LONDON'S GLOBAL UNIVERSITY
Department of Computer Science

Investigating Crowd-Worker Mapping Behaviour on OpenStreetMap During Times of Crisis

Candidate Number: MFLP6

Supervisor: Professor Licia Capra

Department of Computer Science
University College London

May 1st, 2025

This report is submitted as part requirement for the MEng Degree in Computer Science at UCL. It is substantially the result of my own work except where explicitly indicated in the text. The report may be freely copied and distributed provided the source is explicitly acknowledged.

Abstract

In the aftermath of natural disasters, accurate and up-to-date maps are critical for emergency responders to organise relief efforts while directing affected communities to essential services. OpenStreetMap (OSM), an open-source collaborative mapping platform, provides freely accessible geographic information, continuously updated in real-time by volunteers during crises. Despite OSM's importance in disaster response, existing literature lacks systematic, multi-disaster analyses of how volunteer mapping behaviour evolves across the disaster lifecycle. This gap is addressed through a large-scale quantitative analysis of over 51 million changes across 16 disasters, examining temporal, spatial, and content-based trends in mapping activity. We observe that volunteer contributions greatly increase immediately after disasters, accounting for 67.29% of total changes, 77.02% of which represent the creation of new map elements. Contributions subsequently shift towards the enrichment and maintenance of existing elements after the initial response. We find that the timing and intensity of volunteer activity is influenced by pre-disaster mapping completeness, while the disaster type impacts the spatial distribution of changes.

Acknowledgements

First of all, I would like to thank Professor Licia Capra for her invaluable guidance, support and enthusiasm throughout this project. Without her insights and suggestions, which shaped the direction of my progress and helped me to overcome many challenges, this project would not have been possible.

I would also like to thank the OpenStreetMap community for dedicating their time to providing geospatial data in incredible detail, without which this project and many other tools and applications simply would not be the same.

Finally, I must express my great gratitude to my friends, my girlfriend, and my family, for their constant support and encouragement throughout this project and my four years of university study.

Glossary

OSM	OpenStreetMap; a free, open-source, collaborative map of the world that anyone can edit.
HOT	Humanitarian OpenStreetMap Team; a global community supporting mapping in response to disasters.
POI	Point of Interest; a specific location such as a cafe, park, or postbox.
GIS	Geographic Information System; software for storing, analysing and presenting geographic data.
FEMA	Federal Emergency Management Agency; US Federal Agency responsible for coordinating disaster response and recovery efforts.
PII	Personally Identifiable Information; data which can be used to distinguish an individual in a dataset
MAE	Mean Absolute Error; the average absolute difference between predicted and actual values.
MAPE	Mean Absolute Percentage Error; the average absolute difference expressed as a percentage of the actual values.
USGS	United States Geological Survey; the US government agency responsible for monitoring and studying geological hazards.

Contents

Chapter 1

Introduction

In the critical days and weeks immediately following a natural disaster, timely and accurate geographic information is essential for supporting affected communities. Emergency responders rely on detailed and up-to-date maps to coordinate rescue operations, while locals depend on details about essential amenities and support services to recover and rebuild their lives. In this immediate period, commercial mapping platforms can lag behind the rapidly evolving situation on the ground, where roads, buildings, and critical infrastructure have been damaged or destroyed, but maps have not yet been updated to reflect the changes. These platforms can also be inaccessible to those affected by disasters, particularly in areas with limited internet connectivity or access to computers.

1.1 Motivation and Aims

To address these limitations, organisations and individuals turn to OpenStreetMap (OSM)[1], the largest freely accessible collaboratively mapped repository of geographic information covering the entire planet [1]. Much like Wikipedia, one of the largest online sources of factual and historic information [2], where some articles are more detailed and well-referenced than others, OSM exhibits uneven global coverage; urban areas or regions with more active mapping volunteers tend to benefit from extensive and accurate maps, while rural or less developed regions can be comparatively sparse [3]. The platform's key strength is its ability to be updated in real time, both by a global community of volunteer mappers, who in times of crisis collaborate to map and update underrepresented regions, and by mappers on the ground, with specific local knowledge about an area's geography and availability of amenities and services.

Given OSM's established role in disaster response and the availability of not just the current map, but a complete history of every contribution, or **change**, ever made, this project offers a valuable opportunity to contribute to the growing literature on digital crowdwork. The primary aim is to better understand OSM contributor behaviour by analysing and comparing patterns of mapping activity across 16 geographically diverse natural disasters. The disaster types studied include earthquakes, floods, storms, wildfires and landslides, which we hypothesise will influence mapping behaviour in distinct ways,

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

due to their differing impacts on the built environment, the spread of required changes, and the urgency of updated geographic information to support response efforts.

We first study **activity**, which captures the levels and timing of crowdworker engagement across three distinct periods: the **pre**-disaster period (before a crisis, to establish background level of activity), the immediate-disaster **imm** period (the days and weeks following, when the disaster is brought to the world's attention), and the **post**-disaster period (up to a year on from the end of the imm period, following the recovery and long-term aftermath). We then examine the **content** being mapped, specifically which real-world features in the built environment receive the most attention from contributors. In OSM, these features (e.g. roads, buildings, amenities) are represented as **elements**, which can be created, edited or deleted over time. Contributors enrich these elements, and by extension the map, with **tags**, `key=value` pairs which describe an element's function (e.g. `building=hospital`).

For both activity and content, we analyse the distribution of change types - how frequently elements are created, edited and deleted, and investigate how patterns vary across disaster periods, geographic regions, and disaster types.

While literature exists on the use of OSM during times of crisis, few studies compare mapping responses across a large number of disasters, and to our knowledge, none have investigated how activity varies throughout disasters in detail, nor the impact of disaster type on crowdworker behaviour. By performing a large-scale quantitative analysis enabled by computational methods, we seek to identify general patterns across geographic regions and disaster types that previous smaller-scale studies would have struggled to capture. Our findings are particularly relevant for humanitarian organisations such as the Humanitarian OpenStreetMap Team (HOT), which coordinates large-scale mapping efforts through teams of remote volunteers across the world, as well as local contributors on the ground. A better understanding of how mapping activity evolves throughout the disaster response period can inform the allocation of mapping resources, ensuring that maps are not only updated in the immediate aftermath but also maintained and enriched over time. Furthermore, uncovering inequalities and biases in mapping coverage could inform the design of improved workflows and tools to ensure fairer and more effective mapping in future crises.

1.2 Findings

Our analysis reveals several interesting patterns in how mapping activity and content vary during natural disasters. First and foremost, activity varies significantly across disaster periods. In the majority of cases, we observe a sharp spike in user contributions at the start of the immediate period, followed by a rapid decline in activity that eventually stabilises above pre-disaster background levels. However, in other cases, activity is more sustained, lasting into the post-disaster period, with frequent spikes of activity that are not directly correlated with the original disaster. These may reflect ongoing recovery efforts on the ground or the influence of mapping campaigns from groups such as HOT.

Enrichment, the addition of content to the map, emerges as the dominant mapping activity, with most changes involving the creation of new elements. While this is expected in poorly mapped areas pre-disaster, the trend also holds in better-mapped urban and socio-economically developed regions too. In addition to creating new elements, contributors also focus on enriching existing elements through the addition of descriptive tags and maintenance of existing tags, which occur at roughly equal rates. This suggests that, alongside expanding map coverage and richness, contributors prioritise improving the accuracy of existing data, especially during the post-disaster period.

In terms of content, mappers mostly focus their attention on `building` and `highways`, which consistently account for the majority of edits across disasters. These features represent core infrastructure, susceptible to damage during disasters, and prioritised for reconstruction to support recovery efforts. `Amenity` features are also subject to increased mapping in response to disasters, with a focus on hospitals and supplies of drinking water.

1.3 Report Structure

The remainder of this report is structured as follows:

Chapter 2 - Background:
We first provide an overview of OSM, including its history and data model, followed by a literature review on crowdwork on OSM, both during crises and general use. We then look at other crowdsourcing platforms during times of crisis.

Chapter 3 - Research Questions:
Next, we introduce the two main areas of investigation and present the corresponding research questions. Each has a number of sub-questions to guide the analysis.

Chapter 4 - Data:
This chapter begins by discussing the disasters that were selected for investigation, including how the disaster periods and boundaries were defined. It then outlines the source of the data, and details the methodologies for data ingestion, preparation and processing.

Chapter 5 - Methodology:
We then provide a detailed methodology that outlines the techniques and metrics applied to evaluate the variations in mapping activity and content across the disasters.

Chapter 6 - Results and Analysis:
Next, we present the results of the investigation, including visualisations and comparative analyses across disaster phases, regions and types. We highlight key patterns where they emerge, offering insights into the common behaviour of crowdworkers during the crises.

Chapter 7 - Conclusion:
Finally, we present a summary of the project's contributions, a critical evaluation of the methodology and results, and suggestions for future work.

Chapter 2

Background

We first present background information on the growth of OpenStreetMap and humanitarian mapping, as well as wider research into the disaster crowdsourcing field.

2.1 OpenStreetMap

2.1.1 History

In 2004, OpenStreetMap (OSM) was founded by Steve Coast at University College London (UCL) [4] as an alternative to proprietary geographic data sources such as the UK's Ordnance Survey maps that, while extensive, were not freely available. OSM was created as a free, open-source, and collaborative platform [5], which can be edited by anyone with the necessary tools to contribute, similar to Wikipedia. These can include GPS devices for logging routes of paths and roads, a personal computer for collating and uploading data, and, importantly, access to the internet. For ordinary users, the OSM website provides a simple interface to create, edit and delete geographic features, using satellite aerial imagery or local knowledge. Experienced contributors may use tools such as Java OpenStreetMap Editor (JOSM), with advanced features including bulk edits and imports. OSM is widely used as a free alternative to commercial platforms, including Google and Bing Maps, particularly for embedding maps into websites and applications, and can also replace proprietary Geographic Information System (GIS) tools such as ArcGIS, for spatial analysis, visualisation, and research of geographic data.

In 2005, the first OpenStreetMap mapping party was held to map Limehouse in London [6]. These events, where volunteers collaborate to map a specified area, have since become a cornerstone of the OSM community. They not only improve map coverage and promote free geographic data, but also serve to engage newcomers and encourage the most active contributors to socialise and make a meaningful impact through mapping [7].

OSM grew rapidly, and by 2009, surpassed 10,000 monthly active contributors, collectively making over 20 million changes each month [8], meanwhile, the community considered how geographic data from crowdsourcing could assist in response to crises [9]. On January 12, 2010, a magnitude 7.0 earthquake struck Haiti, causing widespread devastation [10]. In response, the OSM community quickly collaborated to map the affected areas using satellite imagery and historic maps. The scale of the disaster highlighted the need for coordinated on-the-ground mapping, leading to the creation of the Humanitarian

OpenStreetMap Team (HOT). Since then, HOT has coordinated reactive and anticipatory mapping campaigns with over 280M changes [11] across a wide range of natural disasters worldwide, facilitated through the HOT tasking manager, an open-source platform that enables collaborative mapping by dividing projects into manageable chunks that can be evenly distributed between volunteers. Initiatives such as the Missing Maps program (founded by HOT) to preemptively map vulnerable and undermapped areas receive support from humanitarian organisations, including the Red Cross and the United Nations, and continue to play a vital role in disaster risk mitigation and preparation [12].

Maps can be accessed online at `https://www.openstreetmap.org/`, and offline through mobile apps such as OsmAnd or MAPS.ME, which enable users to download up-to-date maps for use in areas with limited or no internet connectivity. In disaster-affected areas, these maps can guide civilians to temporary resources such as shelters, relief camps and clean water supplies, as well as essential services that may or may not have been damaged in the disaster, including hospitals, pharmacies and schools. For emergency workers, accurate and current information about road accessibility and building locations is essential for coordinating search and rescue operations and effectively prioritising resources.

2.1.2 Data Model

At the core of OpenStreetMap is a structured but flexible data model that defines how real-world geographic features are represented. As explained by founder Steve Coast, the combination of a simple element model and an open tagging system, rather than a strict hierarchical ontology, allows "mappers to map what they want, when they want" without top-down constraints [13]. Real-world objects are represented in OSM as **elements**, which form the fundamental building blocks of the map. These include roads, buildings, parks, hospitals, or smaller features such as benches or water fountains. Each element is initially created, and may subsequently be edited or deleted by contributors. We refer to these actions: `create, edit` and `delete` as **changes**, which we study in this investigation. All the changes that a user makes in a single session are grouped into a changeset, which should be limited in scope to a specific geographic area.

There are three types of elements in OSM: `node, way and relation`. Nodes represent individual points in space, defined by a longitude and a latitude. They can either represent standalone points of interest (POI), for example, a bench or a tree, or form part of a way. Ways are ordered lists of nodes that can define linear features, such as roads or railway lines, or areas, e.g. a building or forest. Relations are used to define relationships between groups of related nodes and ways. For example, a university could be a relation, made up of multiple buildings and paths (ways) and POIs (nodes). Due to their complexity and unbounded size, relations are more difficult to analyse and are therefore excluded from this investigation. All elements share a common set of required attributes to track their identity and history. These are automatically updated as contributors make changes, and are summarised in Table 2.1.

Attribute	**Description**
`id`	Uniquely identifies the element within its type (a node and a way can both have the same ID number).
`version`	Tracks how many times the element has been modified. The default is 1, increments with each subsequent change.
`changeset`	The ID of the changeset in which the element was last modified.
`timestamp`	The exact date and time when the latest change was made.
`user`	The display name of the contributor who made the change.
`uid`	The unique identifier of the contributor.
`visible`	Indicates whether the element is currently visible on the map. If set to false, the element has been deleted but can still be restored.

Table 2.1: Core metadata attributes of OSM elements

In addition to the required attributes that are populated automatically, elements may optionally include one or more `tags`, composed of user-defined `key=value` pairs. Tags are used to provide additional information beyond a simple geographic position. They can define the purpose of an element or provide additional context. For example:

- `amenity=hospital` - a node that represents a hospital, a type of amenity.
- `highway=residential` – a way that represents a residential road with houses, as opposed to a motorway, which would be tagged as `highway=motorway`.
- `opening_hours=Mo-Su 07:00-23:00` - a POI is open daily from 7 am to 11 pm.

Where a tag key represents a geographic feature, such as a `building`, we define the tag value as the **feature type** e.g. `residential`. There is no fixed dictionary of tags, but conventions for how they should be used are agreed upon in the OSM community and documented online[1]. Their open-ended nature ensures that mappers have the freedom to map the world as it is, without needing to adhere to strict constraints. This flexibility is valuable during times of crisis, allowing users to enrich maps quickly without worrying about rigid tagging constraints. Figures 2.1 - 2.3 show examples of nodes and a way in OpenStreetMap XML format.

[1] https://wiki.openstreetmap.org/wiki/map_features

Node:

```
<node id="8147204370" visible="true" version="1" changeset
    ↪ ="94521963" timestamp="2020-11-20T20:08:07Z"
      user="---" uid="---" lat="51.5238701" lon="-0.1318017
          ↪ ">
    <tag k="amenity" v="bench"/>
</node>
```

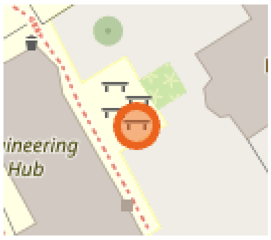

Figure 2.1: Example of a node element representing a bench on the UCL campus.

Way:

```
<way id="4253295" visible="true" version="17" changeset="
    ↪ 153265615" timestamp="2024-06-27T20:40:58Z" user="
    ↪ ---" uid="---">
    <nd ref="698447403"/>
    <nd ref="7921311836"/>
    ...
    <nd ref="9450338311"/>
    <nd ref="2442016051"/>
    <nd ref="185620597"/>
    <tag k="access" v="customers"/>
    <tag k="bicycle" v="dismount"/>
    <tag k="foot" v="customers"/>
    <tag k="highway" v="pedestrian"/>
    ...
    <tag k="name" v="Malet␣Place"/>
    <tag k="oneway" v="no"/>
    <tag k="surface" v="asphalt"/>
</way>
```

Figure 2.2: Example of a way element representing Malet Place, UCL. Several nodes and tags are omitted to save space.

```
<node id="698447403" visible="true" version="6" changeset=
    ↪ "93751961"
      timestamp="2020-11-08T18:05:07Z" user="---" uid="---"
      lat="51.5238191" lon="-0.1328913"/>
```

Figure 2.3: The first node - `id:698447403`, at which the way starts.

Typically, OSM users access the latest version of the map, where only the most recent state of each element is visible. However, to investigate how mapping activity and content vary across disaster periods, we instead analyse OSM history files. These contain every version of each element ever created or edited in a given area; each time an element is changed, for example, to add a tag or move a node, its version is incremented, and a new copy of the element is stored in the history. This allows for the study of activity before, during and after disasters, with each versioned change forming the basis of our analysis.

2.2 OpenStreetMap

2.2.1 Accuracy and Coverage

Early research into OSM focused on assessing the accuracy and coverage of OpenStreetMap, to better understand the reliability of the map. Both Haklay [14] and Fan et al. [15] found that the positional accuracy of mapped features is very high, but the completeness of tag attributes can be lacking. Mashhadi et al. [16] showed that coverage can vary significantly, even within a single city, in their case, London. Herfort et al. [3] later applied a machine learning model, trained to account for real-world uneven spatial coverage, to estimate mapping completeness of over 13,000 cities worldwide. They revealed that 14% of the cities (16% of the population) had over 80% map completeness, while 69% of cities (48% of the population) had completeness below 20%. Completeness tended to be higher for cities in Europe, North America and Central Asia, and lower in Latin America and the Caribbean, East Asia, the Middle East and North Africa.

2.2.2 Maintenance Practices and Biases

Subsequent research into OSM has focused on who contributes to OSM and how the map is maintained over time, reflecting the dynamic nature of the real world. A healthy crowdsourcing platform should reflect a broad base of contributors, rather than being dominated by a small subset of active power users, who may lack local knowledge. Quattrone et al. [17] investigate maintenance practices on OSM, specifically users editing existing elements, rather than creating new ones. The study examines how widespread maintenance is across 117 countries over a year, which types of changes are most frequent, and who performs them. They find that maintenance is mostly carried out by experienced users, who commonly add new tags, especially for elements with few or no tags. Maintenance activity is unevenly distributed, tending to be concentrated in areas where experienced contributors are based: the most frequent maintainers of the map.

2.2.3 OpenStreetMap During Times of Crisis

A growing body of research has explored how OSM is used and contributed to during times of crisis. These studies typically examine contributor activity, spatial coverage and mapped content. Large-scale natural disasters, as well as the COVID-19 pandemic, are frequently studied due to their wide impact and substantial activity generated.

As discussed in [9] and [18], the Haiti 2010 earthquake marked a pivotal moment for the use of crowdsourced geographic data in humanitarian response, proving its effectiveness. In the immediate aftermath, large parts of Haiti lacked detailed maps, particularly the capital, Port-au-Prince. Over 3 weeks, 600 remote contributors built the map from scratch, with support from local individuals; these efforts paved the way for more formal mapping processes and the formation of HOT. Xu et al. [19] subsequently performed a spatial-temporal analysis of mapping activity following Hurricane Matthew's impact on Haiti in 2016. They observed a strong spike in contributions the day the hurricane made landfall, concentrated in areas hardest hit by damage. Roads received more editing attention than buildings, possibly due to their importance for delivering supplies.

The 2015 Nepal Earthquake was also widely studied due to its extensive damage and widespread news coverage. Poiani et al. [20] investigated the response, focusing on the coordination of volunteer contributions. They identified large spikes in mapping activity and active users, including newcomers, with 75% of members making their first edit on the same day they were registered on the HOT project. Edits were primarily distributed in the Kathmandu valley, the most affected area, and Nepal's largest population centre. Coordination was organised by HOT through mailing lists and the Tasking Manager.

The COVID-19 pandemic created an urgent need for accurate, localised geographic data to help the public access health services and combat the spread of the virus. Minghini et al. [21] analysed the Italian OSM community's response during the first wave, highlighting rapid mapping of government-designated red zones, updates to healthcare facilities, and the addition of delivery-related business information. Over 1,000 pharmacies were manually verified and added, avoiding a bulk import due to major location errors in official datasets. New tagging schemes were introduced to reflect changing pandemic rules, and custom maps were created to visualise non-standard COVID data.

Kamptner and Kessler [22] examined four small-scale crisis events, specifically building fires in urban areas, to investigate whether small-scale events attract a greater proportion of local contributors, as well as the variations in mapped features. All events resulted in significant increases in contributions within 10 days of the incident, with edits densely concentrated around the affected building. Surrounding features such as roads, footpaths, bike paths, and shops were frequently edited, indicating that contributors focused not only on the damaged structure but also on the broader local context. The majority of contributors across all events were non-local and globally distributed, with most being experienced mappers rather than newcomers.

While most studies focus on a single or a small number of disasters, Hertfort et al. [23] and Dittus et al. [24] analysed a broader range of events to identify patterns in contributor behaviour. Dittus et al. [24] examined 26 mapping campaigns led by HOT task forces, comparing contributor dynamics between **event-centric** campaigns, organised immediately in response to disasters, and **mission-centric**, aimed at proactively mapping areas. They found that newcomers tended to support mission-centric campaigns more than event-centric ones, where their contributions tended to be of lower quality.

