jan-Hinrik Schmidt and Katharina Johnsen. 2014. On the use of the e-petition platform of the German Bundestag. (2014).
46. Abdul Rehman Shahid and Amany Elbanna. 2015. The Impact of Crowdsourcing on Organisational Practices: The Case of Crowdmapping. In *Proc. ECIS '15*.

47. Sam Smith. 2015. Prefer Typing to Running a Marathon? Take Part in a Mapathon. <http://blogs.redcross.org.uk/international/2015/10/prefer-typing-to-running-a-marathon-take-part-in-a-mapathon/>. (2015).
48. Robert Soden and Leysia Palen. 2014. From crowdsourced mapping to community mapping: the post-earthquake work of OpenStreetMap Haiti. (2014).
49. Kate Starbird and Leysia Palen. 2010. Pass it on?: Retweeting in mass emergency. In *Proc. ISCRAM '10*.
50. Kate Starbird and Leysia Palen. 2011. Voluntweeters: Self-organizing by digital volunteers in times of crisis. In *Proc. SIGCHI '11*. ACM, 1071–1080.
51. Kate Starbird and Leysia Palen. 2013. Working and sustaining the virtual Disaster Desk. In *Proc. CSCW '13*. ACM, 491–502.
52. Humanitarian OpenStreetMap Team. 2016. Featured HOT projects. (2016). <http://hotosm.org/projects>
53. Heli Väättäjä, Teija Vainio, Esa Sirkkunen, and Kari Salo. 2011. Crowdsourced news reporting: supporting news content creation with mobile phones. In *Proc. MobileHCI '11*. ACM, 435–444.
54. Sarah Vieweg, Amanda L Hughes, Kate Starbird, and Leysia Palen. 2010. Microblogging during two natural hazards events: what twitter may contribute to situational awareness. In *Proc. SIGCHI '10*. ACM, 1079–1088.
55. Clay Westrope, Robert Banick, and Mitch Levine. 2014. Groundtruthing OpenStreetMap building damage assessment. *Procedia engineering* 78 (2014), 29–39.
56. Johan Whelan. 2016. So you've thought about becoming a validator. (2016). <https://lists.openstreetmap.org/pipermail/hot/2016-May/011925.html>
57. Thomas Wöhner and Ralf Peters. 2009. Assessing the quality of Wikipedia articles with lifecycle based metrics. In *Proc. WikiSym '09*. ACM, 16.
58. Taha Yasseri, Scott A Hale, and Helen Margetts. 2013. Modeling the rise in internet-based petitions. *arXiv preprint arXiv:1308.0239* (2013).