

AperTO - Archivio Istituzionale Open Access dell'Università di Torino

Work always in progress: Analysing maintenance practices in spatial crowd-sourced datasets

This is a pre print version of the following article:

Original Citation:

Availability:

This version is available <http://hdl.handle.net/2318/1730485> since 2020-02-24T16:12:48Z

Publisher:

Association for Computing Machinery

Published version:

DOI:10.1145/2998181.2998267

Terms of use:

Open Access

Anyone can freely access the full text of works made available as "Open Access". Works made available under a Creative Commons license can be used according to the terms and conditions of said license. Use of all other works requires consent of the right holder (author or publisher) if not exempted from copyright protection by the applicable law.

(Article begins on next page)

Work Always in Progress: Analysing Maintenance Practices in Spatial Crowd-sourced Datasets

Giovanni Quattrone
University College London
London, UK
g.quattrone@cs.ucl.ac.uk

Martin Dittus
University College London
London, UK
martin.dittus.11@ucl.ac.uk

Licia Capra
University College London
London, UK
l.capra@cs.ucl.ac.uk

ABSTRACT

Crowd-mapping is a form of collaborative work that empowers users to share geographic knowledge. Despite geographic information being intrinsically evolving, little research has so far gone into analysing maintenance practices in these domains. In this paper, we quantitatively capture maintenance dynamics in geographic crowd-sourced datasets, in terms of: the extent to which different maintenance actions are taking place, the type of spatial information that is being maintained, who engages in these practices and where. We apply this method to 117 countries in OpenStreetMap, one of the most successful examples of geographic crowd-sourced datasets. Furthermore, we explore what triggers maintenance, by means of an online survey to which 96 OpenStreetMap contributors took part. Our findings reveal that, although maintenance practices vary substantially from country to country in terms of how widespread they are, strong commonalities exist in terms of what metadata is being maintained, by whom, and what triggers them.

INTRODUCTION

Crowd-sourcing has become a successful paradigm for knowledge gathering, where a crowd is mobilised to collect and maintain large repositories of information [16, 3]. The most successful example to date of this paradigm is Wikipedia, with its online community of editors that voluntarily contribute to build and maintain the whole body of knowledge. Another type of knowledge where crowd-sourcing has been widely applied is that of volunteered geographic information, with citizens becoming surveyors, in council-monitoring applications like FixMyStreet;¹ local reporters, as powered by Ushaidi's Crowdmap;² and cartographers, in geo-wikis like Cyclopath³ and OpenStreetMap.⁴ It is the latter type of knowledge that we are interested in this paper.

¹<http://www.fixmystreet.com/>

²<http://www.ushaidi.com/products/crowdmap>

³<http://cyclopath.org/>

⁴<http://www.openstreetmap.org/>

Research has developed methods to quantitatively analyse the accuracy [2, 8, 11, 12, 13, 14, 27], coverage [43, 13, 28], growth [37], and bias [43, 18, 13, 35] of volunteered geographic information. These methods have mainly been applied to OpenStreetMap, making this dataset probably the most widely and deeply investigated geographic information repository to date. However, there is currently a gap in terms of methods to quantitatively capture *maintenance practices* in geographic crowd-sourcing communities like OpenStreetMap. Geographic information is naturally volatile and always evolving (e.g., where a grocery store is today, a coffee shop might be tomorrow); indeed, companies like Google spend several billions of dollars each year just to maintain their proprietary maps up-to-date and to improve their accuracy.⁵ Yet little is known about maintenance practices of geographic crowd-sourced information: whether maintenance takes place at all and, if so, where, about what, and by whom.

Maintenance practices have been analysed in other crowd-sourced knowledge-gathering communities like Wikipedia [20, 10, 9]. However, those methods and findings cannot be directly transferred over to crowd-sourcing communities focused on spatial knowledge. This is because geographic repositories differ from encyclopaedic knowledge ones in two fundamental ways: (i) contributions to geographic databases require knowledge that usually only *the locals* have (i.e., contributors may need to be physically in a place to be able to contribute knowledge); as a result, the emerging dynamics of location-based knowledge gathering and maintenance can be quite different from those emerging in purely online settings. Furthermore, (ii) geographic content is subject to continuous change, due to natural processes such as urbanisation, gentrification, and adverse events; as such, maintenance practices to keep this knowledge up-to-date are likely different from those required to maintain more encyclopedic types of knowledge.

In order to analyse maintenance practices in spatial crowd-sourced datasets, we have developed a method that quantitatively captures: (i) the different types of maintenance actions that take place (i.e., enrichment vs. correction vs. removal of existing information), and how widespread they are; (ii) what type of spatial objects (e.g., schools, hospitals, restaurants) are being maintained; (iii) who is mostly engaged in maintenance practices; and (iv) where such actions are taking place. We have applied this method to OpenStreetMap,

⁵<http://www.wired.com/2014/12/google-maps-ground-truth>

analysing one year of mapping activity in 117 different countries. To further understand (*v*) what triggers maintenance, we have conducted an online survey, to which 96 OSM users responded.

Our findings reveal that maintenance varies substantially from country to country, both in terms of its adoption (i.e., the extent to which it is practised), and in terms of the types of POIs that are being maintained (with all countries focusing on a very small set of POI types, but these being largely different from country to country). Despite big differences in the types of POIs that are being maintained, the tags associated to them that are being added/updated/removed are common across many countries. Our study also revealed that some maintenance actions, such as the addition of new tags to existing spatial objects, are more frequent than other actions, such as the updating or the removal of tags from existing spatial objects. At the moment, these maintenance actions are prevalently done by highly active users. Finally, our online survey revealed that maintenance is often the result of stumbling upon incorrect map information while using external services (e.g., navigation) that rely on the OSM base map. Our findings highlight opportunities for tool developers of crowd-mapping platforms to design specialised tools and work-flows to help users identify what OSM information needs to be maintained and to act upon it.

RELATED WORK

Maintenance practices have been extensively studied in online self-organised communities, most especially Wikipedia. In this domain, researchers often refer to collaborative practices, rather than maintenance ones, intended as the editing activity performed by different editors on the same Wikipedia article, for example to update its content or to improve its quality. In Wikipedia, collaboration has been studied from different perspectives: (*i*) its spread and temporal evolution; (*ii*) the information that is being maintained; and (*iii*) who performs this practice. We review some of the works in each of these themes next.

Spread and temporal evolution. Research has quantified how widespread collaborative practices are in Wikipedia, and how these evolve over time [21, 26]. In this stream of work, researchers found that initially (i.e., when the Wikipedia community was in its early years), most of the editing effort was spent in creating new articles; however, over time other activities (e.g., user coordination, collaborations and discussions) have enormously raised [22, 39].

What information is being maintained. Kaltenbrunner and Laniado analysed the evolution over time of maintenance practices on different topics in Wikipedia [20]. They found that Wikipedia is the most up-to-date encyclopaedia ever seen, and that maintenance is often triggered by external events, with pages about such ongoing events being often edited and discussed on Wikipedia nearly in real-time. On the other hand, articles about historical or scientific facts (i.e., those that are not on people's minds) may take years to reach similar levels of user attention. Similar research done by Ferron et al. [10] observed that articles related to traumatic events often receive many maintenance edits in correspondence with

anniversaries. In [9] the same authors presented a study of activity on different language versions of Wikipedia, specifically during the Egyptian revolution, and found evidence of intensive and rapid participation on articles related to such event.

Who engages in collaboration practices. A study conducted by Laniado and Tasso [24] described the evolution of the user collaboration network in Wikipedia. They found that there exists a nucleus of very active contributors, who seem to spread over the whole wiki, and who interact preferentially with inexperienced users. Other studies that focused on users and their collaborative practises found that the top Wikipedia editors are those who are more involved in article maintenance, revising already existing articles, using quality assurance systems, and invoking community norms [15, 33].