Meanwhile, Herfort et al. [23] present a large-scale quantitative analysis of all humanitarian mapping from the HOT Tasking manager from 2012 to 2020, using regression models to examine relationships between mapping intensity and development indicators. Focusing on roads and buildings, they found that although humanitarian efforts targeted regions with low human development, these areas were still disproportionately less covered on the map relative to their population size. The analysis was focused on the long-term evolution of humanitarian mapping and did not investigate how contributor behaviour varied during individual disasters.

2.3 Other Crowdsourcing Platforms

2.3.1 Twitter

Twitter is particularly active during times of crisis, as shown by Maldeniya et al. [25] who studied tweeting patterns across 200 disaster-affected US communities, examining how user emotions vary, before, during and after disaster strikes. Analysing 200 million tweets from 2 million users, they found that activity largely follows consistent behavioural trajectories, though approximately 25% of events resulted in "substantially heightened level of emotions", with resurgences in fear and anger in areas most severely affected.

Ahmouda et al. [26] compared user contributions to OSM and Twitter during the 2015 Nepal Earthquake and the 2016 Norcia earthquake in Italy, focusing on short and long-term changes in mapping and tweeting behaviour. In both disasters, significant spikes in the number of new ways, nodes, users, and geometry changes occurred around the time of the disaster, with activity usually declining within two months. The composition of contributors differed; over half of the contributors in the Nepal earthquake were local, whereas in Italy, they only made up 13% of the mappers, the rest being international. Geotagged tweet activity also increased after each event, though to a lesser extent, with a more localised focus compared to broader activity on OSM.

2.3.2 Wikipedia

Wikipedia is the world's largest online encyclopedia [27], comprising over 6.9 million articles in English alone [2]. As a free, collaboratively edited resource, Wikipedia plays a central role in the dissemination of information online. Contributors frequently update content in response to recent events, such as the death of public figures [28]; this responsiveness also applies during rapidly evolving natural disasters and health crises.

Ruprechter et al. [29] investigated how the mobility restrictions imposed on populations by the COVID-19 pandemic affected contributions to Wikipedia. Analysing 223 million edits from 2018 to 2020 across 12 languages, they found a significant increase in activity during lockdown periods, particularly in the English edition, along with an influx of newcomers. Only a "negligibly small fraction of edits" were focused on COVID-19, demonstrating a broad increase in engagement across the platform during the crisis.

Finally, Kurek et al. [30] investigated Wikipedia maintenance practices during the ongoing conflict in Ukraine, examining how experienced contributors manage disruption and misinformation, given Wikipedia's value as an authoritative resource worldwide. Editors reported disruption from both Russia and Ukraine-aligned users, including frequent edit wars, an extensively studied phenomenon on Wikipedia [31]. However, they did not feel that this activity was state-sponsored. The editors noted built-in tools on Wikipedia to combat vandalism, including page protections. Despite the benefits of publicly accessible up-to-date maps, the OpenStreetMap community has restricted editing activity in Ukraine, citing concerns that open geographic data can be used to plan attacks on military and civilian infrastructure [32], as well as to reduce the potential for politically motivated editing. These restrictions reflect the sensitive nature of the situation, where the safety of individuals and communities remains a priority.

2.4 Summary

Existing literature on OSM and other platforms during crises has provided valuable insights into mapping activity, the types of features contributed, and the characteristics of contributors. However, spatial and temporal analyses have been limited to studies of individual or small numbers of disasters. In contrast, multi-disaster studies overlook how activity and content vary across the different phases of each disaster. This project addresses this gap by performing a large-scale analysis of mapping behaviour across multiple disasters, focusing on temporal variations in activity and contributed content.

Chapter 3
Research Questions

To address the gaps in the literature, this study investigates how OpenStreetMap contributors responded to a range of natural disasters, by conducting a spatio-temporal analysis of mapping activity and the content they map. The aim is to identify patterns across disaster type and geographic region, to understand whether these factors may influence the intensity of mapping efforts.

Two key research areas are focused on: how mapping activity varies across the lifecycle of a disaster, and what types of contributions users make. Each area is explored through a main research question, supported by a set of sub-questions.

3.1 RQ1: How does mapping activity vary across disasters during times of crisis?

We define **activity** as the number of create, edit and delete changes made by users across a given disaster area, within a defined period. We study activity because it reflects the level and timing of contributor engagement, enabling us to understand how the community responds at different stages of a crisis. Three aspects of activity are studied:

RQ1.1: **How** does the **change activity** vary across disaster periods by change type?
We investigate how the counts of *create*, *edit*, and *delete* changes vary across the pre, imm, and post-disaster periods, aiming to identify spikes and patterns in activity.

RQ1.2: **How** does the **interval between consecutive changes** vary?
We aim to determine whether the frequency of updates to individual elements varies across disaster periods.

RQ1.3: **How** does the **spatial distribution** of activity vary within disaster areas?
We study the location and spread of changes for each disaster, and assess whether it becomes more or less equal across periods.

3.2 RQ2: How does the content contributed to the map evolve throughout disasters?

We define **content** as the geographic features (e.g. building, highway) and their corresponding feature types (e.g. `residential`, `primary`), as well as the nature of these changes. We aim to investigate how mapping priorities shift throughout a disaster by analysing which features are changed over time, what attributes of existing elements are edited, and when contributors change tags.

RQ2.1: **Which geographic features** are most commonly changed during disasters?
We first identify which types of features are most frequently edited by contributors during the different phases of disasters.

RQ2.2: Among the specified geographic features, **which feature types** are most commonly changed?
Having highlighted a set of the most commonly changed geographic features, we then look at which feature types are most commonly the subject of changes. As a reminder, for an element with a tag: `amenity=cafe`, `amenity` is the geographic feature, and `cafe` is the feature type. This reveals in greater detail which places and infrastructure are prioritised in disaster mapping.

RQ2.3: **What attributes** of elements do contributors change in their edits?
We then investigate the specific attributes being modified during edits, such as changes to coordinates, tags, and visibility status. This helps to reveal the main purpose of contributions.

RQ2.4: **When** are changes involving element **tags** made most frequently?
Finally, we compare overall mapping activity, with activity of changes specifically involving element tag additions, edits, or deletions, to assess whether contributors tend to update descriptive information at different periods in the disaster lifecycle.

By answering these questions, we may gain a deeper understanding of how crowd workers support mapping efforts during natural disasters, both in terms of activity patterns and the nature of their contributions.

Chapter 4

Data

This chapter discusses which disasters were investigated and how the disaster areas were defined. It then considers the methodologies for acquiring, preprocessing and filtering the OpenStreetMap historical data.

4.1 Natural Disaster Selection

To perform a meaningful and comparative analysis of crowdmapping behaviour, 16 disasters were selected according to three primary criteria:

1. **Recency**: Disasters took place within the past ten years to ensure relevance to modern OSM mapping practices. An exception was made for the 2010 Haiti earthquake, due to its historical significance as the foundational event for humanitarian mapping on OSM.
2. **Impact**: Disasters were selected based on the availability of clear documentation describing their geographical extent and humanitarian impact, ensuring that each event would have prompted substantial mapping activity on OSM.
3. **Variety**: A diverse selection of disaster types was included: earthquakes, floods, storms, wildfires, and landslides, across four geographically and economically diverse regions: the Americas, Europe and the Middle East, Africa, and Asia.

Individual disasters were identified through sources such as the United States Geological Survey (USGS) Earthquake Hazard Program catalogue, Wikipedia disaster lists, and academic or journalistic reporting. For each event, maps and satellite imagery were used to define an area of interest within each country, capturing the regions most affected by the disaster and focusing analyses on areas with the most expected mapping activity. These boundaries were stored in `GeoJSON` format to support computational analysis. An example of the defined boundary for the 2010 Haiti earthquake is shown in Figure 4.1, and the list of disasters studied, their geographic region, total area of analysis and relevant references are presented in Table 4.1. Six of the disasters studied involved earthquakes, six were storms or floods, two were wildfires, and two were landslides.

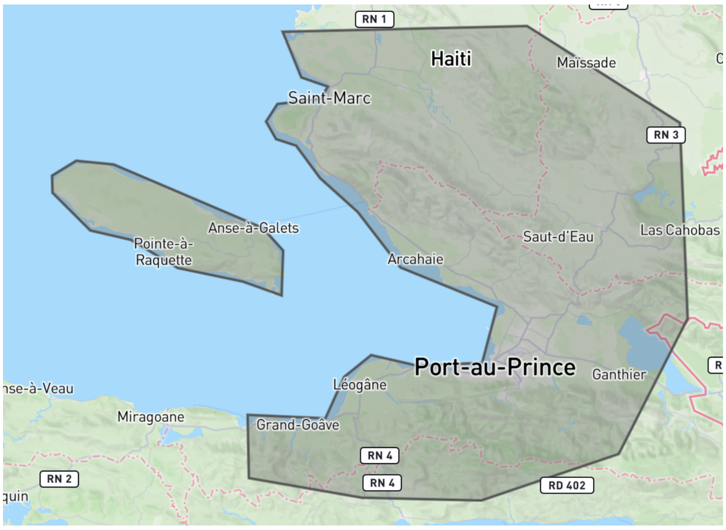

Figure 4.1: Example of a disaster-specific boundary defined in `GeoJSON` format, covering the most affected areas during the 2010 Haiti earthquake.

Table 4.1: List of selected disasters with their geographic region, type, area and citations.

Region	**Disaster**	**Type**	**Area (km^2)**	**Citations**
Americas	Haiti (2010)	Earthquake	8,269.16	[33]
	Texas, US (2017)	Storm (Hurricane)	56,991.57	[34, 35]
	California, US (2020)	Wildfire	14,389.24	[36, 37]
	Haiti (2021)	Earthquake	6,794.69	[38, 39]
Asia	Nepal (2015)	Earthquake	34,204.74	[40]
	Sulawesi, Indonesia (2018)	Flood (Tsunami)	31,561.36	[41]
	Atami, Japan (2021)	Landslide	3.24	[42]
	Pakistan (2022)	Flood	98,600.58	[43]
Europe & Middle East	Attica, Greece (2018)	Wildfire	155.47	[44]
	Izmir, Turkey (2020)	Earthquake	6,367.98	[45]
	Gaziantep, Turkey (2023)	Earthquake	20,483.32	[46]
	Emilia Romagna, Italy (2023)	Flood	4,375.90	[47, 48]
Africa	Freetown, Sierra Leone (2017)	Landslide	26.47	[49]
	Derna, Libya (2023)	Flood (Dam Collapse)	24.80	[50]
	Malawi (2023)	Storm (Cyclone)	20,139.22	[51, 52]
	Morocco (2023)	Earthquake	9,481.95	[53]

4.2 OpenStreetMap Data

This next section outlines the source of historical OSM data, how the periods of a disaster are defined, and the methods employed for data storage, preparation, and filtering.

4.2.1 Data Acquisition and Storage

Historical OSM data was downloaded from Geofabrik's download server[1], which provides full-history extracts in `.osh.pbf` format for many geographic regions and countries. These files contain the complete editing history of all OSM elements, including all attributes referenced in Table 2.1. As these downloads include Personally Identifiable Information (PII) (user IDs and display names), appropriate steps were taken to anonymise the data during processing. No PII was retained beyond what was necessary for aggregate analysis, as user-level data was not necessary for the investigation.

All change in the `.osh.pbf` history files were initially imported into a PostgreSQL database, with PostGIS for spatial processing. Each change corresponds to a versioned modification of an OSM element, e.g. moving a node to a different position, or adding, editing, or deleting a tag. Changes were stored in the `changes` table, for which the schema is shown in Listing 4.1. Information about each disaster was stored in a separate `disasters` table, including the disaster type, date, and GeoJSON boundary geometry. Python scripts were written for all data insertion, querying and analysis; the project GitHub repository URL is included in the appendix.
`Scripts/database/db_bulk_insert.py` was used to import the changes into the database, and all queries can be found in `Scripts/database/db_utils.py`.

```sql
CREATE TABLE changes (
    id SERIAL PRIMARY KEY,
    element_id BIGINT NOT NULL,
    element_type VARCHAR(10),
    edit_type VARCHAR(10),
    timestamp TIMESTAMP,
    disaster_id INTEGER,
    version INTEGER,
    visible BOOLEAN,
    changeset BIGINT,
    tags JSONB,
    building BOOLEAN,
    highway BOOLEAN,
    coordinates GEOMETRY(POINT, 4326),
    uid BIGINT,
    geojson_verified BOOLEAN DEFAULT FALSE
);
```

Listing 4.1: Schema for the `changes` table

```sql
CREATE TABLE disasters (
    id INTEGER PRIMARY KEY,
    country VARCHAR(50)[],
    area VARCHAR(50)[],
    area_geometry GEOMETRY(
        ↪ MULTIPOLYGON, 4326),
    date TIMESTAMP,
    h3_resolution SMALLINT,
    disaster_type VARCHAR(50)
);
```

Listing 4.2: Schema for the `disasters` table

[1] https://osm-internal.download.geofabrik.de/index.html

4.2.2 Disaster Periods

The disaster study periods were defined in line with the US Federal Emergency Management Agency's (FEMA) phases of emergency management [54], allowing mapping activity to be segmented into three distinct phases that reflect the typical lifecycle of a disaster: normal operations (pre), immediate response (imm), and recovery (post). The temporal definitions of each period are shwon in Table 4.2:

Table 4.2: Definitions of each disaster period

Period	**Start**	**End**	**Description**
Pre	365 days before disaster date	Day before disaster date	Baseline period to estimate the typical background rate of mapping activity.
Imm	Disaster date	60 days after disaster date	Intense immediate response period: concentrated rescue and recovery efforts, with possible international media attention.
Post	Day after end of Imm period	365 days after end of Imm period	Longer-term recovery phase, focusing on rebuilding efforts with possible sustained mapping activity.

4.2.3 Pre-processing and Filtering

Having imported the changes into the database, preparation and filtering of the data were performed to ensure only valid changes were retained for analysis. For the time series analysis, up to three years of pre-disaster changes were imported, and all changes made in the immediate and post-disaster periods. Changes made more than three years before the disaster were also included when they represented the most recent prior version of an element, required for determining the interval between and nature of each edit.

Next, all changes made outside the defined `GeoJSON` boundaries for each disaster were removed using `Scripts/database/db_geojson_filtering.py`. This ensured that only geographically relevant edits within the affected areas were retained for analysis. Bulk import filtering was then applied in `Scripts/database/bulk_import_filtering.py` to remove changes likely made through automated imports rather than an individual contributor, as they do not represent the behaviour of a typical human crowdworker. A similar approach to that used in [22] was followed, but with different parameters to provide greater control over what was classified as a bulk import. These were derived through trial and error and may not have captured all automated activity. A changeset and its changes were excluded if:

- It contained more than 5000 changes, completed within a 30-minute window, with over 95% of the changes being of the same type (e.g., all `create` operations).

- It contained more than 3000 changes, completed within a 1-minute window, with over 95% of the changes being of the same type.

A changeset with 1000 edits that would have been excluded in [22] could still reflect intensive human activity, therefore, the more conservative exclusion threshold of 3000 changes was adopted in this study. Initially, the dataset contained **70,286,821** changes, comprising three years of pre, immediate and post-disaster changes, and older changes to preserve prior versions of modified elements. A total of 8,363,284 changes (11.9%) were removed through the GeoJSON boundary filtering, and a further 2,604,296 (3.7%) were excluded by the bulk import filtering. This resulted in a final dataset of **59,319,241** changes across 16 disasters. Of these, **51,147,716** occurred within the study's defined pre, imm, and post-disaster periods, and are summarised in Table 4.3.

Table 4.3: The number of OSM changes across the study's pre, imm, and post-disaster periods, the total number of changes in these phases, and the corresponding change density per km^2. Considerable variation in both change counts and density is observed across disasters, reflecting the differences in disaster impact and scale of response.

Disaster	**Pre**	**Imm**	**Post**	**Total**	**Density (changes/km^2)**
Haiti (2010)	3,684	1,047,263	156,212	1,207,159	145.98
Texas (2017)	499,242	1,314,152	1,159,910	2,973,304	52.17
California (2020)	55,326	12,432	97,351	165,109	11.47
Haiti (2021)	206,992	1,101,494	47,178	1,355,664	199.52
Nepal (2015)	1,116,742	14,485,900	1,071,542	16,674,184	487.48
Sulawesi (2018)	2,210,547	1,000,022	443,217	3,653,786	115.77
Atami (2021)	714	8,980	4,116	13,810	4,262.35
Pakistan (2022)	936,232	758,706	647,163	2,342,101	23.75
Attica (2018)	4,591	5,642	49,850	60,083	386.46
Izmir (2020)	302,147	1,572,247	1,242,850	3,117,244	489.52
Gaziantep (2023)	220,581	8,493,227	1,414,791	10,128,599	494.48
EmiliaRomagna (2023)	206,816	142,688	458,645	808,149	184.68
Freetown (2017)	5,432	23,574	29,137	58,143	2,196.56
Derna (2023)	76	102,268	18,146	120,490	4,858.47
Malawi (2023)	682,785	1,920,144	2,522,595	5,125,524	254.50
Morocco (2023)	22,093	2,429,962	892,312	3,344,367	352.71
Total	**6,474,000**	**34,418,701**	**10,255,015**	**51,147,716**	—

Chapter 5

Methodology

In this chapter, we define the metrics and methodologies used to answer the research questions introduced in Chapter 3. We begin by outlining the techniques and metrics for performing a temporal and spatial analysis of crowdworker activity, folowed by a description of how the content of contributions was assessed.

5.1 RQ1: How does mapping activity vary across disasters during times of crisis?

RQ1.1: How does the change activity vary across disaster periods?
Initially, we computed the **aggregate count of changes** in each period for every disaster, to observe overall trends in contributor behaviour across the disaster periods. **The percentage difference** in total mapping counts between the pre-disaster and the immediate or post-disaster periods was then computed as shown in Equation 5.1.

$$\frac{\text{Count}_{\text{imm/post}} - \text{Count}_{\text{pre}}}{\text{Count}_{\text{pre}}} \times 100 \tag{5.1}$$

Similar percentage difference calculations are used in other sub-research questions. Following a hypothesis-driven approach, we expect mapping activity to increase for all disasters in the immediate period, and remain elevated above the pre-disaster levels in the post-disaster phase, possibly decaying over time.

Next, the proportion of changes by change type (`create`, `edit`, `delete`) was computed to determine whether contributors were primarily enriching the map or maintaining existing elements in each period. Based on our hypothesis, we expect the immediate period would consist mainly of `creates` in previously undermapped areas, and `edits` in well-mapped areas, while the post-disaster period would be dominated by `edit` activity reflecting ongoing maintenance.

We then performed a temporal analysis to examine how the rate of changes by mappers varies at a finer granularity during the three periods. To support this, we adopted a generalised methodology which is applied across several subsequent sub-questions:

- **Intervals**: Each disaster period is divided into fixed-length intervals (daily, weekly, or monthly), depending on the desired granularity.

- **Metric**: For each interval, a metric is computed. In RQ1.1, we calculate the number of `create`, `edit` and `delete` changes as well as the `total`, by binning each change according to its timestamp.
- **Forecasting**: In addition to computing the metric, Prophet [55] time series forecasting is used to generate a prediction for expected behaviour in the post-disaster period, assuming no disaster occurred. Models for each interval and metric are trained on three years of pre-disaster data, excluding the 60 days immediately preceding the disaster to avoid pre-crisis anomalies. Data from periods affected by COVID-19, from January to June of 2020, were also excluded due to the global disruption and outlier effect on mapping, except for California and Izmir due to their pre-disaster overlap. Despite tuning, Prophet occasionally produced implausible forecasts, including negative values, particularly for disasters with low or sporadic pre-disaster activity. Adjustments to changepoints and training windows were explored, yielding minimal improvement. To account for this, all forecasts were clipped at zero to enforce realistic lower bounds.
- **Evaluation**: The Prophet forecasts are then compared with the actual post-disaster results to determine how much each metric has been affected by the disaster, and to what extent it returns to pre-disaster norms. For RQ1.1 forecasted change counts are compared against the observed counts to determine whether post-disaster activity remains heightened.

  To quantify forecast accuracy, we compute the Mean Absolute Error (MAE) and Mean Absolute Percentage Error (MAPE), shown in Equations 5.2 and 5.3,

  $$\text{MAE} = \frac{1}{n}\sum_{i=1}^{n} |y_i - \hat{y}_i| \quad (5.2) \qquad \text{MAPE} = \frac{100}{n}\sum_{i=1}^{n} \left|\frac{y_i - \hat{y}_i}{y_i}\right| \quad (5.3)$$

  where y_i is the actual value, $\hat{y}_i$ is the prediction and n is the number of observations.
- **Visualisation**: Finally, plots of the observed and forecasted metrics for each disaster and interval are generated to interpret the results.

Across all analyses, weekly intervals are used to balance temporal granularity with noise reduction, capturing week-by-week trends while smoothing out daily fluctations.

Metric Summary (RQ1.1):

- Aggregate change counts in each disaster period, and the percentage difference from pre → imm and pre → post.
- Weekly change counts by type (`create`, `edit`, `delete`).
- Prophet-predicted weekly change counts, with MAE and MAPE.

