Analysis of the introduction of a review request mechanism in OSM

Nikola Davidovic
University of Nis, Faculty of Electronic Engineering Nis
Aleksandra Medvedeva 14
Nis, Serbia
nikola.davidovic@elfak.ni.ac.rs

Peter Mooney
Maynooth University
Eolas Building,
North Campus
Maynooth, Ireland
peter.mooney@mu.ie

Leonid Stoimenov
University of Nis, Faculty of Electronic Engineering Nis
Aleksandra Medvedeva 14
Nis, Serbia
leonid.stoimenov@elfak.ni.ac.rs

Abstract

Recently a new option that enables contributors to ask for review of their mapping was introduced into the iD editor in OpenStreetMap (OSM). In this paper we analyse changesets containing this *review_requested=yes* tag. All such changesets that were created in the first two weeks after the option was introduced were analysed. Given that the main review mechanism are comments, written by experienced OSM contributors for the given changeset, the total number of commented changesets of 2.68% was unexpectedly low. This suggests that such reviews were not, immediately, broadly adopted by the OSM community. Moreover, high number of changesets was a part of the HOT-OSM project efforts and only a fraction of them was actually commented by an OSM contributor. We investigate the average types of performed changes, changesets spatial distribution and the used language for the comments. We suggest how changesets could attract proper review and possible improvements to attract wider adoption of the introduced mechanism in OSM.
Keywords: changesets, OSM, review_requested, osmChange.

1 Introduction

Within Volunteered Geographic Information (VGI) OpenStreetMap (OSM) has become a significant subtopic for research and development. Amongst other factors this is in no small part due to its massive adoption among all stakeholders, contributors, software developers but also academics. While OSM already contains huge number of geographic objects OSM contributors are trying to improve the existing data by further editing object geometry or modifying tag descriptions. When volunteers are contributing new data to the OSM, the changes performed by a single contributor over a short period of time are grouped into logical collections called changesets. Each changeset collects and lists all created, modified or deleted elements (objects) or tags of the existing elements. Changesets are then used to update local OSM databases and can serve as incremental differential files (diffs) between two subsequent OSM database states.

For new or novice contributors in OSM the editing process can be challenging and somewhat daunting. Until recently new or novice contributors (or indeed any contributor in OSM) had no mechanism to obtain help or assistance with their mapping effort. In some cases, novice contributors do make significant edits without specific knowledge about the mapping best practices for the objects they are editing. In such cases these edits may be erroneous and require some other OSM contributor to notice them and then "fix" them. In order to overcome this problem, a new tag for OSM changesets called *review_requested*=yes was introduced into the iD OSM editor (Neis, 2017), version 2.4.0, released on 25/08/2017. The idea of this tag was to enable, mostly novice, contributors to set it and thus ask the OSM community for comments on their changesets. These comments would be mostly welcomed from more experienced OSM contributors. The idea of the comments is that they would be used as a guidance or as a verification mechanism for the contributed data. Once comments were provided, the original contributor could then edit some of their data contained in the changeset, based on the received feedback. Clearly the idea for introducing this functionality was to enable inexperienced contributors to receive guidance but also to draw attention of a broader community to some features that required consensus on the mapping approach. Although there is no theoretical limit on the number of changes that a changeset with review request can have, it is expected that such changesets should have a manageable number of changes allowing for a proper review to be performed. Since changesets data are freely available for download, we performed analysis on them.

Based on this changeset review mechanism we have addressed the following research questions in this paper:

- What is the overall volume of these changesets and can the OSM community respond adequately to them?
- Are such changesets constructed in a proper way to be reviewed in a timely fashion?
- Should adding review request to a changeset be somehow constrained in the editor to discourage improper use?

In order to answer these questions, we have conducted an analysis of changesets data between 25/08/2017 and 08/09/2017 corresponding to the first two weeks of the review requested usage. This paper is structured in the following way: after the introduction, a brief literature review is presented with emphasis on mapping procedures and protocols. The third section of the paper describes the analysis process and significant findings within the researched changesets. Finally we conclude the paper with a discussion of our results.