Maintenance/collaboration practices is an active research area also for volunteered geographic information (VGI); however, in this context, current research is mainly investigating how to design tools to facilitate collaboration practices in crisis mapping [1, 7, 23] and little research has gone into analysing these practices more generally. As we shift our attention from encyclopaedic knowledge to spatial knowledge, different collaboration practices may be adopted. In fact, geographic repositories differ from classic encyclopaedic ones in two fundamental ways: *space* and *time*. Specifically, geographic content has an intrinsic spatial dimension, and there is a relationship between the location of a contributor and the type of knowledge that she can offer. Furthermore, compared to the body of knowledge that repositories like Wikipedia maintain, most geographic content is intrinsically volatile and continuously evolving, as a result of natural processes, such as urbanisation. As the nature of content varies, so might the corresponding editing practices. Indeed, a study conducted a few years ago by Mashhadi et al. [29] showed that some properties that typically hold in encyclopaedic type of crowd-sourcing repositories like Wikipedia, do not hold in geographic ones such as OpenStreetMap; for example, it was found that, in the former, the quality of an article depends on how much editing experience its contributors had in the past, while no relationship was found between quality of the map and editing experience of mappers in OSM.

In this paper, we aim to cover this gap, by proposing a method to quantitatively capture maintenance practices in spatial crowd-sourced datasets. Before presenting the method itself, and reporting on the results obtained, we first briefly illustrate the dataset we chose for analysis, provide a working definition of maintenance over such dataset, and spell out the research questions our methods aims to answer.

Note that the crowd-sourcing domain we investigate is fundamentally different from that of crowd-sensing (e.g., Waze⁶). In the latter, the temporal validity of the passively collected information (e.g., GPS) is very short, and indeed one could argue that such data is not supposed to be 'maintained', but rather 'replaced' by fresher data all the time.

DATASET

⁶<https://www.waze.com/>

We chose to apply our method to OpenStreetMap (hereafter OSM), as this is to date the most successful example of spatial crowd-sourced dataset, having been running since 2004, and comprising the largest (and most geographically widespread) user and content base. Furthermore, OSM has been subject to extensive research, so that we can relate our findings to previous studies.

The OSM dataset is freely available to download⁷ and contains the history from 2006 of all edits (over 2.7 billions) performed by all users (over 2 millions) on all spatial objects. In OSM jargon, spatial objects can be one of three types: *nodes*, *ways*, and *relations*. Nodes are single geo-spatial points and typically represent Points-of-Interest (POIs); ways mostly represent roads (as well as streams, railway lines, and the like); finally, relations are used for grouping other objects together, based on logical (and usually local) relationships (e.g., bus routes).

We filter the data in a number of ways before we begin our analysis. Specifically, we restricted our attention to edits of POIs only, i.e., specific point locations described in OSM by latitude/longitude coordinates, plus a variety of attributes (or tags). By focusing on this subset of OSM objects (instead of ways and relations), we aim to capture the actions of a wide range of contributors, from casual mappers to highly-engaged ones; indeed, as Mooney and Corcoran describe: “*Editing or adding tags to objects in OSM is technically one of the simplest operations which contributors can perform as there is very good support in all of the software and web-based editors for this edit action*” [30]. In OSM, a POI edit is represented as a tuple:

$$\langle uid, changeset, tstamp, ver, lat, lon, taglist \rangle$$

where *uid* identifies the user who performed this edit, *changeset* denotes the editing session within which this edit was performed; *tstamp* is the timestamp of when this edit took place; *ver* is a sequential value indicating the edit version of this POI (i.e., *ver* = 1 indicates the POI has just been created, while *ver* > 1 indicates the current edit is an update (i.e., maintenance) of an already existing POI); *lat* and *lon* denote the geographic coordinates of the POI. Finally, *taglist* contains an arbitrary list of attribute-value pairs that further describe the POI; examples of such attributes are ‘name’ (e.g., ‘Hollywood Cafe’), ‘amenity type’ (used to distinguish between different categories of POIs, such as ‘restaurant’, ‘pubs’, ‘school’), address details, opening hours, accessibility considerations, and so on. For the purpose of this study, we consider POIs to be all OSM nodes that have either a *name* or an *amenity* tag at any point in the relevant period.⁸ Finally, we ignored the tag *created_by*, as this is added automatically by editing software and does not reflect user intent.

The second pruning step we performed was time-based. We wanted to avoid the initial phase of OSM, when almost all contributions are creations of new objects, with little to no maintenance work taking place. We thus did not consider in our analysis all POI edits done before January 1st 2014.

⁷<http://www.geofabrik.de/data/download.html>

⁸http://wiki.openstreetmap.org/wiki/Map_Features

From the above dataset, we make an attempt to identify contributions by human editors, while discarding automated contributions representing bulk data imports. In some regions, a significant portion of OSM contributions are automated imports of public domain map data sets, often produced by national mapping organisations or derived from historic map data.⁹ While such data can play an important role in filling gaps on the map, it was not produced by the OSM community of volunteer contributors, and is not representative of human maintenance practices, which is the subject of this study. Imports are not explicitly marked as such in the OSM dataset; we thus needed heuristics to identify them. We applied the same approach used in [35], and marked as imports those edits which came from a single user, in very large quantities (i.e., more than one thousand edits), in a short period of time (i.e., less than one hour), and that were spread over a large geographic area (i.e., in the scale of a whole city).

The final part of this pruning process is to select the geographical areas of the world to analyse. To do so we need first to define a spatial unit of analysis. We expect that maintenance practices are somewhat related to the maturity of the OSM map, and previous research shows that different countries have rather different levels of OSM map maturity [36]. We therefore chose to study maintenance practices at country level. From the above sample, we discard countries with too little OSM editing activity to be meaningfully analysed (i.e., countries with less than one thousand contributions during the period of study). We ended up with a dataset having around 3.4M edits, of 2.7M POIs, done by 80k users, over the 117 countries highlighted in Figure 1. The summary characteristics of our dataset are reported in Table 1.

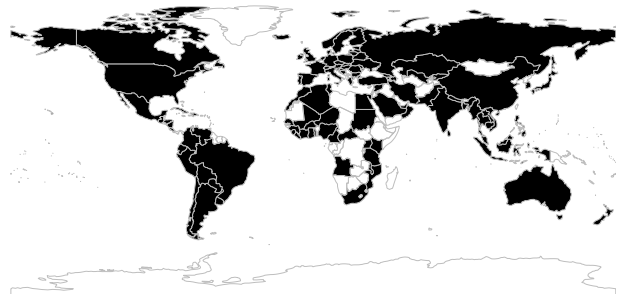


Figure 1. Map of the 117 Countries Under Analysis

	Min	1st Qu.	Median	3rd Qu.	Max.
OSM edits per country	1,021	2,447	6,670	23,605	651,315
OSM POIs per country	912	2,176	5,562	19,368	500,327
OSM users per country	42	117	233	644	15,829

Table 1. Summary Statistics of OSM Features in the 117 Analysed Countries

FORMS OF MAINTENANCE

In order to quantitatively analyse maintenance of OSM information (and, more specifically, OSM POIs), we first need to automatically identify edits, in the OSM edit history, that are representative of such practice. We simply classify as maintenance actions all edits with *ver* > 1. Preliminary analysis

⁹<http://wiki.openstreetmap.org/wiki/Import>

shows that the time interval between two consecutive edits ($ver = n$ and $ver = n + 1$) is one week or longer for 90% of such edits, and that different users perform them.

We then distinguish three different forms of information maintenance, based on the type of *action* that took place since the POI previous version:

- *Add*, maintenance work where at least one new tag has been added to an existing POI (e.g., the tag ‘opening_hours’, along with its associated value, has been added to a restaurant already mapped in OSM).
- *Update*, maintenance work where the value of at least one of the already existing tags associated with a POI has been updated (e.g., the value of the tag ‘amenity’ has been changed from ‘restaurant’ to ‘cafe’, for a POI previously added to OSM).
- *Remove*, maintenance work where at least one tag has been deleted from an existing POI (e.g., the tag ‘is.in’, along with its associated value, has been removed from a POI present in OSM).

Note that the same edit may belong to different action classes (e.g., a single edit can both add a tag and update another). In our study, we will analyse them separately, as the drivers behind such actions can be quite different, and might thus result in different practices. In fact, intuitively, an add action can be seen as a sign of the user intent to *enrich* existing information, and it might be spurred by the emergence of novel location based services that require semantically richer POI information (e.g., opening hours, webpage). Conversely, an update action can be seen as a sign of the user intent to *correct* existing information; this may be the case for POIs that were last edited a long time ago, and thus now contain stale information (e.g., different business name or type), or the case for POIs whose name contains spelling mistakes. Finally, a remove action can be seen as a sign of the user intent to *polish* existing information; this may be the case for POIs that contain some deprecated tags.