RQ1.2: How does the interval between consecutive changes vary?
Next, we investigate how **frequently** individual elements are re-edited across the study periods. This offers insight into the pace of enrichment or maintenance activity; shorter intervals indicate more frequent updates, while longer intervals suggest stability or reduced contributor attention. To achieve this, we first apply the **change differences** methodology introduced in RQ2.3. For each change made in the study period of the disasters, we compute the number of days since the element was last edited or created.

The temporal analysis methodology outlined in RQ1.1 is then followed: each change and the days since the last edit is binned in the intervals. For the changes within each interval, the median number of days between edits is calculated - the metric for this sub-question. As before, Prophet forecasting is applied to model how the edit frequency could have varied in the post-disaster period, had no disaster occurred. Forecasts are compared against actual values using MAE and MAPE, and the results are visualised.
Metric Summary (RQ1.2):

- Weekly median number of days between consecutive changes.
- Prophet-predicted weekly median number of days, with MAE and MAPE.

RQ1.3: How does the spatial distribution of activity vary in disaster areas?
We now take a more detailed look at each disaster area to understand how mapping activity is spatially distributed. The aim is to visualise where changes are **concentrated**, such as urban or rural areas, as well as whether changes follow key infrastructure, e.g. roads or rivers. To achieve this, binned choropleth maps are generated to highlight areas of mapping activity. The map is broken up into a grid of hexagons forming spatial units using Uber's H3 Hexagonal Hierarchical Spatial Index [56], which have configurable resolutions. A resolution of 8 (0.73 km^2) is sufficient for most disasters, providing coverage at the scale of a small town, while a finer resolution of 9 (0.10 km^2) is useful for distinguishing activity within dense urban areas. For each hexagon, the number of changes made to elements whose coordinates fall within its bounds is computed, and the hexagons are coloured accordingly to reflect the density of mapping activity.

As shown in Figure 6.7a in the results chapter, it is difficult to identify which areas were targeted by mapping activity when looking at raw counts. To address this, we instead compute the percentage difference in change counts from pre-disaster to the immediate and post-disaster periods. The resulting choropleth maps are then coloured according to the percentage difference, with a brighter red indicating a positive percentage increase in mapping activity, and a brighter green signifying a greater decrease.

Having computed and plotted the change counts for each spatial unit, it is clear that some areas receive significantly more attention than others. However, visual inspection alone is insufficient for quantifying this variation. A discrete approximation of the Gini coefficient is computed, based on the formulae presented by Dorfman [57], to compare the **inequality** in spatial distribution of mapping activity across disasters and periods. This approach computes inequality from the sorted distribution of change counts across spatial units (H3 hexagons). The coefficient ranges from:

0 → perfect equality (changes evenly distributed across all hexagons), to

1 → perfect inequality (all changes concentrated in a single hexagon).

The coefficient is calculated as shown in Equation 5.4,

$$\text{GINI} = \frac{n + 1 - 2 \cdot \frac{\sum_{i=1}^{n} (n+1-i) \cdot x_i}{\sum_{i=1}^{n} x_i}}{n} \tag{5.4}$$

where x_i is the number of changes in the i-th hexagon, sorted in ascending order, and n is the total number of hexagons.

The Gini coefficients are computed for each disaster, by change type (`create`, `edit`, `delete`, `total`) and period. To assess whether mapping activity becomes more spatially concentrated or distributed in the aftermath of a disaster, we calculate the percentage difference in Gini coefficient between the pre-disaster period, and the immediate/post-disaster periods.

Metric Summary (RQ1.3):

- Total change counts per hexagon in each period and percent difference from pre → imm and pre → post.
- Gini coefficient in each period and percent difference pre → imm and pre → post.

5.2 RQ2: How does the content contributed to the map evolve throughout disasters?

RQ2.1: Which geographic features are most commonly changed?

We begin by investigating which geographic features contributors tend to prioritise throughout the disaster lifecycle. A top-down, hypothesis-driven approach is adopted, focusing on geographical features of the built environment that are particularly relevant in the context of disasters. We first focus on elements with the tag keys `building` and `highway`, which represent physical infrastructure that may be damaged during disasters and are critical to effective crisis response. According to OpenStreetMap's *TagInfo* database [58], these are the top two most frequently mapped elements on the platform. While these features have been examined in prior work, such as by Herfort et al. [23], we go further by examining the variations throughout the disaster, hypothesising that both their activity and proportion of total changes will increase during and after disasters.

We additionally investigate changes to `amenity`-tagged elements, which include facilities such as hospitals, shelters, pharmacies and drinking water, and are hypothesised to become more relevant during disasters, reflecting contributors' focus on essential services. To contrast this, we also examine `leisure`-tagged features, which are expected to receive less attention during and after disasters, reflecting a shift in mapping priorities away from non-essential features during times of crisis.

To assess which geographic features are prioritised during disasters, we count how often each tag key appears across all changes within each period. Since an OSM element can have multiple tags, we focus on the four distinct features which tend not to overlap[1], to ensure that each change reflects a meaningful instance of mapping the given feature.

For each feature, we compute the following metrics:

- **Aggregate usage count**: the number of changes involving the feature for each disaster and across all disasters, segmented by period.
- **Proportion of total changes**: the percentage of all changes in each period that are made to the feature, to compare relative importance over periods.

The aggregate counts are presented in a table, and the proportion of total changes is visualised with a stacked area chart to compare mapping priorities across periods.

RQ2.2: Among the specified geographic features, which feature types are most commonly changed?

Having established which geographic features are most commonly changed in RQ2.1, we now investigate which specific feature types (tag values) of these features are most commonly changed. A similar approach to RQ2.1 is followed, where we count the number of times each tag value appears in a change alongside its corresponding feature key, for example, the number of changes of elements with tags `building=house` or `amenity=cafe`. We then compute the percentage of changes where the `key=value` pair appears, given that the changed element includes the key as a tag, for example, among all changes to elements tagged with `building=*`, we calculate the proportion with `building=house`, or

[1] Less than 5% of elements tagged as `amenity` are also tagged as `building`

building=commercial. These form the main metrics for this sub-question.

The proportions are aggregated across all disasters and analysed separately for each geographic feature; for each, we select the twelve most frequently mapped feature types during the pre-disaster period and track how their relative usage varies across the immediate and post-disaster periods. These changes are also visualised using stacked area plots to observe shifts in mapping priorities over time.

Additionally, we compute the Kendall Tau Rank Correlation Coefficient (τ) [59, p. 752], which measures the similarity between the orderings of feature types across disaster periods for each feature by disaster. By comparing the τ correlations, we identify which geographic features exhibit greater variation in mapping priorities across periods, and which tend to be more consistent. Kendall's τ is used to compare the relative priority of feature types, rather than raw counts, allowing for consistent comparison across disasters with differing mapping volumes.

The coefficient for two ordered lists ranges from:

$-\mathbf{1}$ → perfect disagreement (top 12 feature types in opposite order), to

$+\mathbf{1}$ → perfect agreement (top 12 feature types in same order).

We can expect that the majority of disasters will have similar orderings, with $\tau > 0$.

To compute τ for each disaster, the list of the twelve most changed feature types per geographic feature for each period is selected. τ is then calculated across three period pairs, from the two lists: pre–immediate, pre–post and immediate–post. Since the computation requires identical sets of ranked items, an inner join is performed on the feature types that appear in both periods, excluding any feature types that are not in the top 12 of both periods.

τ is computed using the Python library scipy's implementation of Kendall's τ-b, as defined by Equation 5.5:

$$\tau = \frac{C - D}{\sqrt{(C + D + T_x)(C + D + T_y)}} \tag{5.5}$$

where C is the number of concordant pairs, D is the number of discordant pairs and T_x, T_y are the number of ties in each list respectively.

Once the coefficients have been computed for each period pair across all geographic features and disasters, the median, interquartile range (IQR) and range are calculated for each period. The same statistics are then computed across all periods for the features, yielding an IQR and median for each, with which their consistency may be compared.

Metric Summary (RQ2.2):

- Count and proportion of changes for each feature type within its geographic feature.
- Median, IQR and range of Kendall's τ from feature type orderings.

RQ2.3: What attributes of elements do contributors change in their edits?
So far, we have examined the activity associated with changes to specific geographic features, as well as overall spatio-temporal activity by contributors. We now investigate the underlying attributes of map elements that contributors modify or change in their **edits**. These changes can broadly be understood as either **enrichment** or **maintenance** of the map. Enrichment refers to the addition of new content, such as the creation of new elements, or as we examine edits in this section, the addition of new descriptive tags that provide greater detail about an existing element. Maintenance involves the modification or removal of existing content, as outlined below.

To analyse the attributes that contributors most frequently change in their edits, we first categorise the four types of **edit** changes that a user can make to OSM elements. We focus on deliberate user actions and exclude automated attribute updates to metadata in response to edits. The edit types are:

- **Tag Modifications**: Contributors can create new tags, edit the values of existing tags, or delete them.
- **Node Coordinate change**: Mappers can move nodes on the map, updating their geographic position. This applies to nodes representing POIs and way nodes.
- **Way Geometry change**: The list of nodes that make up a way can be modified; nodes can be added, removed or reordered.
- **Delete Reversion**: Deleted elements remain in the OSM history and can be made visible on the map again.

To support the analysis of the specific contributions made during edits, we introduce the **change differences** methodology. This approach compares each version of an element (represented by a change) with its immediate predecessor to identify the difference, or `diff` between them. A predecessor refers to the same OSM element in the edit history with a version number exactly one less than the current change (i.e., `version`-1). As discussed in Section 4.2.3, if a required predecessor was made more than three years before the disaster, it was retrieved from the full OSM history and imported into the `changes` database table to ensure completeness. These differences were computed in `Scripts/research_tools/change_differences/analyse_change_differences.py`.

Each edit change is paired with its predecessor and a structured `diff` is generated. A code snippet of this process can be found in Listing A.1. The `diff` captures a wide range of attribute metrics, including:

- **Tag changes**: Keys that were created, edited, or deleted.
- **Time between changes**: The number of days since the predecessor change.
- **Coordinate movement**: The geodesic distance between the current and previous coordinates of a node, or the centroids of a way.
- **Way geometry change**: For way elements, the nodes that were added or removed.
- **Visibility restoration**: Whether a previously deleted element was made visible.

The computed diffs were segmented by disaster period and aggregated to produce summary statistics, including the number of edits made to nodes and ways, as well as the proportion of edits involving tag changes, coordinate movements and way geometry modifications. The results of the analysis are used to answer RQ1.2 and RQ2.4.

RQ2.4: When are changes involving element tags made most frequently?
The final sub-research question focuses on the timing of contributors creating, editing and deleting tags. In RQ1.1, it was hypothesised that in the early stages of mapping response, contributors would focus on adding new content to the map, particularly in under-mapped areas. Building on this, we further hypothesise that post-disaster activity may shift towards enriching existing map elements with additional tag information or performing maintenance to ensure the long-term accuracy and completeness of the map.

To examine this behaviour, we apply the generalised methodology from RQ1.1, restricting the analysis to changes where tags are created, edited or deleted only. We then compare the weekly counts of tag-related changes to the overall change activity (RQ1.1), disaggregated by tag change type (tag `creates`, `edits` or `deletes`). Prophet modelling is applied to identify whether tag enrichment and maintenance deviate from pre-disaster trends, and the same visualisation charts are generated to interpret the results.

Chapter 6

Results and Analysis

In this chapter, we present the results of our investigation, organised by the research and sub-research questions. Each section includes quantitative analysis and is supported by visualisations and summary statistics across the pre, immediate, and post-disaster periods. The analysis was supported by an interactive dashboard[1], which enabled the rapid exploration and comparison of pre-generated visualisations across disasters and periods. The site was built with ReactJS and deployed on Netlify's cloud app platform.

6.1 RQ1: How does mapping activity vary across disasters during times of crisis?

RQ1.1: How does the change activity vary across disaster periods?

Figure 6.1a shows that the majority of changes (67.29%) are made immediately in response to disasters, of which 77.02% are `creates`, consistent with the hypothesis that contributors focus on intensively adding new content in the early stages of response. Mapping activity remains heightened in the post-disaster period, with a slightly increased proportion of changes being `edits` and `deletes`, although `creates` remain the dominant activity, as shown in Figure 6.1b.

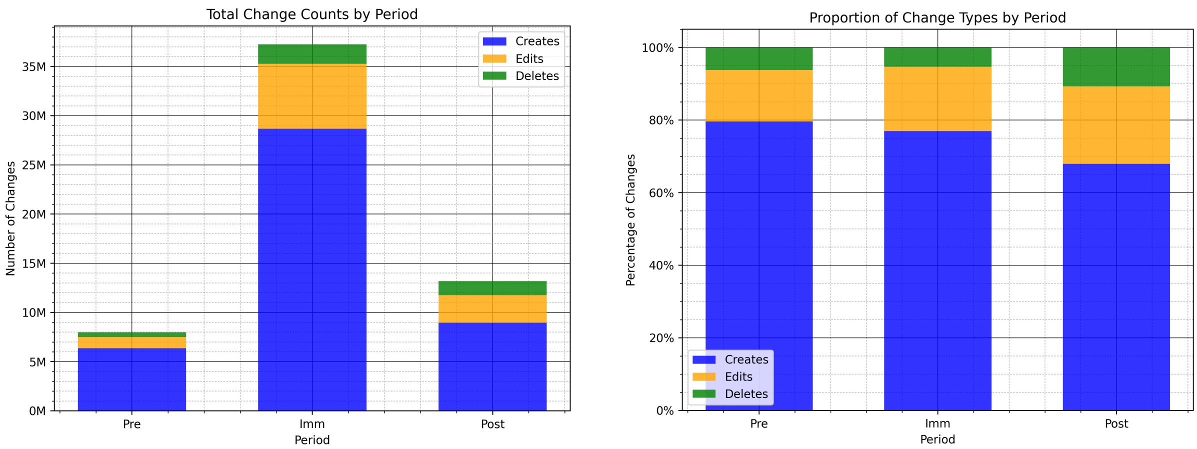

(a) Change counts by period and type for all disasters

(b) Proportion of changes by change type across all disasters

Figure 6.1: Aggregate change counts and change type proportions across periods.

[1] https://osm-times-of-crisis.netlify.app/

A logarithmic chart of the percentage differences in change counts for each disaster between the pre-disaster baselines and immediate and post periods is shown in Figure 6.2. While most experience substantial increases in activity, often exceeding 1000% in the immediate period, activity decreases in Sulawesi, Pakistan, and Emilia Romagna. For these storm and flood events, high and variable mapping activity pre-disaster surpasses that of the immediate period, giving the impression of a decline.

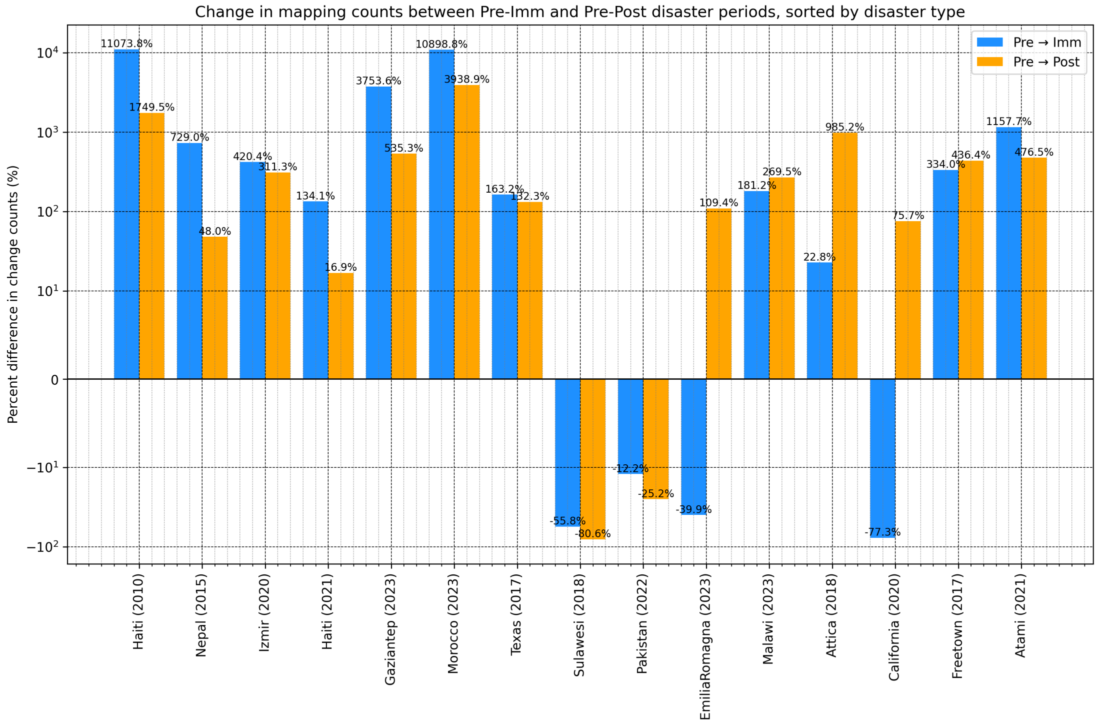

Figure 6.2: Logarithmic chart of percentage differences in change counts from the pre to the immediate, and pre to post disaster periods, ordered by disaster type.

Next, we examine the results of the temporal analysis of weekly change counts by change type, incorporating Prophet modelling to forecast pre-disaster activity had no disaster occurred. Figure A.1 in the appendix presents the charts for all 16 disasters, each showing the observed and forecasted weekly change counts by all change types; these can also be viewed on the interactive dashboard under the **Change Counting** tab[2], with graph style - 'Change count'. Across these, we observe three distinct behavioural patterns:

- A massive spike in the immediate period followed by rapid decay, illustrated by Figure 6.3. This pattern is common in disasters with minimal pre-disaster activity, where external mappers contribute intensively during the early response phase, but then move their focus elsewhere.
- No dominant spike that is higher than pre and post-disaster activity.

[2]https://osm-times-of-crisis.netlify.app/changeCounting?style=counts

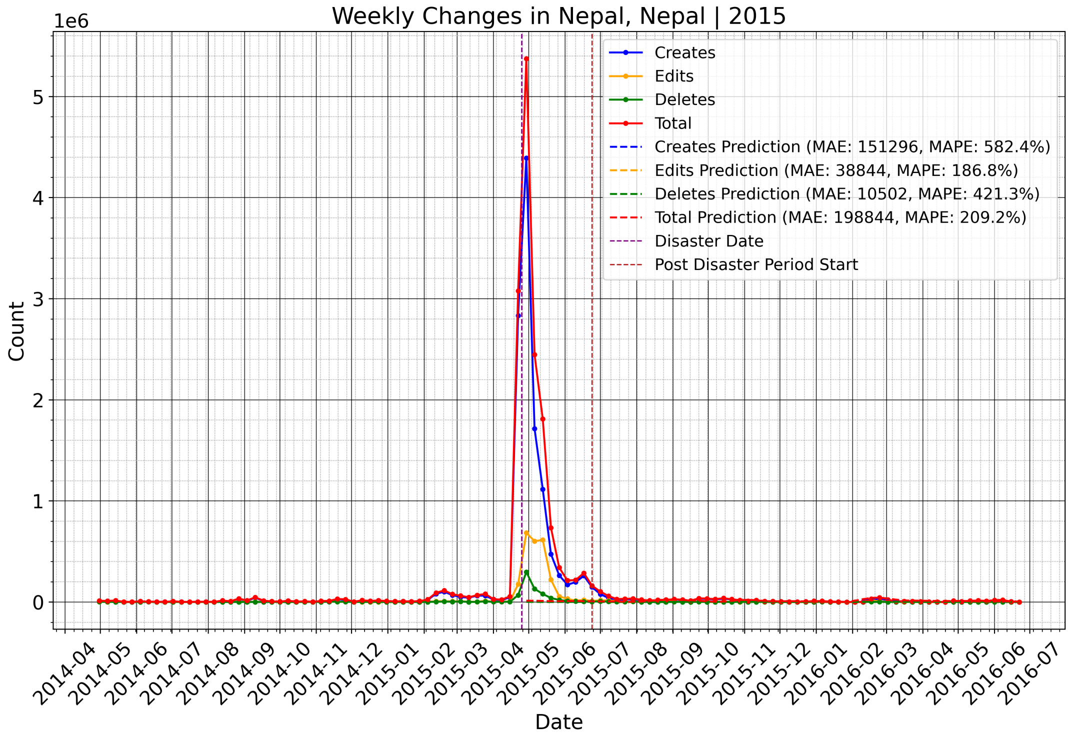

Figure 6.3: Massive peak: Weekly change counts with Prophet modelling for Nepal. Activity is dominated by creates, with a very large spike in the immediate period.

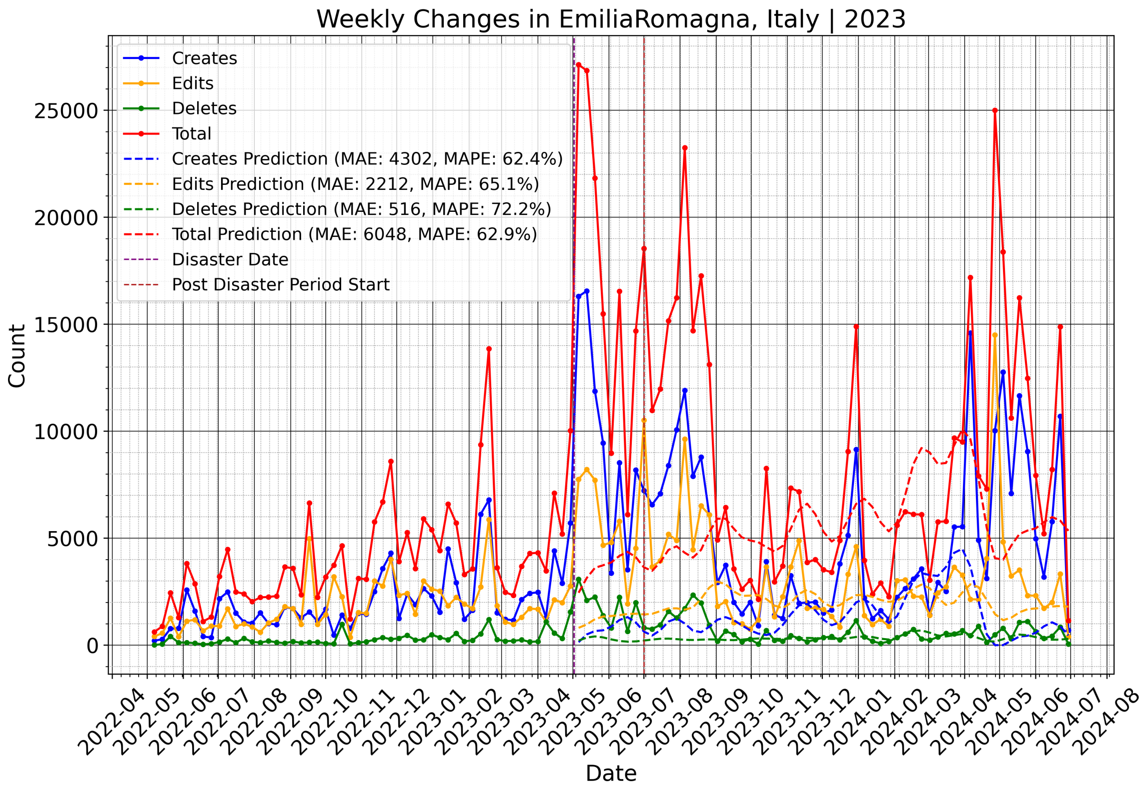

Figure 6.4: Background count and large peak: Weekly change counts with Prophet modelling for Emilia Romagna. A significant amount of activity can be observed, more evenly spread across change types.

- A consistent background level of mapping activity pre-disaster with a clear peak in the immediate period, as in Figure 6.4. This typically occurs in disasters affecting larger or more economically developed areas, where ongoing mapping activity is already present. Post-disaster activity remains variable, though higher than pre.

Table 6.1: Categorisation of disasters according to change count activity trends.

Category	**Disasters**
Massive Spike	Haiti (2010), Nepal, Atami, Izmir, Gaziantep, Freetown, Derna, Morocco
Background Count and Large Peak	Texas, Haiti (2021), Sulawesi, Pakistan, Emilia Romagna, Malawi
No dominant spike	California, Attica

`Create` changes (solid blue line) remain the dominant change across all disasters, particularly during the immediate period. While the majority of disasters conform to these three primary patterns, several exhibit notable differences. Sulawesi and Pakistan experienced very large spikes in mapping pre-disaster, which did not appear to be related to HOT mission-centric mapping or other disasters/significant events. The 2020 California wildfires emerged as a clear outlier throughout this investigation, with no distinct spike in activity during the immediate period, despite an active background count. Although the areas are rural and forested, we hypothesise that they were already well mapped before the disaster, reducing the need for crowdsourced contributions. The fires in Attica also did not result in significant increases in activity. Malawi, Pakistan and Izmir had long tails in mapping post-disaster, driven by sustained engagement from contributors.

Next, we examine the results of the Prophet modelling to assess how post-disaster activity compares to pre-disaster trends. The results are based on the `total` number of changes, rather than separated by change type, to reduce visual clutter, and focus exclusively on the post-disaster period. The results for all 16 disasters are included in the appendix (Figure A.2), and on the dashboard[3]. To illustrate the observed patterns, Figure 6.5 presents one example where post-disaster activity broadly follows the forecasted trend, and three examples that deviate in distinct ways. As summarised in Table 6.2, post-disaster activity tends to only partially align with forecasted trends, with the majority of disasters exhibiting activity above the predicted baseline. This is particularly clear in Emilia-Romagna, where post-disaster activity loosely follows the forecast after the immediate response, but later diverges with additional spikes not captured by the model, unlike Nepal, where activity more clearly returns to the pre-disaster baseline.

[3]https://osm-times-of-crisis.netlify.app/changeCounting?style=counts&type=total

Table 6.2: Categorisation of disasters according to alignment between post-disaster activity and Prophet forecasts.

Category	**Disasters**
Alignment	Haiti (2010), Nepal, Gaziantep
Partial Alignment	Haiti (2021), Morocco, Texas, Atami, Emilia-Romagna, Derna, Malawi, Attica, California
Large Prediction Variations	Pakistan, Freetown
Prophet Overestimates	Sulawesi, Izmir

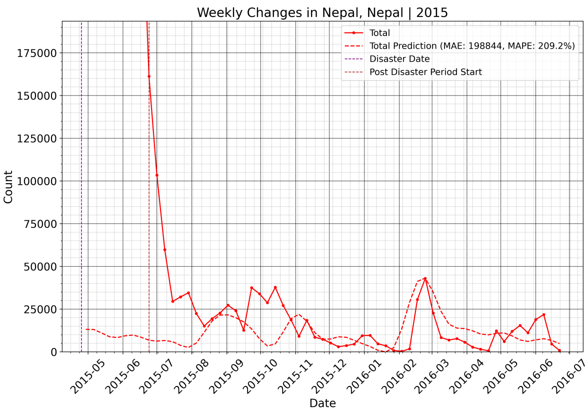

(a) Alignment: Nepal. Post-disaster activity largely follows pre-disaster predictions.

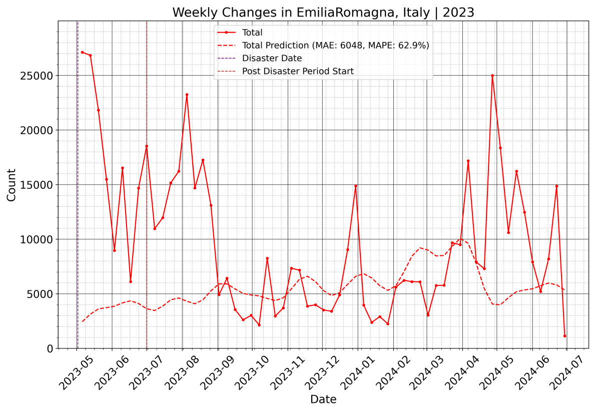

(b) Partial alignment: Emilia Romagna. Activity briefly aligns with the forecast before diverging.

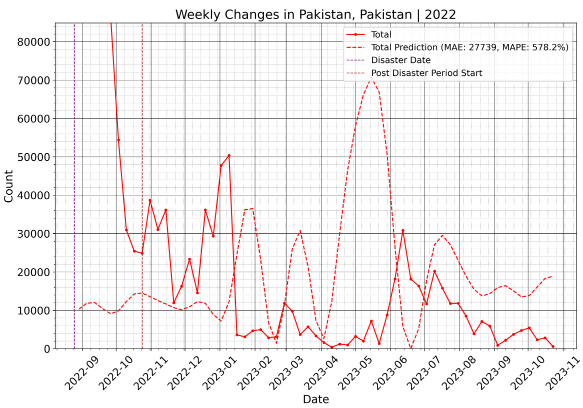

(c) Large prediction variations: Pakistan. High pre-disaster volatility is reflected in the trained time series.

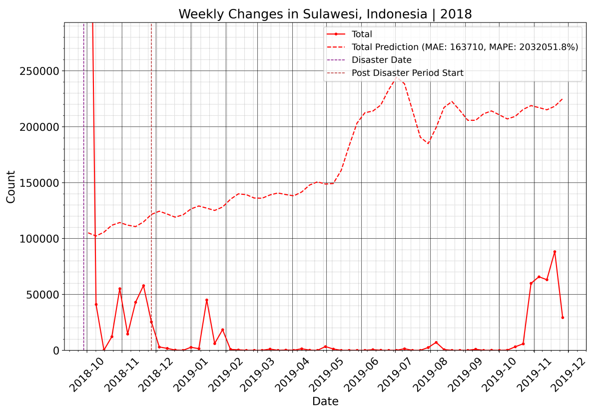

(d) Prophet overestimates: Sulawesi. Elevated pre-disaster activity leads the model to predict continued growth.

Figure 6.5: Examples of post-disaster mapping activity compared to Prophet forecasts.

While most disasters exhibit high percentage error (MAPE) values above 500%, those with consistent pre-disaster mapping activity, such as Texas and Emilia-Romagna, show significantly lower errors (70.7% and 58.6% respectively). It is clear that stable background activity, as found in more economically developed regions, results in more predictable post-disaster mapping. On the other hand, large variations in Prophet predictions for Pakistan and Freetown are a result of pre-disaster volatility, including irregular spikes, being reflected in the trained model.

Across nearly all disasters, the immediate period exhibits the highest mapping activity, often surpassing pre-disaster levels by over 1000%. However, the nature and duration of activity vary significantly by disaster. Earthquakes stand out for triggering the strongest responses, with most events in the massive spike category. In contrast, other disaster types are more inconsistent, particularly wildfires, which failed to trigger a noticeable increase in activity in both cases. Post-disaster mapping activity generally exceeds pre-disaster forecasts, particularly in areas with lower general mapping activity.

RQ1.2: How does the interval between consecutive changes vary?

Following the analysis of overall change activity, we now examine how frequently individual elements are re-edited by measuring the time between consecutive versions of the same element. Across the 16 disasters studied, three distinct temporal patterns emerged, illustrated in Figure 6.6. The full set of charts, Figure A.3, is provided in the appendix, on the dashboard[4], and Table 6.3 categorises each disaster.

Table 6.3: Categorisation of disasters according to post-disaster edit frequency trends.

Category	**Disasters**
Immediate Drop then Linear Growth	Haiti (2010), Nepal, Izmir, Gaziantep, Morocco, Sulawesi, Pakistan, Derna
Approximate Alignment	Izmir, Haiti (2021), Texas, Pakistan, Attica, California, Freetown, Atami
Generally Below Forecast	Emilia-Romagna, Malawi

Pakistan and Izmir exhibit a mixed pattern, with an initial drop and linear growth in intervals, followed by a return to pre-disaster trends, which corresponds with their long tails of sustained post-disaster activity observed in both cases.

[4] https://osm-times-of-crisis.netlify.app/changeCounting?style=avg_days_between_edits

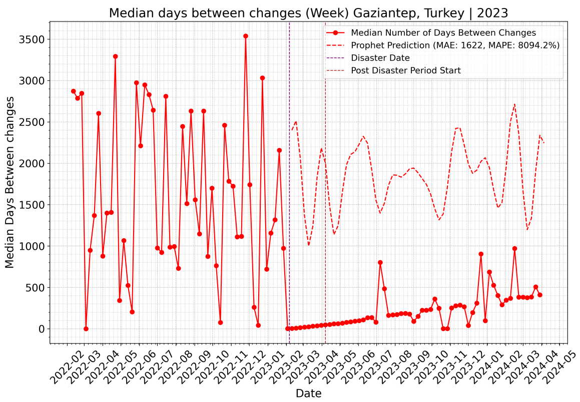

(a) Immediate drop then linear growth: Gaziantep. Post-disaster, the interval between edits drops to zero, then increases linearly, remaining below the forecast. Newly created elements are primarily edited in the immediate aftermath, rather than existing ones: enrichment-focused activity.

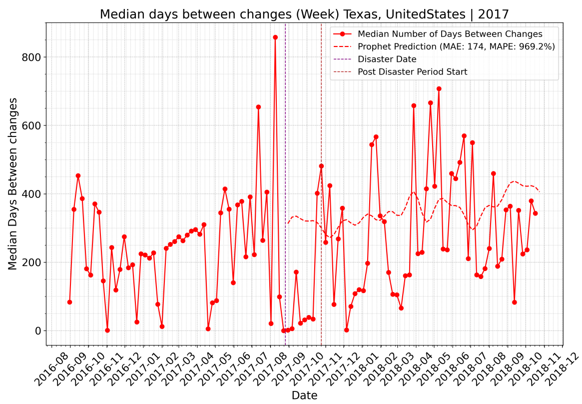

(b) Approximate alignment: Texas. Post-disaster edit intervals approximately follow the forecasted trend, with variation across individual weeks. This suggests a mix of phases of maintenance of existing elements and enrichment of newly created elements.

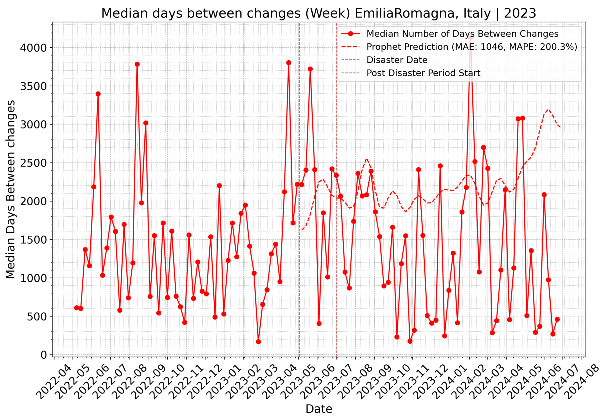

(c) Generally below forecast: Emilia-Romagna. Post-disaster edit intervals are largely shorter than predicted, with occasional weeks aligning with or exceeding the forecast. This suggests a sustained period of frequent edits to newly added and existing elements, indicating ongoing engagement beyond initial response efforts.

Figure 6.6: Weekly edit frequency patterns illustrating the three distinct post-disaster behaviours observed across disasters.

Overall, the interval between consecutive edits appears to align less with disaster type and more with the underlying completeness of mapping in the affected region. As noted by Herfort et al. [3], regions such as North America and Europe tend to be better mapped, whereas lower completeness is common elsewhere. Disasters in less-mapped regions more frequently exhibit the "Immediate Drop then Linear Growth" pattern, reflecting rapid enrichment of newly created elements.

RQ1.3: How does the spatial distribution of activity vary in disaster areas? Following the temporal analysis of contributor activity, we now examine how changes are spatially distributed across disaster areas. The primary goal is to visualise where changes are concentrated, whether in rural or urban areas, and whether these changes align with key infrastructure. Binned choropleth maps of overall change counts and percentage differences (pre→imm, pre→post) were computed, shown in Figure 6.7, and visible on the interactive dashboard under the **Change Density Mapping** tab[5].

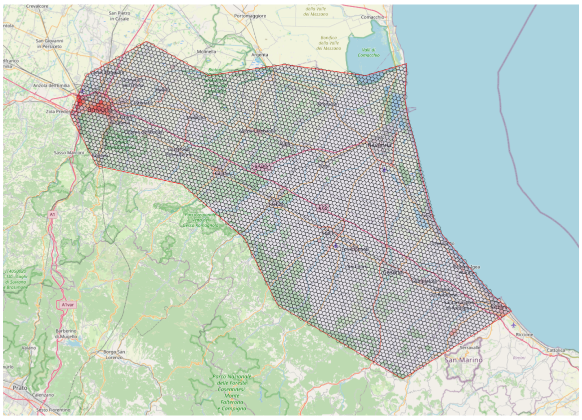

(a) Distribution of raw change counts across Emilia-Romagna, with greater count areas in a darker red. Activity is visible in urban centres, such as Bologna to the west, but is less apparent in other areas.

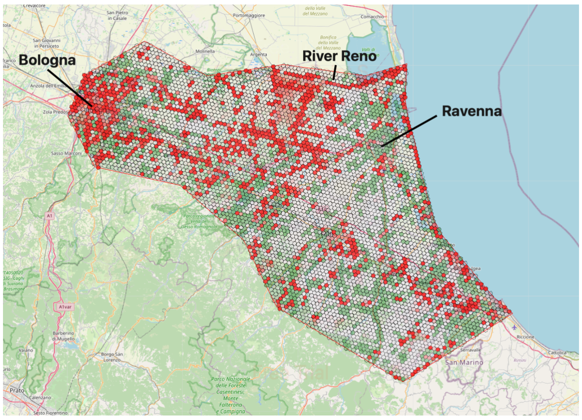

(b) Percentage differences in change counts from pre to post-disaster period. Red indicates an increase in changes, and green a decrease. The distribution of activity is much clearer.

Figure 6.7: Change density mapping for Emilia Romagna, comparing raw change counts and percentage differences between the pre, immediate, and post-disaster periods.

From both maps, it is clear that changes are concentrated around urban areas, although Figure 6.7b better highlights increased mapping activity across the whole region. Activity appears to follow the River Reno to the north, which broke its banks in the floods [48], as well as the city of Bologna, and the cities between it and Ravenna to the east.

Maps comparing pre-disaster activity with the immediate and post-disaster periods are shown in Figures 6.8 to 6.11, with each disaster type represented by a set of maps. Spatial distribution of activity reveals distinct patterns across disaster types. Earthquakes attract widespread attention in the immediate period, followed by a decrease post-disaster. In contrast, storms/floods receive concentrated activity initially, followed by a broader spread. Landslides see concentrated activity in the most affected areas during the immediate period, with subsequent broader mapping post-disaster.

[5] https://osm-times-of-crisis.netlify.app/changeDensityMapping?style=percent&res=8

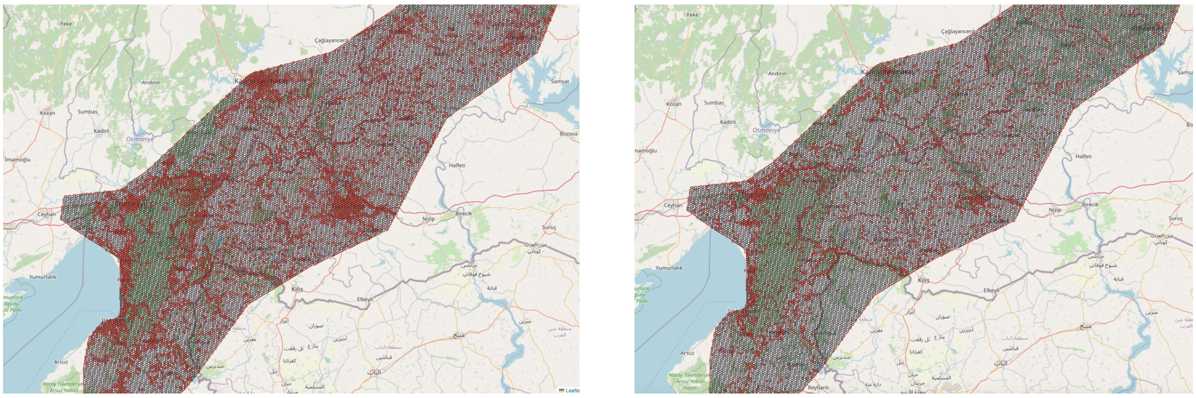

(a) Earthquake: Gaziantep Pre→Imm (b) Earthquake: Gaziantep Pre→Post

Figure 6.8: Percent change in mapping for the Gaziantep earthquake (2023). Mapping is more widespread in the immediate period, particularly in urban areas. Post-disaster period, mapping impacts a smaller total area, but with a greater focus on rural locations.

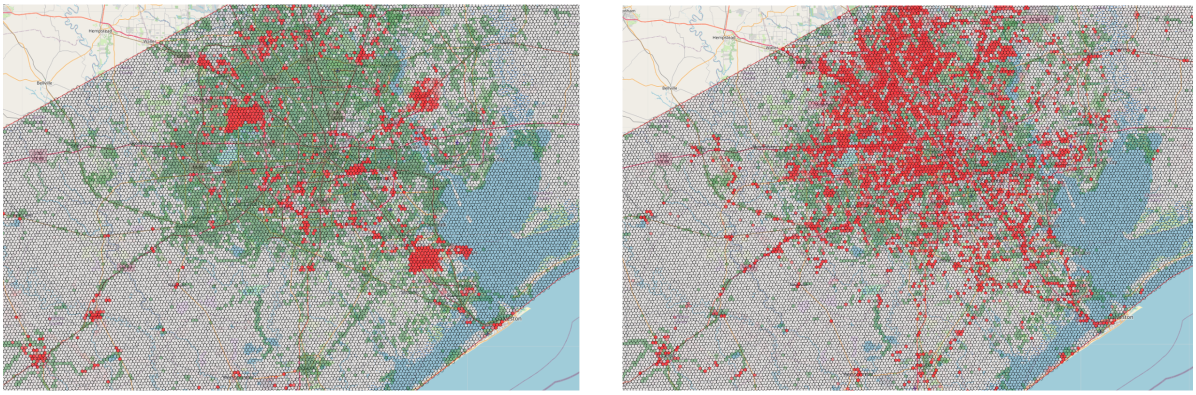

(a) Storm/Flood: Texas, Houston Pre→Imm (b) Storm/Flood: Texas, Houston Pre→Post

Figure 6.9: Percent change in mapping for Texas Hurricane Harvey (2017). Immediate activity is concentrated in specific areas of the city, particularly around waterways and the coast, while post-disaster activity covers a wider area, also following highways.