RESEARCH QUESTIONS AND APPROACH

In this work, we aim to explore the following five research questions about maintenance practices of spatial crowd-sourced datasets:

- RQ1 (Spread)** – How widespread is maintenance work? We begin our exploration by looking at the extent to which such practice is currently taking place across the 117 countries under exam, possibly identifying factors (e.g., map maturity) that correlate with higher (or lower) maintenance practices.
- RQ2 (What)** – What information is being maintained? We then look more specifically at the type of information that is being maintained, to elicit POI information that is commonly maintained across all countries, if any, as well as potential regional differences.
- RQ3 (Who)** – Who is engaged in information maintenance? We next shift our attention to the users performing maintenance edits, to understand whether this practice is evenly

shared among editors, or whether it is undertaken by a select few.

RQ4 (Where) – Where is information maintenance taking place? We then narrow down the spatial dimension of the crowd-sourced information being maintained, to understand whether there are geographic areas, within each country, that naturally attract more information maintenance than others.

RQ5 (Triggers) – Last but not least, we investigate how users decide to engage in maintenance actions – that is, what triggers their decision to maintain existing OSM information.

To answer the first four questions, we defined new metrics and conducted a large-scale quantitative analysis of maintenance practices of over 80k OSM mappers spread across 117 different countries. To answer the fifth and last question, we developed and distributed a questionnaire among OSM contributors, and analysed answers from 96 mappers by means of thematic analysis [4].

METRICS AND RESULTS

RQ1 – How widespread is maintenance?

We started our analysis by quantifying how widespread maintenance practices are across the 117 countries under exam. In this regard, we formulate the following hypothesis.

Hypothesis 1 *Maintenance practices are directly linked to map maturity.*

This hypothesis is inspired by the intuition that crowd workers may first concentrate on adding information to a near empty map, and only later, as the map becomes richer and denser of information, they may start to move towards maintaining what is already there.

To validate this hypothesis, we need first to define a metric that measures how widespread maintenance practices are in a certain country, then another metric that measures the OSM map maturity of that country.

To quantify the former, in each country under exam, we compute the proportion of maintenance work that took place there, relative to the total number of edits (i.e., covering both creation and maintenance of POIs), for the period of study. We name this metric as *Maintenance Ratio* (MR). Formally:

Definition 1 *Let OSM_e be the set of OSM edits for a given country, let $OSM_m \subseteq OSM_e$ be the set of OSM edits devoted to maintaining existing POIs. Then Maintenance Ratio is defined as $MR = \frac{|OSM_m|}{|OSM_e|}$.*

$MR \in [0, 1]$ by definition. Intuitively, the closer this metric is to 1, the higher the proportion of maintenance work in that country; conversely, values close to zero indicate that almost all OSM editing activity is devoted to the creation of new POIs.

We report in Figure 2 a heatmap of MR values in the 117 analysed countries and in Table 2 the same values divided in quartiles. Maintenance practices vary widely: in a quarter

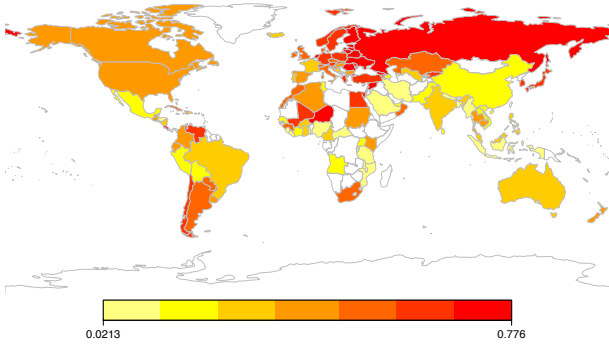


Figure 2. Maintenance Ratio in all Analysed Countries

of the analysed countries, maintenance is almost as frequent as the creation of new POIs ($MR > 0.42$), while there is another quarter of countries where maintenance is a much less widespread practice ($MR < 0.23$). There are also a few countries (e.g., Malawi, Mozambique, and Togo) where MR is almost zero, meaning that, in these countries, crowd workers are almost completely focused on the addition of new POIs, rather than on the maintenance of existing ones.

Min	1st Qu.	Median	3rd Qu.	Max.	Freq. Distr.
0.02	0.23	0.33	0.42	0.77	

Table 2. Summary Statistics of Maintenance Ratio in the 117 Analysed Countries

To estimate OSM map maturity, in each country under exam we compute the number of OSM POIs mapped in that country, normalised by the number of people living in that country.¹⁰ We then compute the Spearman correlation [31] between such proxy and MR . As hypothesised, we do obtain a positive correlation ($\rho = 0.40$, p -value < 0.001), meaning that maintenance practices are more widespread in countries where OSM map maturity is higher. However, the strength of the correlation is not very high. Figure 3 shows the scatter plot between map maturity and MR . As showed, MR is high both in countries with high OSM map maturity (e.g., Netherlands, Germany, Luxembourg), and in countries where OSM map maturity is much lower (e.g., Syrian Arab Republic, Egypt); low values of MR are registered both in countries with low map maturity (e.g., Mozambique, Saudi Arabia, Nigeria), and in countries having a relatively high value of map maturity (e.g., France, Lebanon, Liberia).

We speculate that a variety of other local factors play an important role in the rapid uptake of maintenance practices; these may lower the correlation between map maintenance and OSM map maturity. For example, initiatives like the Humanitarian OpenStreetMap Team (HOT) may have fostered maintenance activities in countries with lower-than-average map maturity, in order to rapidly update the map following

¹⁰We normalise by the number of people living in a country, rather than the surface area of the country, to take into account the fact that some big countries (such as Russia, Canada or Australia) are sparsely populated.

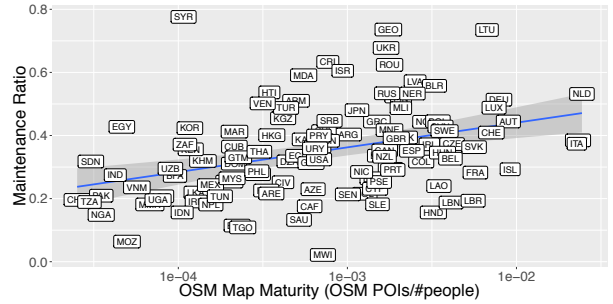


Figure 3. OSM Map Maturity vs. Maintenance Ratio

a natural disaster [32]. Also, rapid urbanisation in some regions may have caused faster than usual changes in the physical world, with consequent maintenance activity taking place on the digital map. Furthermore, in countries with high map maturity, maintenance may not simply be a consequence of refreshing stale information, but may also be triggered by the uptake of add-on services (e.g., location-based recommender systems), that require expanding the existing POI information with semantically richer one.

As different underlying phenomena may trigger maintenance work, we performed an exploratory analysis to observe possibly different forms of maintenance (i.e., add, update or removal tags), when such practice does take place. To this purpose, we define a new metric, *Action Adoption* AA^{act} as the number of maintenance edits of type $act \in \{add, update, remove\}$ over the total number of maintenance edits that occurred in a country. Formally:

Definition 2 Let OSM_m be the set of OSM edits devoted to maintaining existing POIs, let $OSM_m^{act} \subseteq OSM_m$ be the set of OSM edits of action $act \in \{add, update, remove\}$ over the initial set OSM_m , then Action Adoption is defined as $AA^{act} = \frac{|OSM_m^{act}|}{|OSM_m|}$.

$AA^{act} \in [0, 1]$ by definition; a value of such metric close to 1 means that almost all maintenance edits are of action act , and vice versa. Table 3 illustrates quartiles of the Action Adoption AA^{act} metric, $act \in \{add, update, remove\}$, binned over the 117 analysed countries. Results across all quartiles show that the add action (i.e., enriching existing information) is the most common one, usually more common than the update and remove actions combined (i.e., correcting existing information); furthermore, when correcting existing information, it is usually the case of updating an existing tag, rather than removing any. While this is the most common way of performing maintenance across the countries under exam, we also observe a big variance of the AA^{act} metric between the first and fourth quartiles, in each row (action) of Table 3 (i.e., AA^{add} ranges from 0.14 to 0.88, AA^{update} ranges from 0.07 to 0.75, and $AA^{removal}$ from 0.03 to 0.47). This suggests that there exist countries that do not follow the previously mentioned pattern. Indeed, we performed a manual investigation of some such cases and found that, for example, in Haiti, Turkey, and Niger the removal of tags is the most frequent maintenance practice performed; in Oman, Costa Rica

	Min	1st Qu.	Median	3rd Qu.	Max.	Freq. Distr.
AA^{add}	0.14	0.42	0.52	0.59	0.88	
AA^{update}	0.07	0.24	0.32	0.38	0.75	
AA^{remove}	0.03	0.08	0.11	0.18	0.47	

Table 3. Summary Statistics of Action Adoption in the 117 Analysed Countries, divided by Adds, Updates, Removes

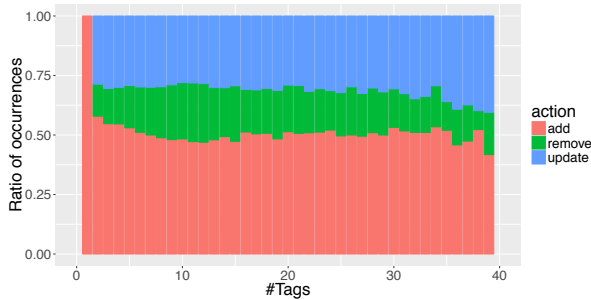


Figure 4. Ratio of Maintenance Actions vs. Number of Tags Associated with OSM POIs

and Azerbaijan the update of tag values is the most frequent action instead.

In an attempt to explain the varying uptake of different forms of maintenance practice, we put forward the following hypothesis:

Hypothesis 2 *The adoption of add, update and removal actions is affected by the information richness of the original OSM POI entry.*

This hypothesis is inspired by the consideration that the add action can be seen as a signal of the user intent to enrich existing information. Such an intent may be driven by the fact that the original entry was not described in enough details. Conversely, an update or a removal action can be seen as signal of the user intent to update existing stale information, or to fix potential mistakes; therefore, these actions may be more likely to take place on OSM POIs that were described by many tags.

To verify our hypothesis, we compute the ratio of times each maintenance action takes place (i.e., add, update, removal) versus the number of tags describing OSM POIs. Results are shown in Figure 4; as expected, we observe that the ‘add’ action is the most frequent one when the number of tags associated with a POI is very low. As the number of tags increases, the ratio of times the ‘add’ action is used falls sharply at first. However, it then stabilises and this action remains the most frequent one even when the number of tags associated with a POI is very high. We speculate that this phenomenon may occur because of the high quality of most of the original entries [11, 13] which do not require any fix. However, in this study we do not investigate the relationship between maintenance practices and map quality, and this result should only be seen as a possible starting point to investigate this important, yet largely unexplored, research direction.

RQ2 – What information is being maintained?

We continued the analysis by investigating what type of information is being maintained.

Hypothesis 3 *By their own nature, certain types of POIs (e.g., a restaurant, a cafe) are maintained more often than others (e.g., a city hall, an hospital).*

To validate this hypothesis, we grouped POIs in each country according to their amenity type (e.g., restaurant, school, hospital, etc.). For each POI type, we have then computed the corresponding Maintenance Ratio (MR) metric. We found a very skewed distribution in each country, with a minority of POI types (less than 10% of all types) being frequently maintained, and several hundreds of POI types receiving near zero maintenance instead. We then looked more closely into the frequently maintained ones, to see if there were commonalities among the analysed countries. Surprisingly, we found almost no overlap, with each country having a distinct set of POI types it maintains. This may either suggest that each country has interest in maintaining different spatial information, or that the same spatial object is being described using different terms in different countries. The latter is possible because, although OSM guidelines suggest a world-wide common taxonomy of amenity types to use,¹¹ in practice mappers are free to use whatever vocabulary they prefer. To find out which of the two is correct, we compared the taxonomy used by each country to describe its most edited POI types with the official OSM taxonomy, and found that more than 90% of the terms used do belong to it. We then restricted our analysis to the POI types that use this official taxonomy (thus dropping the 10% that use non standard terms), looked for commonalities among the 117 countries under exam, and still found almost no overlap. For example, in the Netherlands, the most maintained POI types are ‘restaurant’, ‘cafe’ and ‘place_of_worship’, while in Russia the most maintained ones are ‘clinic’, ‘dentist’ and ‘public_building’. Although we do not know the cause, this result signals that our initial hypothesis, suggesting that some objects require less maintenance than others by their very own nature, is only confirmed when we analyse each country on its own; in fact, different countries maintain distinct types of spatial information.

We then moved our attention from the types of spatial objects that are being maintained, to the set of tags that are being maintained, regardless of the POI type they refer to, and formulated the following hypothesis:

Hypothesis 4 *Certain types of OSM tags (e.g., ‘addr:street’, ‘addr:housenumber’) are more frequently edited during maintenance practices than during the creation of a new OSM object.*

To verify this hypothesis, we define *Tag Adoption* as the ratio of the number of times tag t has been used for a certain action act , over the total number of times action act occurred. Formally:

Definition 3 *Let $adoption_t^{act}$ be the number of times tag t has been used for a certain action $act \in \{add, update, remove\}$, let $|OSM_m^{act}|$ the total number of*

¹¹http://wiki.openstreetmap.org/wiki/Map_Features#Amenity

times action act occurred, then we define Tag Adoption as

$$TA_t^{act} = \frac{adoption_t^{act}}{|OSM_m^{act}|}$$

$TA_t^{act} \in [0, 1]$ by definition; high values of TA_t^{act} indicate that tag t has been frequently used when act took place (e.g., tag ‘addr:street’ has been frequently used during an ‘add’ maintenance practice), and vice versa.

We have computed Tag Adoption TA_t^{act} in each of the 117 countries under exams, for different maintenance actions $act \in \{add, update, remove\}$. We also computed this metric for the ‘creation’ action (i.e., when a POI is added to the map for the first time), to serve as a baseline. As for the case of POI types, in each country and for each action, we found a very skewed distribution, with only a minority of tags (less than 5%) being frequently edited. However, contrary to what we found for POI types, when zooming into this group of frequently edited tags, we found significant overlaps across countries. Table 4 reports both the name of the tags most frequently edited, and the number of countries in which such tags appeared in the list of the 5% most edited ones; the table further distinguishes between a creation edit (top left of Table 4) and the three different types of maintenance edits (add, update, remove). For readability, only the top ten most globally adopted tags for each action are reported.

<i>Creation</i>		<i>Maintenance</i>	
<i>Adding a tag</i>		<i>Adding a tag</i>	
Tag	# Countries	Tag	# Countries
name	117	name	108
amentiy	114	addr:street	44
place	78	addr:city	43
shop	64	wikipedia	31
addr:street	42	addr:housenumber	29
source	38	name:en	27
addr:city	31	addr:postcode	26
highway	22	source	24
addr:housenumber	21	operator	24
natural	16	name:ru	24
<i>Maintenance</i>			
<i>Updating a tag</i>		<i>Removing a tag</i>	
Tag	# Countries	Tag	# Countries
name	117	name	79
place	106	amenity	70
amenity	78	source	28
opening_hours	52	fixme	27
wikipedia	45	highway	26
shop	39	place	24
addr:street	31	building	21
source	30	note	18
name:en	27	is_in	16
website	24	wikipedia:en	16

Table 4. Top Ten Globally Adopted Tags for Each Action

As hypothesised, when a new POI is added to the base map, tags ‘name’ and ‘amenity’ are almost always filled in; conversely, when maintenance actions are performed, then other tags are involved. Although we cannot be sure of the rationale for these tags to be globally maintained, we can draw some interesting observations. Let us consider each maintenance action in turn, starting with the *addition* of tags (which, as seen before, is by far the most frequently performed maintenance practice worldwide). Aside from adding names to POIs that did not have one before, this practice seems

to focus on address details (e.g. ‘addr:street’, ‘addr:city’, ‘addr:housenumber’, ‘addr:postcode’). This corroborates the intuition that information maintenance is often subject to external drivers, such as the integration of location-based services over the base map,¹² which do require address information to operate effectively.