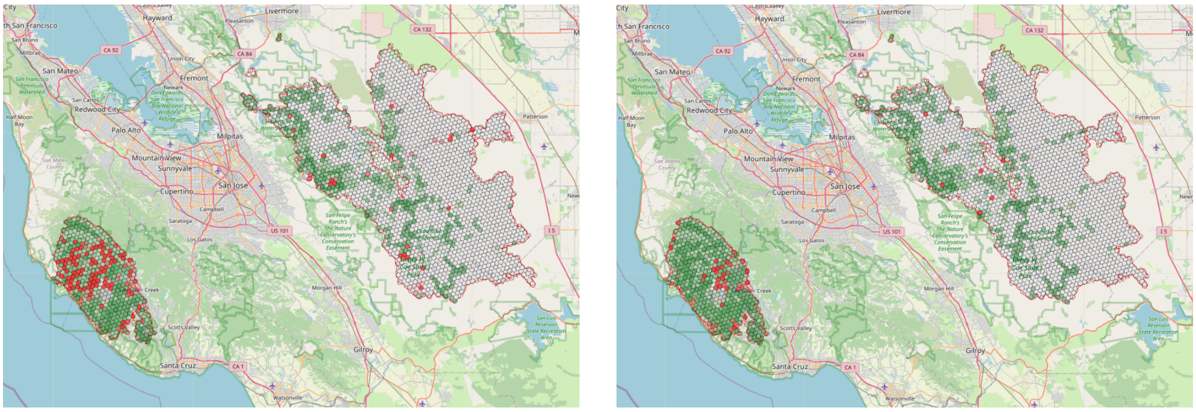

(a) Wildfire: California, Bay Area Pre→Imm (b) Wildfire: California, Bay Area Pre→Post

Figure 6.10: Percent change in mapping for the California wildfires (2020). Contrary to the findings of RQ1.1, we observe greater activity in the immediate period, affecting a greater area. The activity is limited to forested areas.

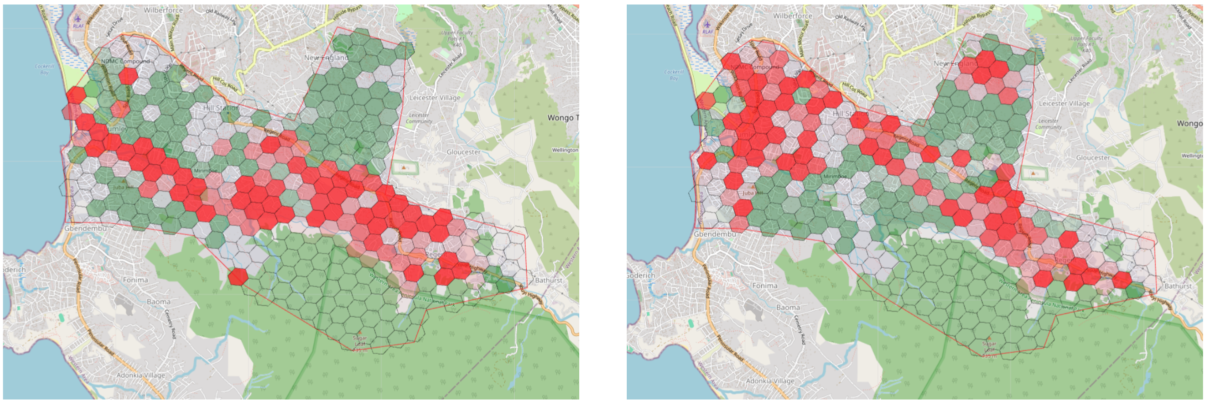

(a) Landslide: Freetown Pre→Imm (b) Landslide: Freetown Pre→Post

Figure 6.11: Percent change in mapping for the Freetown landslide (2017). Activity is concentrated in the most affected areas during the immediate period [49], and more widespread across the town post-disaster, highlighting the shift in focus of relief efforts.

Unlike other disasters, the clearly defined boundaries of the California wildfires were used [60], resulting in 133 disjoint areas of activity, located in rural, wooded regions. Activity was more widespread in the immediate period, contrasting with the lack of a significant spike in the RQ1.1 charts. This could be due to the inclusion of only two fire boundaries in Figure 6.10, with other affected areas not shown exhibiting lower levels of activity. The spatial distribution of changes across disasters does not appear to be significantly influenced by pre-existing mapping activity or the geographic region, contrasting with the overall activity observed in RQ1.1.

To quantitatively compare spatial distributions, the Gini coefficient was computed for each disaster and period. The coefficient measures inequality, with 1 indicating that all changes are made within a single spatial unit, and 0 all changes are evenly spread across the disaster area. The percentage differences in Gini from the pre to the immediate and post-disaster periods are presented in Figures 6.12 and 6.13.

We observe that in the immediate period, variations in the concentration of changes are more pronounced, with both larger increases and decreases in inequality compared to post-disaster: mapping activity becomes more concentrated in some areas, and more dispersed in others. As expected, post-disaster activity is more similar to pre, after the heightened activity in the immediate period. However, the variations in inequality appear to be independent of disaster type, geographic region, area size and urbanisation, therefore, other more localised factors may influence how the concentration of changes varies. Investigating the inequality by separate change types yielded similar findings.

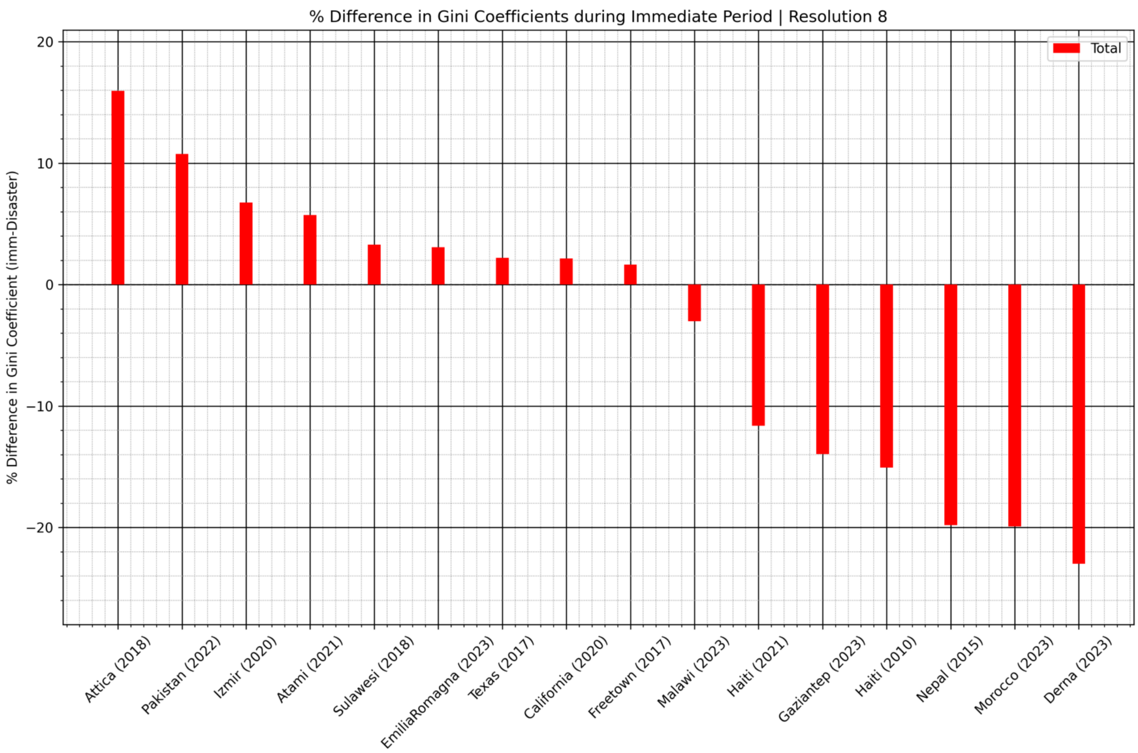

Figure 6.12: Percentage difference in Gini coefficient, Pre→Imm. Attica has the greatest increase in inequality, with activity becoming more concentrated in the immediate period. Meanwhile, Derna, Morocco and Nepal become more equally spread.

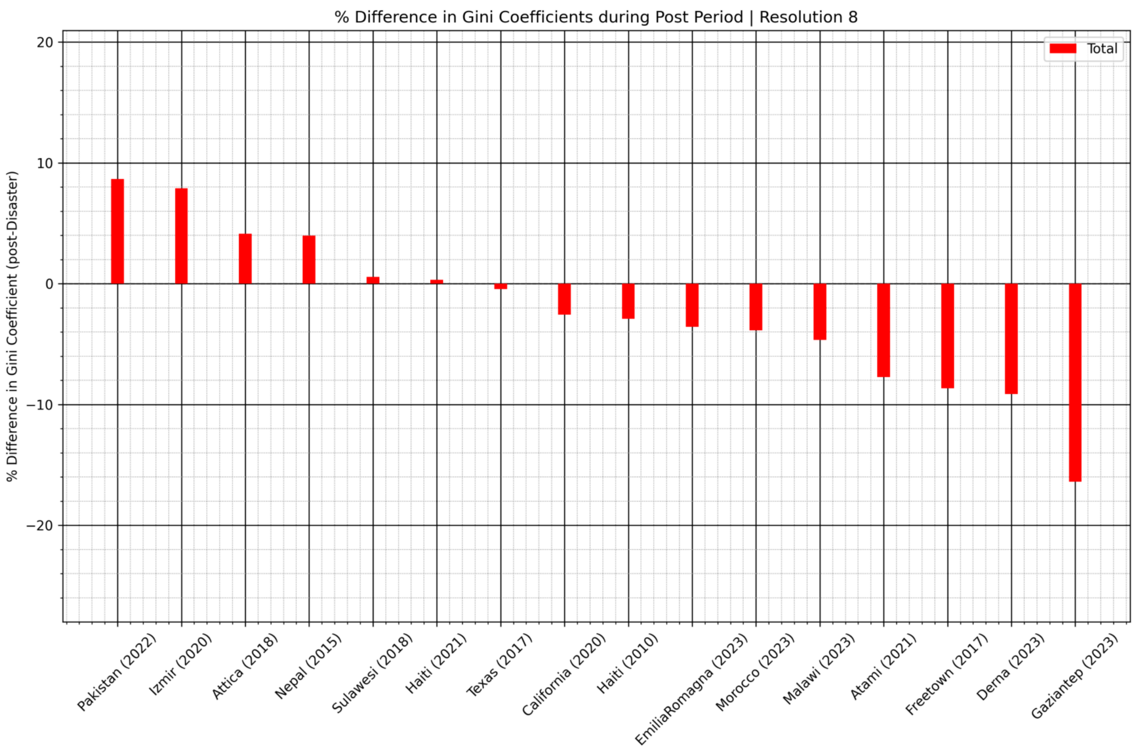

Figure 6.13: Percentage difference in Gini coefficient Pre→Post. The magnitude of change is generally lower compared to the immediate period, suggesting a return to a more standard distribution of mapping activity post-disaster.

6.2 RQ2: How does the content contributed to the map evolve throughout disasters?

Having examined how mapping activity varies across disaster periods to address RQ1, we shift focus toward the content and nature of contributions to answer RQ2.

RQ2.1: Which geographic features are most commonly changed?
Following the selection of the four key geographic features of interest: `building`, `highway`, `amenity`, and `leisure`, the frequency of changes across all disasters and periods, as well as the proportion of changes within each disaster phase, were computed, as on the dashboard[6], and shown in Table 6.4. To clearly illustrate the relative importance of these features across periods, a stacked area chart in Figure 6.14 is also presented.

Index	Feature	Count	All periods (%)	Pre (%)	Imm (%)	Post (%)
0	`building`	6,475,267	10.267	9.654	16.823	18.438
1	`highway`	2,357,930	3.740	6.915	3.348	9.619
9	`amenity`	170,605	0.271	0.340	0.123	0.884
49	`leisure`	39,788	0.063	0.113	0.043	0.227

Table 6.4: Number of changes and percentage share where a changed element included a tag corresponding to each selected geographic feature key, across disasters and periods. Index refers to the key's ranking among the top 50 keys across all changes.

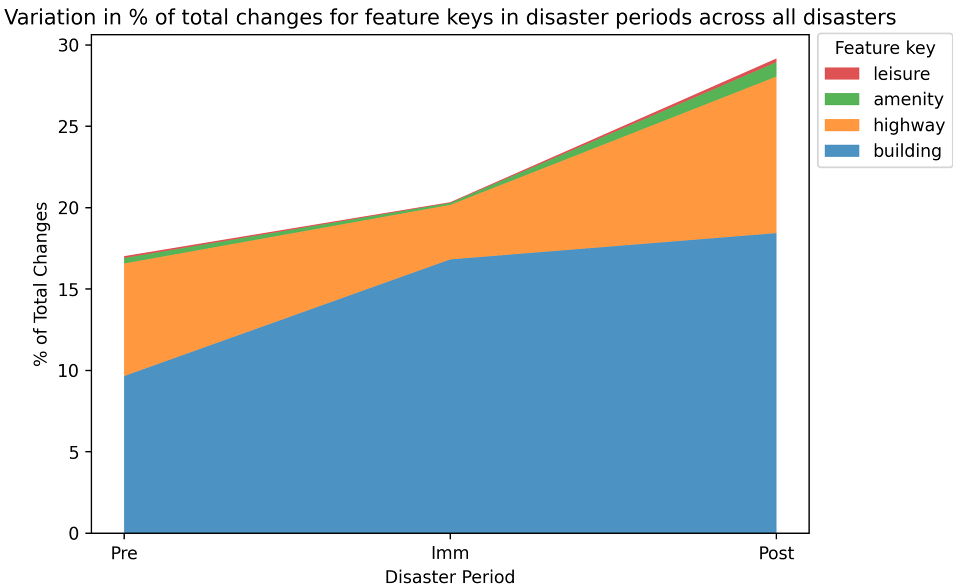

Figure 6.14: Stacked area chart of relative proportions of geographic features by period.

[6] https://osm-times-of-crisis.netlify.app/tagKeysValues

As expected, `building` and `highway` were by far the most frequently changed features across periods, accounting for nearly 30% of all tagged changes in the post-disaster period, and growing in dominance compared to pre-disaster. Interestingly, the dominance of `highway` contracts in the immediate period before recovering, suggesting a temporary shift away from mapping roads.

Although `amenity` and `leisure` were less commonly changed features, meaningful conclusions can still be drawn about the growing importance of `amenity` in the post-disaster period, as evidenced by the distributions presented in Table 6.4. Both contracted substantially during the immediate period, by approximately 2.7×, relative to their pre-disaster proportions. However, post-disaster, `amenity` recovered more strongly, increasing by a factor of 7.2×, whereas `leisure` grew by only 5.3×. Overall, this indicates that `amenity` features became more important in the aftermath of disasters, reflecting a shift in contributor focus toward mapping essential services critical for recovery efforts, rather than non-essential `leisure` areas.

Across disaster types and geographic regions, no consistent pattern was observed in the geographic features being mapped. However, Emilia-Romagna stood out, with higher proportions of `highway`, `amenity`, and `leisure` changes than most other disasters. This likely reflects its well-mapped, developed status, although high urbanisation alone did not guarantee similar patterns elsewhere.

RQ2.2: Among the specified geographic features, which feature types are most commonly changed?

Having examined how frequently the four geographic features were changed across disasters and periods, the investigation next focuses on the specific feature types most commonly mapped within each feature. As a reminder, these correspond to the `value` in `key=value` tag pairs. Following a similar approach to that used in RQ2.1, the top 12 values for each key in the pre-disaster period were identified. Their usage proportions across the disaster periods were then computed and are presented in Figure 6.15. The generic value `building=yes` was excluded from the analysis, as it overwhelmingly dominated the other `building` feature types.

Significant growth in `building=house` was observed during the immediate period, likely due to the creation of clear individual house elements that were previously unmapped, while `building=residential` becomes more prominent post-disaster, signifying the subsequent enrichment of the map with greater detail about the nature of buildings after the urgent response phase. The increased mapping of `highway=path` and

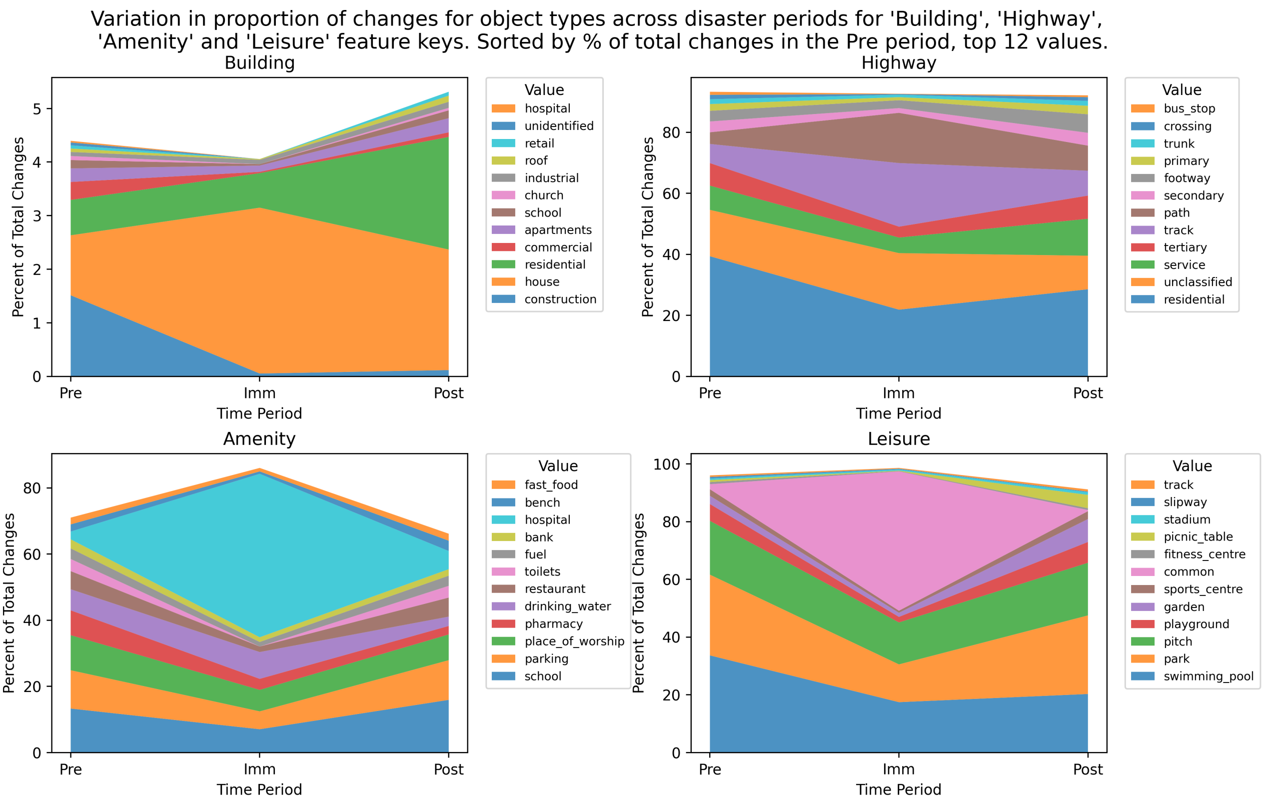

Figure 6.15: Stacked area charts of relative proportions of the top 12 tag values from the pre-disaster period, for the four geographic features by disaster period.

`highway=track` during the immediate period likely stems from the need to establish access routes for emergency responders, particularly in rural or remote areas where paths and minor tracks may not have previously been mapped.

The immediate period shows a major increase in mapping of `amenity=hospital`, indicating a strong focus on updating and adding emergency medical infrastructure over other amenities. These were particularly frequently mapped in response to the Haiti 2010 and Nepal 2015 earthquakes. Interestingly, `amenity=pharmacy` decreases from 7.51% to 3.34% of amenity-related changes during the immediate response period, suggesting that pharmacies were a lower priority relative to hospitals. Mapping `amenity=drinking_water` also increased during the immediate period, driven by the need to identify safe water sources following infrastructure damage or supply contamination. This was particularly evident in the Americas, Emilia Romagna, and Nepal, as shown in Figure 6.16, encompassing both regions where tap water is typically reliable (United States, Emilia-Romagna), and where tap water is not safe to drink.