Let us consider the *update* action next. We previously hypothesised that this maintenance action is a signal of the user intent to either update existing stale information, or to fix potential mistakes. To support the validity of these hypotheses, we have then focused on the two most updated tags worldwide (i.e., the ‘name’ of a POI, and its ‘place’) and analysed how the values of these tags have been changed. For the tag ‘name’, we computed the Levenshtein distance between its value before and after the update, to measure the extent to which the initial tag value had been altered. We found that, in most cases, the Levenshtein distance was equal to the unit, meaning that only one character had been changed (e.g., the name of a coffee shop had been changed from “Tim Horton’s” into “Tim Hortons”). This interesting finding corroborates the intuition that most ‘update’ work aims to fix small spelling mistakes in the names of already existing POIs. More rare is the action of changing the name of a POI into a complete different name (high values of Levenshtein string distance), suggesting that the OSM map does not contain much stale information. For the tag ‘place’, the Levenshtein distance between the pre- and post- update values was very high instead (e.g., from ‘suburb’ to ‘neighbourhood’, or from ‘neighbourhood’ to ‘quarter’). To understand the semantics behind this kind of changes, we examined the OSM wiki page which describes the tag ‘place’;¹³ we learned that the OSM community suggests using values for this tag taken from a specific taxonomy rather than free text; the taxonomy is described in details, suggesting what terms can be interchangeably used, and what spatial relationship exists between groups of equivalent terms (e.g., the value ‘neighbourhood’ indicates a particular location that is generally smaller than ‘suburb’ and ‘quarter’). To automatically measure the extent to which the semantics of the tag ‘place’ had been altered after each update, we grouped all its possible values into five different semantic classes, as suggested in the OSM wiki page. Then, we coded the change that occurred during the update using a binary value: 1 when the new tag value belonged to a different semantic class than the old one, and 0 otherwise. The former suggests a major semantic change had taken place (e.g., ‘neighbourhood’ changed into ‘town’); the latter suggests a smaller refinement instead (e.g., ‘neighbourhood’ changed into ‘quarter’). We found that, in over 90% of cases, the changes were of the latter type. Similar to what we found when analysing updates on tag ‘name’, most of the times the action ‘update’ leads to small refinements of the spatial information already contained in OSM, and only much less frequently to major corrections of possibly stale or plain wrong information.

We finally consider the ‘remove’ tag maintenance action. In half of the considered countries, the most removed tags are

¹²<https://blog.openstreetmap.org/2015/02/16/routing-on-openstreetmap-org/>

¹³<http://wiki.openstreetmap.org/wiki/Key:place>

‘name’ and ‘amenity’. We did some manual inspection to understand in what cases these tags were removed; we discovered that, in all the considered cases, these tags had been removed because they were improperly used and did not follow the official guidelines provided by the OSM community¹⁴ (e.g., the tag ‘name=bar’ was used instead of the more appropriate tag ‘amenity=bar’). This result confirms what we previously hypothesised, that is, this type of maintenance action is a signal of the user intent to polish existing OSM information in order to be more compliant with the official OSM taxonomy.

RQ3 – Who engages in information maintenance?

We now move our attention from what information is being maintained to who takes charge of performing maintenance work. Previous studies of OSM have shown that there exists a small set of highly engaged (expert) users who are responsible for the majority of the mapping [35]; we so make the following hypothesis.

Hypothesis 5 *Maintenance practices are performed especially by highly engaged and expert users.*

One might expect this to be the case for various reasons, ranging from motivation (i.e., the same drivers that make them map extensively may also drive them to maintain extensively), to knowledge (e.g., having previously contributed a lot of information, they might know what information is most stale and in need of updates), to skills (e.g., updating existing information may require users to have acquired a certain skill-set first, as was observed in other crowd sourcing communities like Wikipedia [26, 33, 40]).

To verify this hypothesis, we first grouped users within each country into five different classes of engagement (or expertise). We measured user’s engagement using two alternative proxies: (i) *NumEdts*, that is their total number of OSM edits; and (ii) *ActDays*, that is the number of days during which they performed OSM editing activity. Summary statistics of the number of users per each class, across all countries, are reported in Table 5. In the following, since results are very similar when using either of the two proxies, we only report results obtained with the *ActDays* one.

<i>NumEdts</i>		<i>ActDays</i>	
Class	#Users	Class	#Users
(0,1]	25,235	(0,1]	50,177
(1,10]	36,295	(1,5]	20,442
(10,100]	15,335	(5,10]	4,074
(100,1k]	3,927	(10,100]	6,100
(1k,10k]	606	(100,1k]	605

Table 5. Summary Statistics of Classes of Users

We computed the Maintenance Ratio (*MR*) metric defined before, but on a per user basis rather than on a per country basis. Figure 5 shows the quartiles (in yellow) and the frequency distributions (in green) of *MR* for each class of users. As hypothesised, the more experienced the users are, the more effort they devote to maintaining existing POIs compared to the effort they spend to edit new ones. Interestingly,




Figure 5. ActDays Vs. Maintenance Ratio

Action	Min	1st Qu.	Median	3rd Qu.	Max
(all)	-0.01 **	0.13 **	0.17 ***	0.22 ***	0.33
Add	0.00 **	0.21 ***	0.26 ***	0.32 ***	0.44
Update	0.11 **	0.29 ***	0.34 ***	0.39 ***	0.46
Remove	0.17 ***	0.28 ***	0.33 ***	0.41 ***	0.54

Table 6. Summary statistics of the Spearman correlations between user expertise (measured as the total number of active days in editing OSM) and *MR*, binned for the 117 countries under exam. Stars indicate the percentage of correlations that are statistically significant within each quartile (p -values < 0.01): 0% ‘ ‘ 25% ‘*’ 50% ‘**’ 75% ‘***’ 100%

a similar behaviour can be found in other crowd-based knowledge production platforms such as Wikipedia. Indeed, previous research has found that the most engaged Wikipedia editors are those who revise available entries [33], making them both more accurate and complete.

To see whether this result is consistent across all 117 countries under exam, we zoomed-in and computed, within each country, the Spearman correlation between a user *MR* and its engagement level (measured with the proxies defined above). The first row of Table 6 reports the computed correlations, divided in quartiles, when the proxy *ActDays* was used. As shown, we found statistically significant positive correlations between our proxy of user expertise and Maintenance Ratio *MR* in three quarters of the analysed countries – i.e., the Spearman correlation ρ between the user *ActDays* and their *MR* ranges from 0.13 to 0.33 in three quarters of the analysed countries. This suggests that the overall behaviour showed in Figure 5 is taking place globally.

For completeness, we further computed the above correlations on a per action basis. Results are shown in Table 6 (rows 2–4) and they confirm that, in almost all countries under exam, experienced users are those who devote more effort to maintaining existing POIs, compared to less experienced ones. This behaviour is even more pronounced for ‘update’ and ‘remove’ actions, where the correlations are stronger.

RQ4 – Where is maintenance taking place?

We next analyse where maintenance is taking place. Previous research by Muki *et al.* [14] has found that there is a strong correlation between the quality of the data contained in OSM, and the number of editors which operate in a given area. This preliminary finding drove us to hypothesise the following:

Hypothesis 6 *The higher the number of OSM editors in an area, the higher the amount of maintenance work that takes place there.*

¹⁴<http://wiki.openstreetmap.org/wiki>

To validate this hypothesis, we sub-divided each country in grids of different sizes (i.e., 50km × 50km and 100km × 100km). Then, for each cell of the grid, we computed the corresponding level of maintenance ratio MR and the number of OSM contributors editing it. Finally, for each country, we measured the correlation between these two values. Since we are dealing with geographic data, we had to address the problem of spatial auto-correlation, which we found to be indeed high in several countries. To overcome this problem, we used the method introduced by Clifford *et al.* [5] to address auto-correlation in spatial datasets. The first row of Table 7 reports summary statistics of the correlations for the 117 different countries under exam, when using a 100km × 100km grid. Similar results were obtained when adopting a 50km × 50km grid and are thus omitted.

	Min	1st Qu.	Median	3rd Qu.	Max
# All Users	-0.37 *	-0.12	0.09 *	0.32 ***	0.75
# Experienced Users	-0.28	-0.06 *	0.20 **	0.32 ***	0.77

Table 7. Summary statistics of the Clifford correlations between number of contributors and MR , binned for the 117 countries under exam. Stars indicate the percentage of correlations that are statistically significant within each quartile (p -values < 0.01): 0% ‘ ’ 25% ‘*’ 50% ‘’ 75% ‘***’ 100%**

As shown, in around half of the analysed countries, there is a statistically significant positive correlation between the number of contributors and Maintenance Ratio, thus supporting our hypothesis. However, we have also found an equal number of countries where our hypothesis does not hold. We have attributed this last result to the fact that, as previously observed, not all users undertake maintenance practices; rather, maintenance seems to be predominantly performed by experienced users only. Previous research also shows that experienced editors in OSM (also called ‘power-users’) do not operate everywhere within a country; rather, they often concentrate their mapping activities in few specific territories/areas [35]. We thus corrected our initial conjecture and instead hypothesised the following:

Hypothesis 7 *The higher the number of ‘experienced’ OSM editors in an area, the higher the amount of maintenance work that takes place.*

To verify this, we calculated the number of experienced users editing each cell of the grid, then correlated this value with the maintenance ratio MR associated with the same cell. To this purpose, we classified OSM editors as ‘experienced’ if they have mapped OSM for more than 10 days (the last two groups shown in Table 5). The second row of Table 7 reports summary statistics of the obtained correlations for the 117 countries under exam. Interestingly, in this second case we found both higher and more statistically significant correlations than the previous case where any OSM editor was taken into account. This result suggests that spatial information is being maintained in areas where many experienced users operate. For example, in the UK, these areas are around major urban areas (i.e., London, Liverpool and Manchester).

RQ5 – What are the factors that trigger maintenance?

The statistical analysis conducted so far tells us that experienced OSM mappers are those undertaking the majority of

Q1. In what country do you live?
Q2. In what country do you usually edit objects in OSM?
Q3. Approximately how many edits have you made in OSM?
Q4. How often do you rely on OSM as a map?
Q5. What tool do you mainly use to edit objects in OSM?
Q6. Have you ever maintained existing objects in OSM, for example to correct mistakes or to add more information?
If the answer to Q6 is ‘Yes’, answer Q7-Q9
Q7. How do you identify what OSM objects need maintenance?
Q8. Please give us an example of maintenance work you did, and why you did it
Q9. Have you used special tools to maintain existing objects in OSM?
If the answer to Q6 is ‘No’, answer Q10
Q10. Why have you never maintained an object already existing in OSM?

Table 8. Survey Questions

maintenance work. Despite different countries being engaged in maintenance activity at different levels, and targeting different types of POIs, we found strong commonalities in terms of how information on existing spatial objects is being maintained (i.e., using what action, on what tags). What we do not know yet is what triggers a maintenance action, that is, how do these experienced mappers know that maintenance work is required?

In an attempt to answer this last question, we created an online survey that asked OSM contributors the following questions: first, we asked some background information, in terms of the country they live in, the country they most frequently map, the tools they use for mapping, and approximately how many edits they have done so far (we provided a link to a tool to automatically extract this information from OSM, without them having to reveal us their username). Multiple-choice answers were provided to minimise efforts when filling in this part of the survey. Second, we asked respondents whether they had ever done any maintenance work in OSM. If not, we asked why; if yes, we asked for examples of maintenance work they had done. Free-text entries were used in this case. The full list of questions we asked is showed in Table 8.

We conducted this online survey during the months of April and May 2015, by sending a link to it to 205 OSM contributors via OSM direct messages. Information about who edits in OSM is publicly available; we randomly selected who to contact, so to cover the whole spectrum of OSM users based on: (i) their country of residence, which should be representative of the countries where OSM contributors operate, and (ii) their level of experience, based on past editing activity, as defined in Table 5. Of the 205 OSM mappers we contacted, 96 took part in our survey. Table 9 shows how answers were distributed among the respondents.

Respondents

Figure 6 depicts the number of survey participants (y axis) over the number of OSM contributors (x axis) per country. As showed, the distribution of participants’ country of residence is fairly representative of the countries where OSM mappers operate, with countries such as Germany, United States of

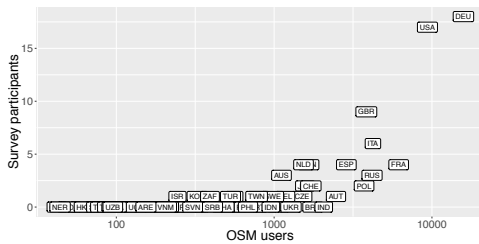


Figure 6. Number of Survey Participants and OSM Users per Country

America, United Kingdom, Italy and France being heavily represented.

Unsurprisingly, we did not get answers uniformly distributed across all user bins of user expertise (see distribution of Q3 in Table 9); rather, the majority of users who participated in our study were highly engaged OSM editors, with experience of both creating and maintaining information in OSM. Although not representative of the whole OSM contributor base in terms of user expertise, we next elaborate on the themes that emerged from their answers, as they still offer important insights that can be leveraged to understand what are the factors that trigger OSM contributors (or, at least, the type of contributors who answered our survey) to maintain OSM information.

Triggers

When asked how they identify what objects need maintenance (Q7), and to provide examples of maintenance actions they had done (Q8), the majority (45 out of 96 participants – see Table 10) reported engaging in such practice as a consequence of stumbling upon incorrect information, while using OSM as a map service within third-party applications (mostly navigation services). For example, respondents wrote: “I always use OSM based maps for driving and regularly review my tracks, if there is a significant discrepancy in map vs. track I correct it”; “When I hike, I compare the world with the map on my phone”, and “I update usually only after I see that there is a difference between maps and real world”. Some respondents (15 out of 96 – see Table 10) reported to pro-actively check whether OSM map information was correct and up-to-date instead. If mistakes were found, they then engaged in maintenance activity, with the aim to then experience a good quality-of-service in the companion applications and services they use. For example, a user wrote “I update the map [...] going around the city and comparing what I see with the OSM map”).

In both cases, going online and maintaining OSM spatial objects seemed to be triggered by experiencing incorrect information *offline* (i.e., in the physical world). There were also participants (28 out of 96 – see Table 10) who declared that maintenance was triggered by *online* events instead: “I notice things that are flagged by various Q&A tools”, “Usually I use *osmose.openstreetmap.fr* which suggests me what may need to be revised”. One participant reported that she actively reported in some local OSM communities what needed to be updated in her region: “A few weeks ago I upgraded pharmacies in my region adding the tag ‘dispensing’ which

Q1. In what country do you live?	
Germany	18
United States of America	17
United Kingdom	9
Italy	6
Canada	4
France	4
Netherlands	4
Spain	4
Australia	3
Russia	3
Japan	2
Poland	2
Switzerland	2
Others	18
Q2. In what country do you usually edit objects in OSM?	
Mainly in the same country where I live	85
Elsewhere	11
Q3. Approximately how many edits have you made in OSM?	
Less than 10	2
Between 10 and 100	18
Between 100 and 1000	44
More than 1000	32
Q4. How often do you rely on OSM as a map	
Never. I always use Google Maps or other non-OSM maps	1
Sometime. Once a month or even less	41
Very often. At least once a week	54
Q5. What tool do you mainly use to edit objects in OSM?	
OSM official website (iD)	54
Third party applications (e.g., JOSM)	36
Mobile editors	3
Other	3
Q6. Have you ever maintained existing objects in OSM?	
Yes	95
No	1
Q7. How do you identify what OSM objects need maintenance?	
<i>[Some extracts]</i>	
“I check the map against what I have seen with my own eyes. If information is missing or incorrect, I’ll edit.”	
“By explore the map for things that I know best and detect if there are some issue on the map.”	
Q8. Please give us an example of maintenance work you did, and why you did it	
<i>[Some extracts]</i>	
“I always use OSM based maps for driving and regularly review my tracks, if there is a significant discrepancy in map vs. track I correct it.”	
“Completing info about shops, when cycling it’s good to know where you could find repair-shop, or café/restaurant to drink/eat something.”	
Q9. Have you used special tools to maintain existing objects in OSM?	
I used my usual editor	79
I used other apps (e.g., WheelMap, My Opening Hours)	6
I played Mapping games (e.g., Kort, MapRoulette)	6
Others	4
Q10. Why have you never maintained an object already existing in OSM?	
I would have liked to, but I found it difficult to identify what needed to be updated	1

Table 9. Answer Distributions of the Users who Participated in our Study

<i>Main factors that trigger maintenance</i>	
Experiencing incorrect information while using OSM	45
Noticing errors flagged by various Q&A tools	28
Pro-actively checking the OSM information	15
Others	8

Table 10. Distributions of the Main Factors that Trigger Maintenance

distinguishes them from drugstores. I did it because it is important for those who want to create services based on OSM. I reported this fact on various local OSM communities asking the help of other editors”.