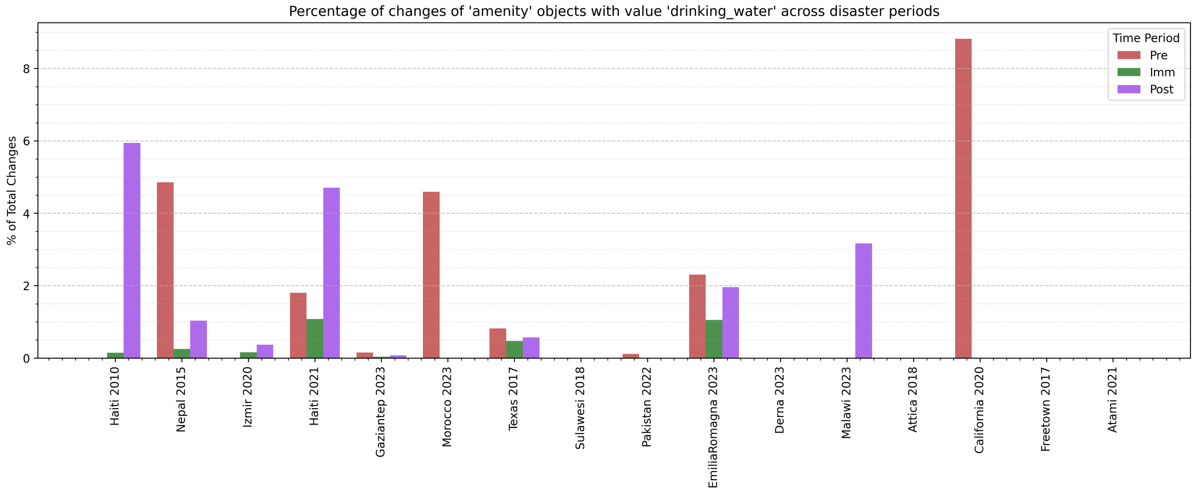

Figure 6.16: Percentage of changes involving `amenity=drinking_water` across periods for each disaster.

Finally, the growth of `leisure=common` is observed, primarily driven by them being marked for use as possible helicopter landing sites in Nepal. Relatively few changes were made to `leisure` elements across disasters, affirming that recreational features were not a major focus during disaster response efforts. Overall, other than the patterns highlighted, disaster type and geographic region did not appear to have a strong or consistent influence on the feature types mapped across disasters.

To compare the consistency of mapping priorities across disaster periods, the Kendall rank correlation coefficient was computed based on the ordering of the most frequently mapped feature types within each geographic feature key across the periods. A coefficient of 1 indicates perfect agreement across periods, and -1 perfect disagreement. Table 6.5 reports the median of the coefficients computed across the period pairs: pre–immediate, pre–post and immediate–post. This median reflects the typical consistency in the ordering of feature types for each geographic feature across periods. In addition, the minimum, maximum, and interquartile range (IQR) are reported for each feature to capture the spread and variability of the correlation scores.

Feature Key	**Min τ**	**Max τ**	**Median τ**	**IQR**
`amenity`	0.000	0.756	0.548	0.220
`building`	0.242	1.000	0.636	0.426
`highway`	0.236	0.909	0.716	0.173
`leisure`	-1.000	1.000	0.536	0.532

Table 6.5: Kendall's τ statistics for feature type ordering consistency across periods.

Among the features analysed, `highway` demonstrated the most stable mapping priorities, with the highest median Kendall's τ and the narrowest IQR, while `leisure` had the lowest median and the widest IQR. The lower consistency observed for `leisure` is likely due to the relatively small number of changes made to `leisure`-related elements across disasters. Most disasters are clustered above τ values of 0.5 for `building` and `highway`, with a wider cluster generally above 0.25 for `amenity` and `leisure`. Mapping in Sulawesi (2018 earthquake and tsunami) exhibited the least consistent value orderings across disaster periods, while the other disasters showed moderate variability in mapping priorities.

RQ2.3: What attributes of elements do contributors change in their edits? Following the spatio-temporal analysis of activity and mapping priorities across geographic features and feature types, this sub-research question examines the specific attributes of map elements that contributors modify. The change differences methodology was applied, comparing each edited element with its predecessor to identify the nature of modifications. Across all disasters and periods, 11,006,905 edit changes were made. The majority, 87.43%, affected node elements, and the remaining 12.57% were made to ways. This is expected, as nodes far outnumber ways in OSM; however, many node edits also contribute to modifications of way geometries, as ways are structured from nodes.

Only 8.43% of edits involved tag changes, including the creation, modification, or deletion of tags. The majority of edits, 78.86%, moved nodes, with a mean distance of 3.68 meters. Additionally, 7.26% of edits modified way geometries by adding, reordering, or removing nodes. 0.13% of edits were delete reversions, restoring previously deleted elements to make them visible on the map once again. Such edits typically occur in response to a contributor mistakenly deleting an element or to repair instances of vandalism. The remaining 5.32% of edits are observed in the change data but do not correspond to any of the categories listed. These edits may involve maintenance of other element attributes not captured in the change difference analysis, requiring further investigation.

We hypothesise that the relatively small proportion of tag edits, compared to the large number of node movements, results from the fact that moving a node is a much simpler task that can be performed quickly by any contributor, either remotely using satellite imagery or with local knowledge. It is also easy to move many nodes efficiently, especially when editing ways. In contrast, editing tags to enrich or maintain elements requires local knowledge and greater care to ensure accuracy. Tag edits are rarer and considered more valuable than simple positional updates. Given their importance, we specifically investigate the patterns of tag change activity in RQ2.4.

RQ2.4: When are changes involving element tags made most frequently?
The final sub-research question investigates when contributors `create`, `edit` and `delete` tags across periods. In answering RQ1.1, we observed that element `creates` dominated activity, particularly during the immediate period, as shown in Figure 6.1. In contrast, when focusing only on tag-related changes, a different pattern emerges, presented by Figure 6.17.

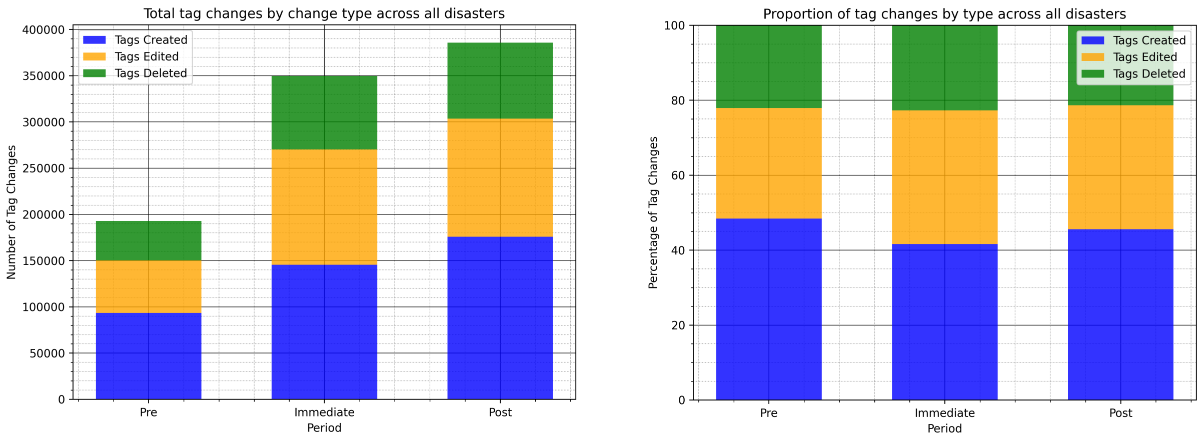

(a) Change counts by period and type for all disasters

(b) Proportion of tag changes by tag change type across all disasters

Figure 6.17: Aggregate change counts and change type proportions across periods.

Although overall mapping activity declines after the immediate period (Figure 6.1a), tag changes become more prevalent post-disaster (Figure 6.17a), making up a greater proportion of edits compared to during the immediate response, suggesting that enrichment and maintenance of existing map elements becomes more of a focus once the immediate threat and response to the disaster has passed. Furthermore, compared to overall changes, the distribution of tag change types is much more even across periods, with a greater proportion of `deletes` and `edits`, rather than being dominated by `creates`. This reflects the fact that while the underlying map elements often remain unchanged, the information that describes them is more likely to be updated as situations change.

To further examine temporal patterns, the weekly tag change counts were compared to the overall weekly change counts by applying the generalised time-series methodology from RQ1.1. The complete set of weekly tag change charts for all disasters is included in the appendix (Figure A.4) and on the dashboard[7]. The focus on tag changes during the post-disaster period is best exemplified in Figure 6.18, in which the change counts and tag change counts over weekly intervals are compared.

[7] https://osm-times-of-crisis.netlify.app/changeCounting?style=tag_changes

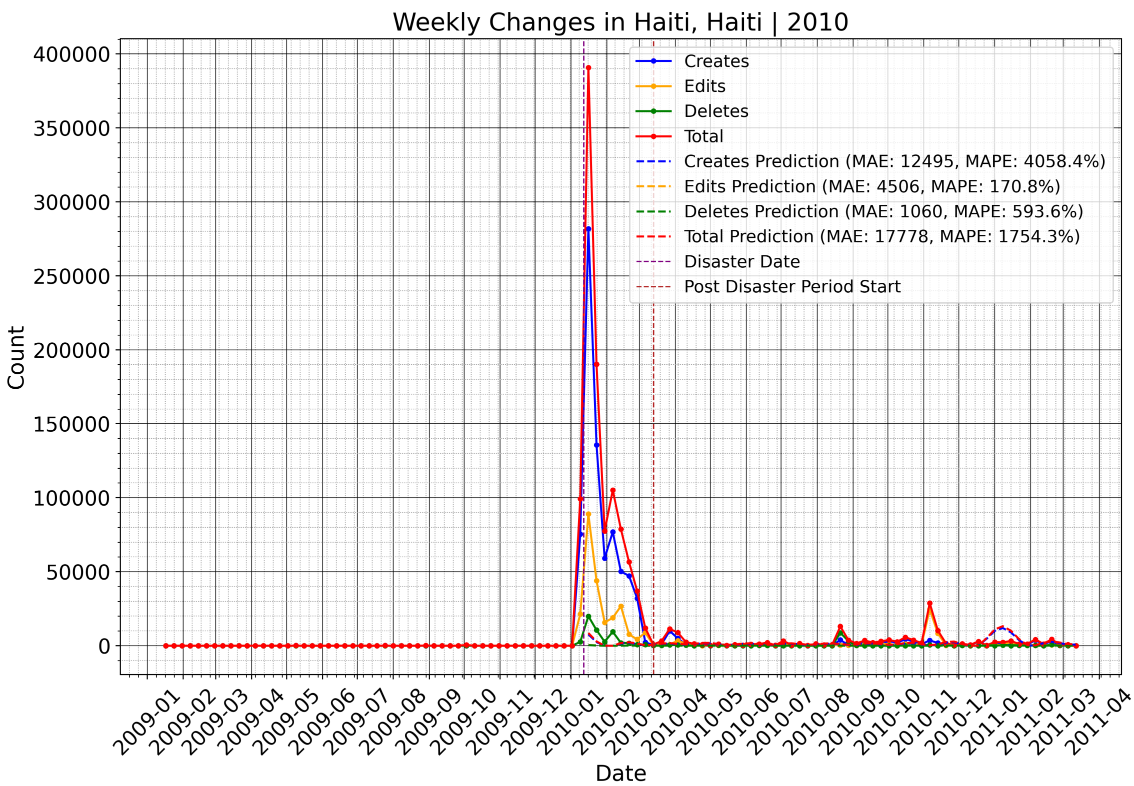

(a) Overall weekly change counts for Haiti 2010.

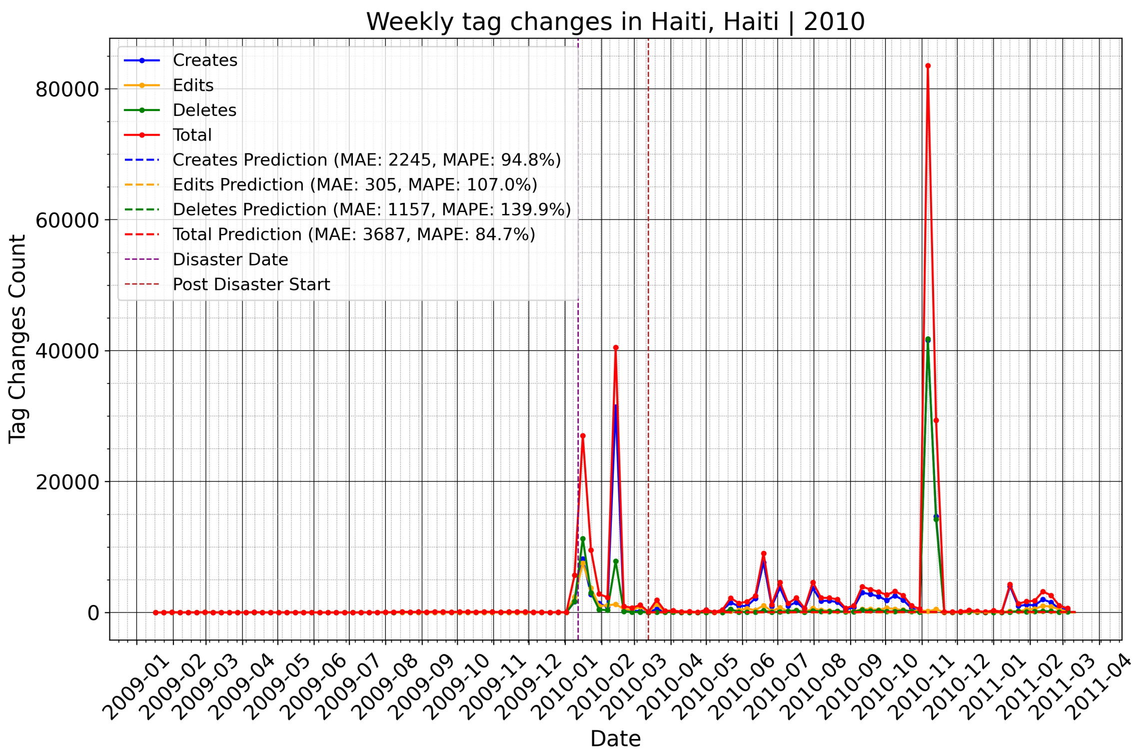

(b) Weekly tag change counts for Haiti 2010.

Figure 6.18: Comparison of weekly overall change and tag change counts for Haiti 2010.

In contrast to overall changes (Figure 6.18a), which peak sharply during the immediate response, post-disaster tag changes (Figure 6.18b) align with smaller bumps in overall activity, including a significant spike in November 2010, while the immediate period spike is less pronounced.

Across the disasters studied, similar patterns of tag changing behaviour were observed, where post-disaster tag activity was frequently higher than any peak during the immediate period. During the immediate response period, contributors primarily focused on creating new elements, while post-disaster, there is a greater emphasis on enrichment and maintenance. Deletion of tags is much more prevalent than deletion of elements, illustrating the correction and refinement of tags created during the immediate response period, many of which may not be necessary. It is possible that less experienced contributors initially created tags, which were subsequently edited or removed by more experienced mappers to improve map accuracy and consistency. The Prophet forecasts for tag changes tended to align much more closely with the actual results than for overall changes, typically with an MAPE below 200% compared to around 500% for overall activity, highlighting greater stability of tag editing behaviour across periods.

Interestingly, while California did not experience a major visible spike in overall mapping activity during the immediate period, clear spikes in tag creations and deletions were recorded, suggesting that the area was already well mapped and primarily required information enrichment. Although tag changes tended to become more prominent in the post-disaster period across most disasters, no specific pattern was observed relating to disaster type or geographic region.

Chapter 7
Conclusion

7.1 Summary

This project investigates how volunteer contributions to OSM evolve throughout disasters, addressing a gap in large-scale analyses of crowdmapping behaviour. Analysing over 51 million changes across 16 disasters, surges in mapping activity were observed during the immediate response periods of disasters, driven by the creation of new elements. While activity declined after the immediate period, it consistently remained at or above pre-disaster levels. The strongest spikes were typically associated with regions with lower pre-disaster activity, often corresponding to differences in geographic region and development. The intervals between successive re-edits to elements varied across disasters, with patterns linked to existing mapping completeness, rather than disaster type, as less-mapped regions exhibited more rapid post-disaster editing. However, disaster type was found to influence the spatial distribution of changes: earthquakes exhibited widespread mapping during the immediate period with more concentrated activity post-disaster, whereas storm/flood mapping was initially concentrated before becoming more dispersed across affected regions.

Analysis of the geographic features affirmed that `building` and `highway` were the most commonly mapped features, consistent with previous literature, with houses and paths increasingly mapped in the immediate period. `Amenity` features also grew in significance during the immediate period, particularly hospitals and drinking water sources, highlighting the increased focus on mapping essential services for disaster response and recovery. Most edit changes involved moving nodes, with tag changes representing a small proportion of overall edits; however, tag changes became more prevalent post-disaster, reflecting a shift towards the maintenance and enrichment of existing map elements.

Overall, while we find that disaster type influences certain aspects of mapping behaviour, the pre-existing completeness of the map is equally, if not more important. This highlights the value of pre-emptive mission-centric mapping efforts led by organisations such as HOT, as improving baseline mapping completeness in disaster-prone regions ensures that users have immediate access to critical geographic information, while volunteers can focus on the enrichment and maintenance of the map to reflect changes in response to the disaster.

7.2 Critical Evaluation and Future Work

Several limitations affected this project, which we acknowledge for transparency and to suggest possible future improvements and areas of research.

The first concerns the heterogeneity of the crises studied. The characteristics of the disasters were widely varied, including geographic area, socio-economic status, the extent of the area affected, human impact, and other factors. This diversity made it more challenging to identify concrete patterns, as each disaster presented distinct features. Differences in baseline mapping completeness, local population density, and computer literacy likely influenced volunteer responses, alongside the nature of the disaster itself. This is an inherent difficulty when investigating a varied set of disasters; future work could either focus more closely on a small number of comparable crises, as in previous research, or expand to include an even larger number of disasters.

In some disasters, the definition of geographic boundaries was sometimes too restrictive, particularly in urban areas. This may have excluded relevant mapping activity from the analysis, limiting the spatial distribution findings. Larger boundaries could have been defined, incorporating both the disaster areas and more surrounding regions. Additionally, we note that for the Gini coefficient computations, resolution 8 Uber H3 spatial units were consistently used across all disasters. However, for small urban areas, these were too coarse, often covering most of the disaster area, particularly in Atami. Resolution 9 spatial units were also explored, but are not presented in the report, as inequalities were generally similar.

Another limitation was the imperfect bulk import filtering process applied in this investigation. Although the concept has been discussed in previous literature, a reliable method for classifying changes as bulk imports or human-made changes has not been well defined. While an attempt was made to filter these changes, we could not guarantee that there weren't more non-human changes, especially as we erred on the side of caution to avoid accidentally removing valid contributions. Future work could focus on improving the identification of bulk imports, possibly by investigating their characteristics more closely and identifying their features, or training machine learning models to classify changes as human-made or automated.

Prophet modelling of change counts was generally reliable for most disasters, however, the model struggled with disasters characterised by pre-disaster spikes in activity and otherwise low background levels, often predicting similarly large spikes post-disaster. The model had to have its outputs clipped at zero, as predicted counts would occasionally be negative, despite parameter tuning and adjustment of the model. Future work could explore alternative time series forecasting tools, such as ARIMA, or strategies to smooth pre-disaster activity for more accurate predictions.

When applying the change differences methodology, we found that 5.3% of changes could not be accounted for in the analysis. This suggests that there were changes whose differences were not captured, indicating a potential limitation in the methodology, as it is unlikely that these modifications would occur outside the scope of the OSM change history data.

Future research could examine the spatial distribution of changes in greater detail, particularly investigating how proximity to the centre of different types of natural disasters, such as the distance to an earthquake epicentre or the path of a hurricane, influences the density and intensity of mapping activity. Full spatial time-series analyses could also be conducted to capture the evolution of the spread of changes over time, rather than aggregating the activity across the disaster periods. Finally, building on this project's findings that post-disaster periods see significant tag enrichment and maintenance activity, future studies could investigate the extent to which these edits improve the quality and completeness of the map, which specific tags are changed, and who the volunteers are responsible for post-disaster activity.

Bibliography

[1] P. Mooney, M. Minghini *et al.*, "A review of openstreetmap data," *Mapping and the citizen sensor*, pp. 37–59, 2017. [Online]. Available: https://doi.org/10.5334/bbf.c

[2] Wikipedia Contributors, "Size of Wikipedia," https://en.wikipedia.org/wiki/Wikipedia:Size_of_Wikipedia, 2025, Accessed: 2025-04-18.

[3] B. Herfort, S. Lautenbach, J. Porto de Albuquerque, J. Anderson, and A. Zipf, "A spatio-temporal analysis investigating completeness and inequalities of global urban building data in OpenStreetMap," *Nature Communications*, vol. 14, no. 1, p. 3985, 2023. [Online]. Available: https://doi.org/10.1038/s41467-023-39698-6

[4] M. Haklay and P. Weber, "Openstreetmap: User-generated street maps," *IEEE Pervasive computing*, vol. 7, no. 4, pp. 12–18, 2008. [Online]. Available: https://doi.org/10.1109/MPRV.2008.80

[5] OpenStreetMap Contributors, "Openstreetmap about page," https://www.openstreetmap.org/about, 2025, Accessed: 2025-04-14.

[6] Saul Alber, "[Openstreetmap] Map Limehouse," https://lists.openstreetmap.org/pipermail/talk/2005-July/000917.html, 2005, Accessed: 2025-04-14.

[7] D. Hristova, G. Quattrone, A. Mashhadi, and L. Capra, "The life of the party: Impact of social mapping in OpenStreetMap," in *Proceedings of the international AAAI conference on web and social media*, vol. 7, no. 1, 2013, pp. 234–243. [Online]. Available: https://doi.org/10.1609/icwsm.v7i1.14416

[8] OpenStreetMap Contributors, "OpenStreetMap Stats," https://wiki.openstreetmap.org/wiki/Stats, 2025, Accessed: 2025-04-14.