We note that only a minority of participants in our survey referred to issues reported in Q&A and group discussion tools as triggers for maintenance. It is worth investigating the rationale for this in the future, as previous research in other crowd-based knowledge production services like Wikipedia has actually revealed that experienced Wikipedia editors regularly engage with Q&A and group discussion tools, and use them to inform how to correct existing entries [33, 15].

Tools

Finally, we investigated whether OSM contributors who perform maintenance use dedicated tools for this task. Indeed, a variety of tools have been developed to support this specific practice. Examples include: ‘mapping games’, such as Kort¹⁵ and MapRoulette,¹⁶ that organise challenges dedicated to correcting potential errors in OSM (e.g., “Welcome to MapRoulette! You will be working on this challenge: Italian Wrong Addresses”); and ‘validators’, such as Osmose,¹⁷ that aim to automatically detect potential issues in OSM data (e.g., inconsistent tag usage within the same POI), and expose such inconsistencies to OSM contributors for manual resolution.

We asked participants of our study whether they use special tools to perform information maintenance in OSM (Q9). The majority (79 out of 96 participants) declared that they do not; rather, information maintenance is a practice they perform within their normal OSM editing tool (with OSM official tool iD¹⁸ and third party tool JOSM¹⁹ being the most commonly adopted ones). Among the tools explicitly dedicated to supporting maintenance work, only validators (and Osmose in particular) were mentioned by our respondents. This might be linked to the fact that users currently engaged in maintenance work are overall very active mappers already, so their motivation to curate OSM information is probably intrinsically high, and does not need to be stirred by external incentives like mapping games.

DISCUSSION

From a theoretical perspective, this work has presented a method to make visible the otherwise hidden maintenance practices of self-organised communities of practice interested in gathering and maintaining geographic knowledge. We

¹⁵<http://www.kort.ch/>

¹⁶<http://maproulette.org/>

¹⁷<http://osmose.openstreetmap.fr/en/map/>

¹⁸<http://ideditor.com/>

¹⁹<https://josm.openstreetmap.de/>

have applied this method to a specific community and data type (OpenStreetMap and its POIs). However, we believe the same method can be used to study other data types within OpenStreetMap (e.g., ways and relations), as well as other crowd-mapping platforms, such as CrowdMap and FixMyStreet, for comparative studies. The method can also be reapplied to the same community and data type over time, in order to capture changes in behaviour, for example, as might be induced by major updates of the tools offered to support this practice.

From a practical perspective, our findings highlight opportunities for new research studies and for the development of more specialised tools and work-flows to support maintenance practices. For example:

RQ1 (Spread) – We found that, in most of the analysed countries, the action to add new tags to existing spatial objects is more frequent than the updating or the removal of existing tags. This finding calls for interesting follow up questions: are ‘update’ and ‘removal’ actions rare because of the high quality of the original OSM entries? Or are they rare because current tools do not help users identify what OSM information needs to be rectified? New research studies are required to shed light onto these questions.

RQ2 (What) – We found that the most commonly added/updated/removed tags are the same across many countries. This could be a consequence of mainstream editors – such as iD and Potlatch 2²⁰ – giving users the option to select tags from a pre-filled list which reflects the official OSM taxonomy, rather than encouraging them to create their own tags by means of, for example, a free-form text box. However, based on our results, current tools could be further improved. For example, our analysis reveals which tags are most commonly edited on a per action (‘add’, ‘update’, ‘remove’) basis; tool developers may explore how to build tools that recommend the appropriate tags to edit, when users engage in these actions.

RQ3 (Who) – We found that maintenance actions are prevalently done by experienced users. From our survey, the only participant who had never done information maintenance reported that she “would have liked to perform it, but found it difficult to identify what needed to be maintained”. Designing tools that support non expert users in these practices is thus a direction worth investigating. New follow up studies are needed to understand how to effectively involve and facilitate this class of users in undertaking maintenance work.

RQ4 (Where) – We found that the higher the number of ‘experienced’ OSM editors in an area, the higher the amount of maintenance work that takes place. In order to maintain information also in areas where experienced mappers do not naturally attend to, new techniques need to be investigated to focus mappers’ attention. Previous research on the Cyclopath geo-wiki of bike routes in Minnesota has showed that simply highlighting on the map those areas

²⁰<https://www.openstreetmap.org/login?referer=%2Fedit%3Feditor%3Dpotlatch2>

that have not been edited in a long while was sufficient to steer the community to go check and edit in those areas [34]. Similar techniques could be used, for example, to automatically highlight in OSM areas that have undergone fast urbanisation, areas affected by natural disasters, and so on. However, Cyclopath operates on a relatively small geographic area (a city), within a very specific community of interest (cyclists), where one may assume social capital to be high. It is an open question whether map highlights are sufficient to motivate and steer map maintenance in broader and more ‘general purpose’ geographic communities like OSM.

RQ5 (Triggers) – Most of our survey participants reported to engage in maintenance practice as a consequence of stumbling upon incorrect information while using OpenStreetMap. However, performing information maintenance mainly as a consequence of finding mistakes in the field may result in severe limitations in terms of coverage: the OSM objects located in the areas physically covered by active mappers may naturally end up being maintained, while all others may quickly fall by the wayside. This trend is supported by previous results showing that, in many of the analysed countries, the maintained OSM objects are mainly located in areas covered by experienced users. New tool functionalities need to be explored to overcome this limitation; as an example, approaches worth investigating are: (i) the (semi) automatic assignment of maintenance tasks to contributors, as done in other online crowdsourcing communities [6, 17, 25, 38], and (ii) incentive strategies, as used in crowd-sensing communities to encourage users to share their data [19, 42, 41].

CONCLUSION, LIMITATIONS AND FUTURE WORK

In this paper, we have proposed a method to quantify maintenance work carried out in knowledge production communities, where knowledge has a distinct spatial natural, and where it naturally evolves over time. We have applied our method and metrics to OpenStreetMap in particular, one of the most successful examples of geographic crowd-sourced datasets.

The results presented in this study suffer from some limitations that leave open opportunities to explore. First, we would like to point out that, with this study, we aimed to elicit maintenance practices as they naturally emerge in self-organised communities like OSM. What we did not do is to relate these practices to metrics of map quality. For example, we made no assumption about the number of spatial objects that need to be maintained in each area, and the absence of maintenance that we have quantitatively observed in some regions may simply be due to the fact that there are no spatial objects that need to be maintained there. Investigating the relationship between maintenance practices and map quality is an important research direction that we aim to explore in the future.

Another limitation of this work is that all results have been computed at country level. A more fine-grained study, for example at regional or city level, could be able to reveal more localised mapping dynamics than those we detected with this

work. Also, this work has focused on one OSM data object, that is OSM POIs. If other OSM data types were considered (e.g., OSM roads), different maintenance practices may emerge. We leave both investigations to future work; however, we note that the applicability of the method proposed in this paper withstands.

REFERENCES

1. Jennings Anderson, Robert Soden, Kenneth M Anderson, Marina Kogan, and Leysia Palen. 2016. EPIC-OSM: A Software Framework for OpenStreetMap Data Analytics. In *2016 49th Hawaii International Conference on System Sciences (HICSS)*. IEEE, 5468–5477.
2. J.J. Arsanjani, C. Barron, M. Bakillah, and M. Helbich. 2013. Assessing the Quality of OpenStreetMap Contributors together with their Contributions. In *Proc. of AGILE*.
3. Daren Brabham. 2013. *Crowdsourcing*. MIT Press.
4. Virginia Braun and Victoria Clarke. 2006. Using thematic analysis in psychology. *Qualitative Research in Psychology* 3, 2 (2006), 77–101.
5. P. Clifford, S. Richardson, and D. Hemon. 1989. Assessing the significance of the correlation between two spatial processes. *Biometrics* 45, 1 (1989), 123–134.
6. Djellel Eddine Difallah, Gianluca Demartini, and Philippe Cudré-Mauroux. 2013. Pick-a-crowd: tell me what you like, and i’ll tell you what to do. In *Proceedings of the 22nd international conference on World Wide Web*. ACM, 367–374.
7. Martin Dittus, Giovanni Quattrone, and Licia Capra. 2016. Analysing volunteer engagement in humanitarian mapping: building contributor communities at large scale. In *Proc. of CSCW*. ACM.
8. H. Fana, A. Zipfa, Q. Fub, and P. Neisa. 2014. Quality assessment for building footprints data on OpenStreetMap. *International Journal of Geographical Information Science (IJGIS)* 28, 4 (2014), 700–719.
9. Michela Ferron and Paolo Massa. 2011a. Collective memory building in Wikipedia: the case of North African uprisings. In *Proceedings of the 7th International Symposium on Wikis and Open Collaboration*. ACM, 114–123.
10. Michela Ferron and Paolo Massa. 2011b. Studying collective memories in Wikipedia. *Journal of Social Theory* 3, 4 (2011), 449–466.
11. J.F. Girres and G. Touya. 2010. Quality assessment of the French OpenStreetMap dataset. *Transactions in GIS* 14, 4 (2010), 435–459.
12. M.F. Goodchild. 2007. Citizens as Sensors: the World of Volunteered Geography. *GeoJournal* 69, 4 (2007), 211–221.

13. M. Haklay. 2010. How good is volunteered geographical information? A comparative study of OpenStreetMap and Ordnance Survey datasets. *Environment and Planning B: Planning and Design* 37, 4 (2010), 682–703.
14. M.M. Haklay, S. Basiouka, V. Antoniou, and A. Ather. 2010. How Many Volunteers Does it Take to Map an Area Well? The Validity of Linus Law to Volunteered Geographic Information. *The Cartographic Journal* 47, 4 (2010), 315–322.
15. A. Halfaker, R.S. Geiger, J. Morgan, and J. Riedl. 2013. The Rise and Decline of an Open Collaboration System: How Wikipedia’s reaction to sudden popularity is causing its decline. *American Behavioral Scientist* 57, 5 (2013), 664–688.
16. Jeff Howe. 2006. The Rise of Crowdsourcing. *Wired* (2006).
17. Panagiotis G Ipeirotis and Evgeniy Gabrilovich. 2014. Quiz: Targeted crowdsourcing with a billion (potential) users. In *Proceedings of the 23rd international conference on World Wide Web*. ACM, 143–154.
18. K. Ishida. 2012. Geographical Bias on Social Media and Geo-local Contents System with Mobile Devices. In *Proc. of HICSS*. 1790–1796.
19. Luis G Jaimes, Idalides J Vergara-Laurens, and Andrew Raji. 2015. A survey of incentive techniques for mobile crowd sensing. *IEEE Internet of Things Journal* 2, 5 (2015), 370–380.
20. Andreas Kaltenbrunner and David Laniado. 2012. There is no deadline: time evolution of Wikipedia discussions. In *Proceedings of the Eighth Annual International Symposium on Wikis and Open Collaboration*. ACM, 6.
21. Brian Keegan, Darren Gergle, and Noshir Contractor. 2012. Do editors or articles drive collaboration?: multilevel statistical network analysis of wikipedia coauthorship. In *Proceedings of the ACM 2012 conference on computer supported cooperative work*. ACM, 427–436.
22. Aniket Kittur, Bongwon Suh, Bryan A Pendleton, and Ed H Chi. 2007. He says, she says: conflict and coordination in Wikipedia. In *Proceedings of the SIGCHI conference on Human factors in computing systems*. ACM, 453–462.
23. Marina Kogan, Jennings Anderson, Robert Soden, Kenneth M Anderson, and Leysia Palen. 2016. Collaboration in OpenStreetMap: A Network Analysis of Two Humanitarian Events. In *Proc. of CHI*.
24. David Laniado and Riccardo Tasso. 2011. Co-authorship 2.0: Patterns of collaboration in Wikipedia. In *Proceedings of the 22nd ACM conference on Hypertext and hypermedia*. ACM, 201–210.
25. Hongwei Li, Bo Zhao, and Ariel Fuxman. 2014. The wisdom of minority: Discovering and targeting the right group of workers for crowdsourcing. In *Proceedings of the 23rd international conference on World Wide Web*. ACM, 165–176.
26. Jun Liu and Sudha Ram. 2011. Who does what: Collaboration patterns in the wikipedia and their impact on article quality. *ACM Transactions on Management Information Systems (TMIS)* 2, 2 (2011), 11.
27. I. Ludwig, A. Voss, and M. Krause-Traudes. 2011. A Comparison of the Street Networks of Navteq and OSM in Germany. *Advancing Geoinformation Science for a Changing World* 1, 2 (2011), 65–84.
28. Afra Mashhadi, Giovanni Quattrone, and Licia Capra. 2013. Putting ubiquitous crowd-sourcing into context. In *Proceedings of the 2013 conference on Computer supported cooperative work*. ACM, 611–622.
29. Afra Mashhadi, Giovanni Quattrone, Licia Capra, and Peter Mooney. 2012. On the accuracy of urban crowd-sourcing for maintaining large-scale geospatial databases. In *Proceedings of the Eighth Annual International Symposium on Wikis and Open Collaboration*. ACM, 15.
30. P. Mooney and P. Corcoran. 2013. Analysis of Interaction and Co-editing Patterns amongst OpenStreetMap Contributors. *Transactions in GIS 2013* (2013).
31. J.L. Myers and A.D. Well. 2003. *Research Design and Statistical Analysis (2nd ed.)*. Routledge. 508 pages.
32. Leysia Palen, Robert Soden, T Jennings Anderson, and Mario Barrenechea. 2015. Success & scale in a data-producing organization: The socio-technical evolution of OpenStreetMap in response to humanitarian events. In *Proceedings of the 33rd annual ACM conference on human factors in computing systems*. ACM, 4113–4122.
33. Katherine Panciera, Aaron Halfaker, and Loren Terveen. 2009. Wikipedians are born, not made: a study of power editors on Wikipedia. In *Proceedings of the ACM 2009 international conference on Supporting group work*. ACM, 51–60.
34. Reid Priedhorsky, Mikhil Masli, and Loren Terveen. 2010. Eliciting and focusing geographic volunteer work. In *Proceedings of the 2010 ACM conference on Computer supported cooperative work*. ACM, 61–70.
35. Giovanni Quattrone, Licia Capra, and Pasquale De Meo. 2015. There’s no such thing as the perfect map: Quantifying bias in spatial crowd-sourcing datasets. In *Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work & Social Computing*. ACM, 1021–1032.
36. G. Quattrone, A. Mashhadi, and L. Capra. 2014a. Mind the Map: The Impact of Culture and Economic Affluence on Crowd-Mapping Behaviours. In *Proceedings of the 17th ACM conference on Computer supported cooperative work & social computing*. ACM, 934–944.

37. Giovanni Quattrone, Afra Mashhadi, Daniele Quercia, Chris Smith-Clarke, and Licia Capra. 2014b. Modelling growth of urban crowd-sourced information. In *Proceedings of the 7th ACM international conference on Web search and data mining*. ACM, 563–572.
38. Markus Rokicki, Sergej Zerr, and Stefan Siersdorfer. 2015. Groupsourcing: Team Competition Designs for Crowdsourcing. In *Proceedings of the 24th International Conference on World Wide Web*. International World Wide Web Conferences Steering Committee, 906–915.
39. Bongwon Suh, Gregorio Convertino, Ed H Chi, and Peter Pirolli. 2009. The singularity is not near: slowing growth of Wikipedia. In *Proceedings of the 5th International Symposium on Wikis and Open Collaboration*. ACM, 8.
40. Dennis M Wilkinson and Bernardo A Huberman. 2007. Assessing the value of cooperation in wikipedia. *arXiv preprint cs/0702140* (2007).
41. Dejun Yang, Guoliang Xue, Xi Fang, and Jian Tang. 2012. Crowdsourcing to smartphones: incentive mechanism design for mobile phone sensing. In *Proceedings of the 18th annual international conference on Mobile computing and networking*. ACM, 173–184.
42. Dejun Yang, Guoliang Xue, Xi Fang, and Jian Tang. 2015. Incentive mechanisms for crowdsensing: Crowdsourcing with smartphones. (2015).
43. D. Zielstra and A. Zipf. 2010. A Comparative Study of Proprietary Geodata and Volunteered Geographic Information for Germany. In *Proc. of AGILE*.