[9] R. Soden and L. Palen, "From crowdsourced mapping to community mapping: The post-earthquake work of OpenStreetMap Haiti," in *COOP 2014-Proceedings of the 11th International Conference on the Design of Cooperative Systems, 27-30 May 2014, Nice (France)*. Springer, 2014, pp. 311–326. [Online]. Available: https://doi.org/10.1007/978-3-319-06498-7_19

[10] R. DesRoches, M. Comerio, M. Eberhard, W. Mooney, and G. J. Rix, "Overview of the 2010 Haiti earthquake," *Earthquake spectra*, vol. 27, no. 1_suppl1, pp. 1–21, 2011. [Online]. Available: https://doi.org/10.1193/1.3630129

[11] Humanitarian OpenStreetMap Team, "HOT Tasking Manager," https://tasks.hotosm.org/, 2025, Accessed: 2025-04-14.

[12] S. Scholz, P. Knight, M. Eckle, S. Marx, and A. Zipf, "Volunteered geographic information for disaster risk reduction—the missing maps approach and its potential within the red cross and red crescent movement," *Remote Sensing*, vol. 10, no. 8, p. 1239, 2018. [Online]. Available: https://doi.org/10.3390/rs10081239

[13] S. Coast, "How OpenStreetMap is changing the world," in *International Symposium on Web and Wireless Geographical Information Systems*. Springer, 2011, pp. 4–4. [Online]. Available: https://doi.org/10.1007/978-3-642-19173-2_2

[14] M. Haklay, "How good is volunteered geographical information? A comparative study of OpenStreetMap and Ordnance Survey datasets," *Environment and planning B: Planning and design*, vol. 37, no. 4, pp. 682–703, 2010. [Online]. Available: https://doi.org/10.1068/b35097

[15] H. Fan, A. Zipf, Q. Fu, and P. Neis, "Quality assessment for building footprints data on OpenStreetMap," *International Journal of Geographical Information Science*, vol. 28, no. 4, pp. 700–719, 2014. [Online]. Available: https://doi.org/10.1080/13658816.2013.867495

[16] A. Mashhadi, G. Quattrone, and L. Capra, "Putting ubiquitous crowd-sourcing into context," in *Proceedings of the 2013 conference on Computer supported cooperative work*, 2013, pp. 611–622. [Online]. Available: https://doi.org/10.1145/2441776.2441845

[17] G. Quattrone, M. Dittus, and L. Capra, "Work always in progress: Analysing maintenance practices in spatial crowd-sourced datasets," in *Proceedings of the 2017 ACM Conference on Computer Supported Cooperative Work and Social Computing*, 2017, pp. 1876–1889. [Online]. Available: https://doi.org/10.1145/2998181.2998267

[18] M. Zook, M. Graham, T. Shelton, and S. Gorman, "Volunteered geographic information and crowdsourcing disaster relief: a case study of the Haitian earthquake," *World Medical & Health Policy*, vol. 2, no. 2, pp. 7–33, 2010. [Online]. Available: https://doi.org/10.2202/1948-4682.1069

[19] J. Xu, L. Li, and Q. Zhou, "Spatial-temporal analysis of OpenStreetMap data after natural disasters: a case study of Haiti under Hurricane Matthew," *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, vol. 42, pp. 1401–1406, 2017. [Online]. Available: https://doi.org/10.5194/isprs-archives-XLII-2-W7-1401-2017

[20] T. H. Poiani, R. D. S. Rocha, L. C. Degrossi, and J. P. De Albuquerque, "Potential of collaborative mapping for disaster relief: A case study of OpenStreetMap in the Nepal earthquake 2015," in *2016 49th Hawaii International Conference*

on System Sciences (HICSS). IEEE, 2016, pp. 188–197. [Online]. Available: https://doi.org/10.1109/HICSS.2016.31

[21] M. Minghini, A. Sarretta, and M. Napolitano, "Openstreetmap contribution to local data ecosystems in COVID-19 times: Experiences and reflections from the Italian case," *Data*, vol. 7, no. 4, p. 39, 2022. [Online]. Available: https://doi.org/10.3390/data7040039

[22] E. Kamptner and F. Kessler, "Small-scale crisis response mapping: comparing user contributions to events in openstreetmap," *GeoJournal*, vol. 84, no. 5, pp. 1165–1185, 2019. [Online]. Available: https://doi.org/10.1007/s10708-018-9912-1

[23] Herfort, Benjamin and Lautenbach, Sven and Porto de Albuquerque, João and Anderson, Jennings and Zipf, Alexander, "The evolution of humanitarian mapping within the OpenStreetMap community," *Scientific reports*, vol. 11, no. 1, p. 3037, 2021. [Online]. Available: https://doi.org/10.1038/s41598-021-82404-z

[24] M. Dittus, G. Quattrone, and L. Capra, "Mass participation during emergency response: Event-centric Crowdsourcing in Humanitarian mapping," in *Proceedings of the 2017 ACM conference on computer supported cooperative work and social computing*, 2017, pp. 1290–1303. [Online]. Available: https://doi.org/10.1145/2998181.2998216

[25] D. Maldeniya, M. De Choudhury, D. Garcia, and D. M. Romero, "Pulling through together: social media response trajectories in disaster-stricken communities," *Journal of Computational Social Science*, vol. 6, no. 2, pp. 655–706, 2023. [Online]. Available: https://doi.org/10.1007/s42001-023-00209-8

[26] A. Ahmouda, H. H. Hochmair, and S. Cvetojevic, "Analyzing the effect of earthquakes on OpenStreetMap contribution patterns and tweeting activities," *Geo-spatial Information Science*, vol. 21, no. 3, pp. 195–212, 2018. [Online]. Available: https://doi.org/10.1080/10095020.2018.1498666

[27] F. Lemmerich, D. Sáez-Trumper, R. West, and L. Zia, "Why the world reads Wikipedia: Beyond English speakers," in *Proceedings of the twelfth ACM international conference on web search and data mining*, 2019, pp. 618–626. [Online]. Available: https://doi.org/10.1145/3289600.3291021

[28] B. C. Keegan and J. R. Brubaker, "'Is' to'Was' Coordination and Commemoration in Posthumous Activity on Wikipedia Biographies," in *Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work & Social Computing*, 2015, pp. 533–546. [Online]. Available: https://doi.org/10.1145/2675133.267523

[29] T. Ruprechter, M. Horta Ribeiro, T. Santos, F. Lemmerich, M. Strohmaier, R. West, and D. Helic, "Volunteer contributions to Wikipedia increased during COVID-19 mobility restrictions," *Scientific reports*, vol. 11, no. 1, p. 21505, 2021. [Online]. Available: https://doi.org/10.1038/s41598-021-00789-3

[30] L. Kurek, C. Budak, and E. Gilbert, "Wikipedia in Wartime: Experiences of Wikipedians Maintaining Articles About the Russia-Ukraine War," *arXiv preprint arXiv:2409.02304*, 2024. [Online]. Available: https://doi.org/10.1145/3711107

[31] T. Yasseri, R. Sumi, A. Rung, A. Kornai, and J. Kertész, "Dynamics of conflicts in Wikipedia," *PloS one*, vol. 7, no. 6, p. e38869, 2012. [Online]. Available: https://doi.org/10.1371/journal.pone.0038869

[32] OpenStreetMap Contributors, "Russian–Ukrainian war," https://wiki.openstreetmap.org/wiki/Russian%E2%80%93Ukrainian_war, 2025, Accessed: 2025-04-18.

[33] USGS, "Earthquake Hazards Programme: M 7.0 - 10 km SE of Léogâne, Haiti," 2010, accessed: 2025-04-20. [Online]. Available: https://earthquake.usgs.gov/earthquakes/eventpage/usp000h60h/executive

[34] Texas Division of Emergency Management, "Hurricane Harvey DR-4332," 2017, accessed: 2025-04-20. [Online]. Available: https://tdem.texas.gov/disasters/hurricane-harvey

[35] L. Hamel, B. Wu, M. Brodie, S.-C. Sim, and E. Marks, *An Early Assessment of Hurricane Harvey's Impact on Vulnerable Texans in the Gulf Coast Region: Their Voices and Priorities to Inform Rebuilding Efforts*. Henry J. Kaiser Family Foundation, 2017. [Online]. Available: https://www.kff.org/report-section/an-early-assessment-of-hurricane-harveys-impact-on-vulnerable-texans-in-the-gulf-coast-region-appendices/

[36] California Department of Forestry and Fire Protection, "2020 Fire Season Incident Archive," 2020, accessed: 2025-04-20. [Online]. Available: https://www.fire.ca.gov/incidents/2020

[37] BBC: The Visual and Data Journalism Team, "California and Oregon 2020 wildfires in maps, graphics and images," 2020, accessed: 2025-04-20. [Online]. Available: https://www.bbc.co.uk/news/world-us-canada-54180049

[38] USGS, "Earthquake Hazards Programme: M 7.2 - Nippes, Haiti," 2021, accessed: 2025-04-20. [Online]. Available: https://earthquake.usgs.gov/earthquakes/eventpage/us6000f65h/executive

[39] B. J. Reinhart, A. Reinhart, and R. Berg, "National Hurricane Center Tropical Cyclone Report: Hurricane Grace," *HNational Hurricane Center*, 2021. [Online]. Available: https://www.nhc.noaa.gov/data/tcr/AL072021_Grace.pdf

[40] USGS, "Earthquake Hazards Programme: M 7.8 - 67 km NNE of Bharatpur, Nepal," 2015, accessed: 2025-04-20. [Online]. Available: https://earthquake.usgs.gov/earthquakes/eventpage/us20002926/executive

[41] ——, "Earthquake Hazards Programme: M 7.5 - 72 km N of Palu, Indonesia," 2018, accessed: 2025-04-20. [Online]. Available: https://earthquake.usgs.gov/earthquakes/eventpage/us1000h3p4/executive

[42] F. IMAIZUMI, N. OSANAI, S. KATO, M. KOIKE, K. KOSUGI, Y. SAKAI, H. SAKAGUCHI, Y. SATOFUKA, S. TAKAYAMA, T. TANAKA *et al.*, "Debris flow disaster in atami, japan, in july 2021," *International Journal of Erosion Control Engineering*, vol. 15, no. 1, pp. 1–6, 2022. [Online]. Available: https://doi.org/10.13101/ijece.15.1

[43] J. Wang, K. Li, L. Hao, C. Xu, J. Liu, Z. Qu, X. Yan, M. M. Sajjad, and Y. Sun, "Disaster mapping and assessment of pakistan's 2022 mega-flood based on multi-source data-driven approach," *Natural Hazards*, vol. 120, no. 4, pp. 3447–3466, 2024. [Online]. Available: https://doi.org/10.1007/s11069-023-06337-8

[44] Copernicus, "Forest Fires in Attika- Greece," 2018, accessed: 2025-04-20. [Online]. Available: https://mapping.emergency.copernicus.eu/activations/EMSR300/

[45] USGS, "Earthquake Hazards Programme: M 7.0 - 13 km NNE of Néon Karlovásion, Greece," 2020, accessed: 2025-04-20. [Online]. Available: https://earthquake.usgs.gov/earthquakes/eventpage/us7000c7y0/executive

[46] ——, "Earthquake Hazards Programme: M 7.8 - Pazarcik earthquake, Kahramanmaras earthquake sequence," 2023, accessed: 2025-04-20. [Online]. Available: https://earthquake.usgs.gov/earthquakes/eventpage/us6000jllz/executive

[47] CIMA Research Foundation, "The Italian floods of May 2023: a scientific analysis," 2023, accessed: 2025-04-20. [Online]. Available: https://www.cimafoundation.org/en/news/the-italian-floods-of-may-2023-a-scientific-analysis/

[48] L. Cremonini, P. Randi, M. Fazzini, M. Nardino, F. Rossi, and T. Georgiadis, "Causes and impacts of flood events in Emilia-Romagna (Italy) in May 2023," *Land*, vol. 13, no. 11, p. 1800, 2024. [Online]. Available: https://doi.org/10.3390/land13111800

[49] ERCC - Emergency Response Coordination Centre, "ECHO Daily Map of 16 August 2017," 2023, (Freetown) Accessed: 2025-04-20. [Online]. Available: https://erccportal.jrc.ec.europa.eu/ECHO-Products/Maps#/maps/2196

[50] ——, "ECHO Daily Map of 15 September 2023," 2023, accessed: 2025-04-20. [Online]. Available: https://erccportal.jrc.ec.europa.eu/ECHO-Products/Maps#/maps/4641

[51] H. Kadzuwa, "Post cyclone freddy disaster mapping – malawi," Scotland Malawi Partnership, Edinburgh, Scotland, Tech. Rep., 2023. [Online]. Available: https://www.scotland-malawipartnership.org/assets/resources/Cyclone-Freddy-Mapping-Final-Report.pdf

[52] Map Action, "Accessible areas and number of affected/displaced people (as of 13/03/23)," 2023, accessed: 2025-04-20. [Online]. Available: https://maps.mapaction.org/dataset/86aa74d0-7acf-40b3-9735-429b283b8809

[53] USGS, "Earthquake Hazards Programme: M 6.8 - Al Haouz, Morocco," 2023, accessed: 2025-04-20. [Online]. Available: https://earthquake.usgs.gov/earthquakes/eventpage/us7000kufc/executive

[54] FEMA, "2.3. Incident Response and Recovery Generic Operation Phases," 2024, accessed: 2025-04-20. [Online]. Available: https://www.fema.gov/oet-tools/chemical-incident-consequence-management/2/3

[55] S. J. Taylor and B. Letham, "Forecasting at scale," *The American Statistician*, vol. 72, no. 1, pp. 37–45, 2018. [Online]. Available: https://doi.org/10.1080/00031305.2017.1380080

[56] Uber Engineering, "H3: Uber's Hexagonal Hierarchical Spatial Index," 2018, accessed: 2025-04-21. [Online]. Available: https://www.uber.com/en-GB/blog/h3/

[57] R. Dorfman, "A formula for the Gini coefficient," *The review of economics and statistics*, pp. 146–149, 1979. [Online]. Available: https://doi.org/10.2307/1924845

[58] OpenStreetMap Contributors, "Taginfo: Tag statistics for OpenStreetMap," 2018, accessed: 2025-04-22. [Online]. Available: https://taginfo.openstreetmap.org/keys

[59] W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery, *Numerical Recipes: The Art of Scientific Computing*, 3rd ed. Cambridge University Press, 2007.

[60] California Department of Forestry and Fire Protection, "California State GeoportalCalifornia Fire Perimeters (all)," accessed: 2025-04-25. [Online]. Available: https://gis.data.ca.gov/datasets/CALFIRE-Forestry::california-fire-perimeters-all/explore

Chapter A

Appendix

A.1 Resources

- The GitHub repository for the project can be found at: https://github.com/FabianBindley/osm-times-of-crisis.
- The interactive dashboard used to explore the results, visualisations, and maps can be accessed at: https://osm-times-of-crisis.netlify.app/.

This project made use of the following Python libraries for data processing, analysis, and visualisation:

Library	**Purpose**
`osmnx`	Initial experiments with visualising OSM data.
`geopandas`	Initial experiments with generating GeoJSON boundaries by geocoding locations.
`matplotlib`	Plotting charts and visualisations.
`osmium`	Processing OpenStreetMap history files to capture the attributes of element changes
`geojson`	Handling GeoJSON format geographic data.
`folium`	Creating interactive HTML based maps of spatial distribution of changes.
`prophet`	Time-series forecasting of mapping activity.
`scipy`	Statistical analysis, including Kendall's Tau
`geopy`	Calculating geodesic distances between coordinates.

Table A.1: Python libraries used in the project implementation.

A.2 Charts for All Disasters

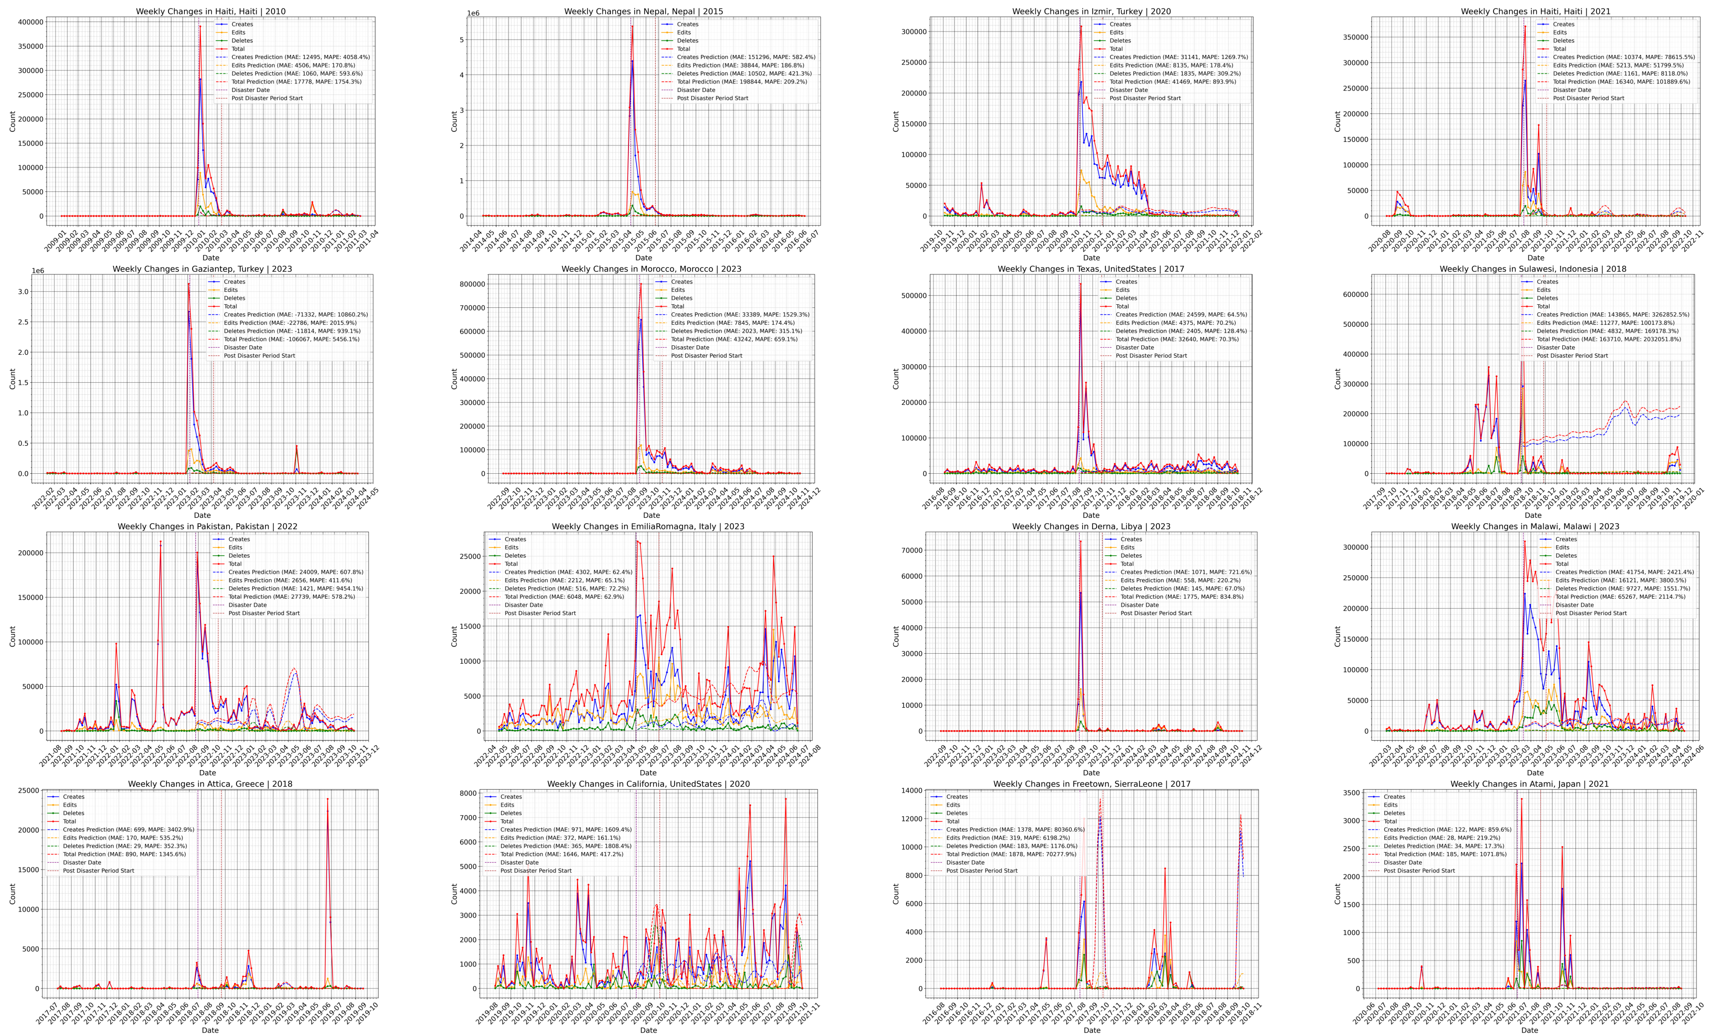

Figure A.1: Weekly change counts for all disasters. Ordered by disaster type

Figure A.2: `Total` change counts in the post-disaster period for all disasters, Prophet forecast and actual results.

Figure A.3: Weekly median number of days between consecutive changes for all disasters, Prophet forecast and actual results.

Figure A.4: Weekly tag change counts for all disasters. Post-disaster spikes frequently align with smaller peaks in overall weekly change counts, presented in Figure A.1.

A.3 Code Snippets

```
def diff_changes(curr, prev, way_handler, missing_way_count):
    diff = {
        "element_id": curr["element_id"],
        "edit_type_curr": curr["edit_type"],
        "edit_type_prev": prev["edit_type"],
        "element_type": curr["element_type"],
        "version_curr": curr["version"],
        "timestamp_curr": curr["timestamp"],
        "tags_created": [],
        "tags_edited": [],
        "tags_deleted": [],
        "coordinate_distance_change": None,
        "made_visible": False,
        "timestamp_between_edits": None,
        "way_nodes_created": [],
        "way_nodes_deleted": [],
        "way_length_change": None,
    }

    diff["coordinate_distance_change"] = round(
        compute_coordinate_distance_change(curr["coordinates"], prev["
            ↪ coordinates"]), 4)

    diff["timestamp_between_edits"] = curr["timestamp"] - prev["timestamp"]

    curr_tags = curr.get("tags", {})
    prev_tags = prev.get("tags", {})

    if curr_tags != prev_tags:
        diff['tags_created'] = [key for key in curr_tags if key not in
            ↪ prev_tags]
        diff['tags_deleted'] = [key for key in prev_tags if key not in
            ↪ curr_tags]
        diff['tags_edited'] = [key for key in curr_tags if key in prev_tags and
            ↪  curr_tags[key] != prev_tags[key]]
```

```
    if prev["visible"] == False and curr["visible"] == True:
        diff["made_visible"] = True

    if curr["element_type"] == "way":
        curr_key = (curr["element_id"], curr["version"])
        prev_key = (curr["element_id"], curr["version"] - 1)

        if curr_key in way_handler.way_data and prev_key in way_handler.way_data:
            curr_way_nodes = way_handler.way_data[curr_key]["nodes"]
            prev_way_nodes = way_handler.way_data[prev_key]["nodes"]

            diff['way_nodes_created'] = [node for node in curr_way_nodes if node not in prev_way_nodes]
            diff['way_nodes_deleted'] = [node for node in prev_way_nodes if node not in curr_way_nodes]
        else:
            missing_way_count += 1

    return diff, missing_way_count
```

Listing A.1: Function to compute the diffs between successive OSM element versions for the **change differences** implementation.

Chapter B

Project Reports

B.1 Project Plan

Project Plan for Final Year Project

Candidate Number: MFLP6

8th November, 2025

1 Project Title: Investigating crowd-workers mapping activity during times of crisis.

2 Supervisor: Professor Licia Capra

3 Aim, Objectives and Research Questions

3.1 Aim

The primary aim of this project is to analyse and understand the behaviour of crowd-workers contributing to OpenStreetMap (OSM) during times of crisis. Specifically, the project will investigate how mapping activities vary in response to natural disasters that impact the physical geography of environments.

We consider natural disasters because they are usually constrained to a given area, and take place over a set period of time, with initial impacts over a few days, but lasting consequences and responses over weeks, months or years.

We will consider multiple natural disasters from all over the world, to determine if mapping behaviour is consistent across different countries and cultures, and identify if the type of the disaster impacts behaviour.

3.2 Objectives

To achieve our aim, the following objectives will need to be met over the course of the project:

- Understand and analyse the existing literature involving contributions to OSM, to identify what hasn't yet been studied in terms of contributors' response during crises.
- Identify a set of research questions that when investigated support the aim of the project.
- Select a diverse set of natural disasters, especially including those that took place recently that haven't yet been widely studied. The geographical area being investigated should be constrained according to public reports on the disaster, and the historical edit data for the given area should be retrieved from OpenStreetMap.
- Develop a methodology to store and analyse the historical edit data, allowing for analysis of:
  - The timeline of edits throughout the disaster's response.
  - What type of edit is made to the map (new node, edit existing node, delete node)
  - What is the subject of the edit? (House, Shop, Restaurant, Infrastructure)
- Present the results of each research question clearly and concisely, explaining relevant findings. Compare and contrast findings from different disasters, identifying patterns and clustering in behaviour

3.3 Research Questions

Based on the aim and objectives, we propose 3 research questions which we will aim to answer in this project, to answer the overall question of what behaviour do crowd-workers exhibit and how might it vary?

RQ1: How does the mapping frequency in an area vary before, during and after disasters

RQ2: How does map content vary before, during and after disasters?

RQ3: How consistent is mapping behaviour across different natural disasters?

4 Expected Outcomes/Deliverables

The following deliverables can be expected upon project completion

- Comprehensive literature review of crowd-sourced mapping initiatives, their outputs and their impacts.
- Detailed and reproducible methodology allowing for future work investigating natural disasters. Automated tooling and setup explanations to enable this.
- Compelling results which present a deeper understanding of how crowd workers support in mapping during crises.
- In the case that there are biases or differences in mapping behaviour, make suggestions for how crowd workers could better provide support in mapping during crises.

5 Work Plan

- **October - Start November:** Make progress on literature search and understanding project requirements and scope. Write Project plan
- **November:** Agree project requirements and scope, continue literature review, begin work on collating and organising OSM data to allow for analysis. Agree on natural disasters to be considered.
- **December:** Refine analysis methodology, initially working with 1 or 2 natural disasters. Consider findings and decide
- **January:** Expand analysis to more disasters.
- **February:** Final analyses, start writing report
- **March:** Writing Report
- **April:** Complete Final Report

6 Ethics Review

All data that will be used is publicly available, and any personally identifiable information will be removed. No individual will be identifiable in this investigation and only aggregate statistics will be calculated. We affirm that an ethics review will not be necessary.

B.2 Interim Report

Interim Report for Final Year Project

Candidate Number: MFLP6

January 24th 2025

1 Project Title in October Research Project Plan: Investigating crowd-workers mapping activity during times of crisis.

2 Current Project Title: Investigating crowd-worker mapping activity during times of crisis.

3 Supervisor: Professor Licia Capra

4 Current Progress

Progress on the project has generally been smooth, with occasional bumps, which were successfully overcome. We break the progress down by term; as a reminder, the research questions are as follows:

Before, during and after disasters:
RQ1: How does the mapping frequency in an area vary?
RQ2: How does map content vary?
RQ3: How consistent is mapping behaviour across different natural disasters?

4.1 Term 1

Initially, I focused on defining the project requirements and purpose, and performed some early investigations into OpenStreetMap (OSM) historical change data, to understand OSM's data model. These investigations included counting the overall number of changes recorded, and those within a defined period, (pre-diaster, imm-diaster, post-disaster).

In parallel, I began reviewing the literature relevant to my project, in particular looking at how OSM has been used during times of crisis, with most papers focusing on 1 or 2 localised disasters. The existing literature highlights the importance of OSM being up to date to assist emergency responses and details some of the methods and behaviour of crowd-workers. Next, I considered which disasters I wished to investigate; after deliberation and research, I settled on the following list:

Americas

- 2010 - Haiti Earthquake
- 2017 - Hurricane Harvey (US)
- 2020 - California Wildfires (US)
- 2021 - Haiti Earthquake and Hurricane Grace

Asia

- 2015 - Nepal Earthquake
- 2018 - Sulawesi Earthquake and tsunami (Indonesia)
- 2022 - Pakistan Floods
- 2024 - Wayanad Landslide (India)

Europe and the Middle East

- 2020 - Izmir Earthquake (Turkey)
- 2023 - Emilia Romagna Floods (Italy)
- 2023 - Turkey/Syria Earthquake
- 2024 - Valencia Floods (Spain)

Africa

- 2022 - Derna Dam Collapse (Libya)
- 2023 - Cyclone Freddy (Malawi)
- 2023 - Morocco Earthquake

Next, as manually reading in OSM `.osh` history files is computationally slow, I set up a PostgreSQL database with PostGIS following a methodology explained in [1]. The changes for a select few disasters were loaded into the database, and filtered to remove all changes occurring more than a year before, and a year + a month after the disaster. I manually defined the boundary of the disasters according to information about the disasters' impacts available online, focusing on the areas most affected, and filtered out any changes that did not fall within this boundary. The final part of preprocessing involves removing changes made during a bulk import, as these are typically largely automated, and do not reflect normal OSM user behaviour. The PostgreSQL allows for very efficient and fast querying and counting of changes at scale, essential for at the time of writing, handling approximately 27 million changes. We then began to investigate how change counts vary across different disaster periods by change type (create, edit, delete), plotting charts similar to that of figure 1.

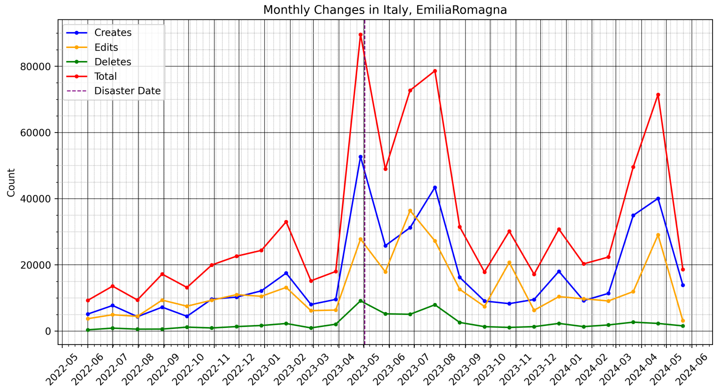

Figure 1: Plot of monthly changes by change type, for the Emilia Romagna floods of 2023

Next, we examined how change counts varied at a local level with disaster areas, by aggregating them into Uber H3 hexagons [2]. These are a system of hexagons which cover the entire planet, with smaller areas at higher resolutions, allowing disaster areas to be broken down into equally sized chunks. At resolution 7, each hexagon covers approximately the area of a town's neighbourhood, allowing the distribution of changes within the disaster area to be analysed. We initially computed the raw change counts for disaster areas and plotted them on a map, however because of the high concentration of change counts in urban areas, we decided to also compute the % difference in the number of counts in the immediate and post-disaster periods, compared to pre-disaster.
Figures 2 and 3 demonstrate the areas experiencing changes.

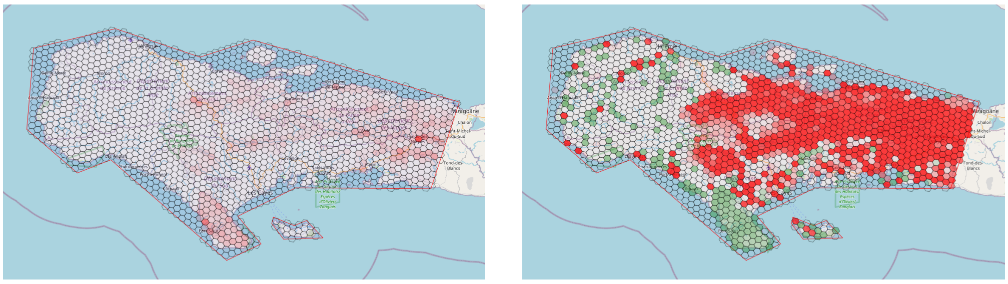

(a) Haiti Earthquake 2021 - Raw density of counts

(b) Haiti Earthquake 2021 - Percentage difference in change count from pre-disaster period to post-disaster.

Figure 2: Comparison of change count density mapping: raw counts, and % difference for Haiti 2021

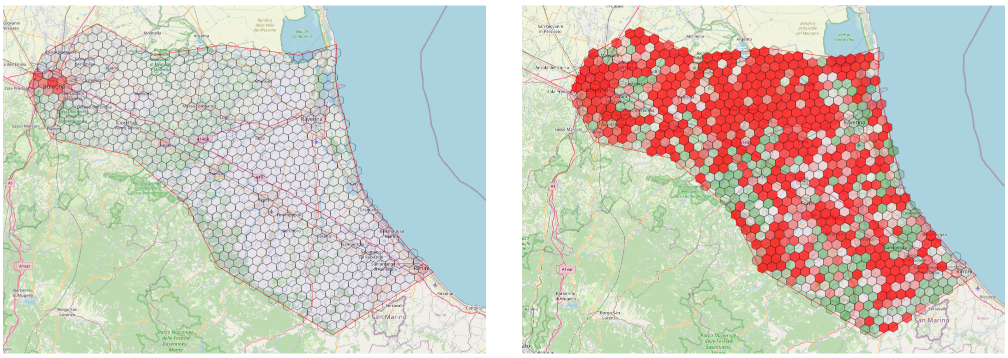

(a) Emilia Romagna Floods 2023 - Raw density of counts

(b) Emilia Romagna Floods 2023 - Percentage difference in change count from pre-disaster period to post-disaster.

Figure 3: Comparison of change count density mapping: raw counts, and % difference for Emilia Romagna 2023

The next task was to define the requirements and metrics to be used for the analysis; these were more formally defined for RQ1 and RQ2, but with room to manoeuvre where necessary. To display the lower-level count maps, a project visualisation website was instantiated using ReactJS and Vite and hosted locally. This website has proved invaluable, as it allows us to view maps and charts for different disasters, and quickly adjust settings eg: the hexagon resolution, the period, or the map style. The website was soon extended to include the change counting graphs (eg: figure 1), and ongoing development has added new visualisations and metrics as the project has progressed.

4.2 Term 2

We started term 2 by computing the Gini coefficient for the lower-level counts, representing the inequality in the spread of changes across disaster areas. Next, we experimented with time-series analysis modelling with Prophet, to compare how different behaviour is after a disaster, and to what extent the behaviour returns to pre-disaster levels. There were some difficulties with this modelling due to variability in counts in training data in the pre-disaster period, therefore investigation into it will continue.

Having completed the bulk of the research for RQ1, we moved onto RQ2, by first investigating the tags used on OSM. The TagInfo database was initially consulted to better understand the most commonly used tags on the platform. Tags describe information about an object, formed by `key=value` pairs. Examples of keys include the object's type: `building, highway, natural`, what an object does: `amenity, service, landuse`, or details about an object: `name, addr:housenumber, opening_hours`. OSM does not have a strictly defined tag hierarchy, an amenity does not have to be a building for example, but conventions defined as the platform has evolved guide users on how tags should be used.

We then analysed the frequency of tag keys appearing in changed objects across disasters and periods, noting that a given object may have multiple keys. The results shown in table 1 demonstrate that buildings and highways are the most commonly changed physical objects, with other keys mostly representing additional information about objects.

Index	Key	Count	% of Total Changes For Period
0	`building`	2,449,841	9.020%
1	`highway`	597,686	2.201%
2	`source`	501,427	1.846%
3	`name`	218,585	0.805%
4	`landuse`	130,276	0.480%
5	`natural`	116,955	0.431%
6	`surface`	107,134	0.395%
7	`damage:event`	76,262	0.281%
8	`leaf_type`	69,723	0.257%
9	`genus`	67,986	0.250%
10	`species`	67,736	0.249%
11	`genus:it`	65,554	0.241%
12	`species:it`	63,771	0.235%
13	`amenity`	54,866	0.202%
14	`addr:street`	52,895	0.195%
15	`leaf_cycle`	51,449	0.189%
16	`species:wikidata`	50,856	0.187%
17	`waterway`	48,976	0.180%
18	`attribute_source_date`	42,384	0.156%
19	`attribute_source_type`	41,091	0.151%
20	`addr:housenumber`	39,261	0.145%

Table 1: The most common tag keys, their counts, and their percentage of total changes for the period.

Similarly, the frequencies of values of some of the most common keys (**building, highway, source, name, surface, amenity, landuse, waterway, natural**) were calculated to identify the most common objects being changed. Table 2 shows the most common amenities in the pre-disaster period, while table 3 presents the imm-disaster period.

Index	Key	Count	% of Total Changes For Period
0	school	2,142	12.510%
1	parking	1,890	11.038%
2	place_of_worship	1,823	10.647%
3	drinking_water	1,590	9.286%
4	restaurant	976	5.700%
5	toilets	858	5.011%
6	bench	487	2.844%
7	pharmacy	443	2.587%
8	hospital	423	2.470%
9	gambling	380	2.219%
10	bicycle_parking	369	2.155%

Table 2: Pre-disaster: most common values for **amenity**, their counts, and their % of total changes.

Index	Value	Count	% of Total Changes For Period
0	hospital	4,444	41.932%
1	drinking_water	2,362	22.287%
2	school	1,007	9.502%
3	place_of_worship	550	5.190%
4	parking	282	2.661%
5	restaurant	169	1.595%
6	bicycle_parking	130	1.227%
7	bench	98	0.925%
8	fuel	97	0.915%
9	cafe	88	0.830%
10	shelter	86	0.811%

Table 3: Imm-disaster: most common values for **amenity**, their counts, and their % of total changes.

In the 30-day imm-disaster period, we observe a significant increase in changes relating to **hospital** and **drinking_water** objects, as can be expected when responding to a disaster. We seek to identify other object **key=value** pairs with significant differences compared to pre-disaster.

To compute an aggregate metric for the difference in the type of objects being changed, across the range of keys above, as well as the values associated with each key, we experimented with computing a range of coefficients that represent the rank order or count of each key/value during 2 distinct periods, such as pre and imm-disaster. This approach allows us to understand how similar the ordering of the keys/values and their counts are. We compute the Kendall rank correlation coefficient for the rank order and cosine similarity/Pearson correlation coefficient for the magnitude of counts.

In many cases, performing computations on changes or generating charts is done independently for each disaster. To greatly improve performance, multiprocessing was introduced, allowing a separate process to be created for each disaster. This approach significantly reduced the overall computation time, as the computations for each disaster can run in parallel without conflict or the need for processes to wait for each other.

Finally, the project visualisation website was published online at osm-times-of-crisis.vercel.app, deployed on Vercel. Anyone can view the charts and results, which are updated on pushes to the project's GitHub repository. The website will remain available indefinitely, providing future researchers with a resource to benefit from and build upon our work.

5 Remaining Work

A significant amount of work remains to complete this project, but we do not anticipate any major blockers, and aim to deliver a complete final report, at least an initial draft, by the end of Term 2. We outline an ambitious yet flexible project timeline to guide our progress, expecting that adjustments may be made to accommodate any challenges.

5.1 Weeks 3 and 4

- Formally decide on the keys which we will use for the value analysis: essentially which type of objects will we be looking at. I am expecting that `amenity`, will be one of them, other possible options include `highway`, `building`, `leisure`, `emergency` and more. These keys should describe the purpose of objects in the real world, rather than provide supporting information.
- Investigate the keys/values across all disasters, and individually, to determine whether object types are more frequently associated with a certain edit type (create, edit and delete).
- Begin investigating what users are actually changing when they make changes. Creates and deletes are self-explanatory, but edits to an object could be users changing any of a number of things:
  - Creating or deleting tags (tag key changes)
  - Updating existing tags (tag value changes)
  - Moving a node, updating its coordinate (measure the distance between old and new coordinates)
  - Adding, editing or deleting nodes defining a way

  A `diff` function will need to be defined, to compare the differences between a change, and its previous version (according to its version number).

  If the previous version of a change is not stored in the PostgreSQL database, it can be individually retrieved from the `.osh` file and added to the database for subsequent reuse.
- Continue analysis of Prophet modelling results, to ensure that they can be relevant for answering RQ1.
- Update longer running scripts with multiprocessing by `disaster_id` to reduce computation time.
- Once analysis mostly finalised, import the datasets for all remaining disasters, and define the geographical boundaries of the affected disaster areas.

5.2 Week 5 and Reading week

- Begin work on RQ3 by comparing results between different disasters, to determine if there are any differences that could be attributed to the disaster type or geographic region (Americas, Asia etc).
- Agree with Professor Capra on the methodology, results and analysis which will be presented in the report, and draft a plan detailing these.
- Continue reading papers to prepare for the literature review and look for a few more that are relevant.
- Draft plans for the introduction, preparatory work, literature review and requirements/research agenda.
- Extend the methodology plan by adding details about the design + implementation.

5.3 Weeks 6 and 7

- Write:
  - Introduction
  - Literature review
  - Requirements/research agenda

5.4 Weeks 7 and 8

- Write:
  - Methodology and design + implementation
  - Results
  - Analysis

5.5 Weeks 9 and 10

- Complete the project's first draft, ready to be checked by Professor Capra.
- Make changes according to suggestions, to improve the report, and finalise the draft.

5.6 Holidays, up to submission deadline

- Catch up with project progress if it has slipped beyond the dates above.
- Ensure the codebase and repository are ready to be shared, with instructions and guidance provided in the form of a README. Also, ensure that the visualisation site has all the necessary content and is well-presented.
- Submit the project.

References

[1] E. Kamptner and F. Kessler, "Small-scale crisis response mapping: comparing user contributions to events in openstreetmap," *GeoJournal*, vol. 84, no. 5, pp. 1165–1185, 2019. [Online]. Available: https://doi.org/10.1007/s10708-018-9912-1

[2] Uber Engineering, "H3: Uber's Hexagonal Hierarchical Spatial Index," 2018, accessed: 2024-12-04. [Online]. Available: https://www.uber.com/en-GB/blog/h3/