2 Literature Review

In OSM the mapping rules are not strictly imposed. Although it is expected from contributors to review general recommendations for mapping, it is not required that they follow some strict rules or procedures. Nevertheless, there are numerous findings in the literature where data quality, of the VGI and OSM data in some specific areas is of at least similar quality to those produced by NMAs and CSCs (Ludwig et al. 2011, Graser et al. 2014, Neis and Zielstra. 2014). This usually means that in some areas, contributors take both positional and thematic accuracy into consideration. But OSM also shows variations in the level of data quality between different geographical areas. Even though high quality OSM data exist, there are numerous examples where contributors are not following recommendations that can be found on the OSM Wiki pages and in the mailing lists. Such practices contribute to the high inhomogeneity in OSM data. While much attention is drawn to inhomogeneity in the context of positional accuracy attention is also to the context of proper tagging (Davidovic et al. 2016). Therefore, contributions towards superimposing protocols for VGI and OSM data collection, in order to increase collected data quality, have emerged lately (Mooney et al 2016, Minghini et al. 2017). The idea is not only to introduce strict rules for mapping contributions but to make proper trade-off by applying the mapping protocol and at the same time keeping the interesting part of the mapping that attracts people to contribute in the first place.

In the absence of protocols, one of the existing processes that keeps quality in some regions on a high level is the fact that senior contributors are constantly reviewing and improving data. This is essentially the *unofficial loopback mechanism* that has been always used in OSM. It is because in each region there were few senior mappers (Neis, 2014) who, among other things, were reviewing existing data in their surroundings and thus improving their positional and thematic accuracy. But these reviews were happening only in cases when such senior mappers were particularly interested in some specific objects. Large amounts of contributed data can be submitted to OSM without review by senior mappers. It has been discussed for some time in OSM that there is a need for a user driven or automatic data validation technique which triggers a reviewing mechanism of contributed data.

3 Analysis

We have performed the analysis on a large number of changesets. Changesets can be obtained from the (PlanetOSM, 2018). The Latest Weekly Planet XML File contains all the changesets that were made from the day one of the OSM. Changesets within this file are represented in XML format where each changeset can have multiple attributes that define its date and time of creation and closing, user who created it, number of changes that it contains and the affected bounding box. Changesets include tags and, for this research, especially interesting tags are semicolon delimited set of *hashtags* and *changesets_count*. Hashtags are strings starting with # sign that are used to add thematic context to the mapped object. For example, objects mapped within HOT-OSM project for specific geographical area would contain hashtag #hotosm- appended with the name of the target geographical area. Changesets count is the number of changes that user has created. Typically, these are added to the changesets created by users with under 100 edits total. Of the most importance for this research was the tag *review_request=yes*.

3.1 The analysis process

The analysis was performed mainly by using custom developed shell scripts. First preprocessing step was to cut off the part of the changesets file before the 25th of August. The analysed changesets were the ones whose editing started in the period between 2017-08-25T20:31:25Z and 2017-09-08T20:29:59Z. The total number of such changesets in this period was 417,194. The second step involved extracting only the changesets which contained tag *review_request=yes* with a total of 10,799. The total number of such changesets divided by the number of days within the timeframe of the analysis gives 771 requested reviews per day. Given the worldwide distribution of the changesets and the number of senior contributors that are active within the OSM on a daily basis, we find that it is highly unlikely that so many reviews per day will be adequately answered by the community. Having all the changesets data, next step was getting all the related discussions from the OSM API. The process of discussions retrieval was not overloading the OSM API since there were too few changesets with discussions included (based on the *comments_number* XML attribute).

The changeset is a logical grouping of changes in the OSM, but it does not contain the actual changes. The changes are split in three possible changeset operations: create, modify and delete, which can be performed over a node, way or relation and are stored in *osmChange* diff files. The diff files with all the actual changes can be obtained from Planet OSM replication diffs area for each minute, hour or day in the past. Based on the date when the changeset was created a shell script calculated the file name and then downloaded and extracted the proper daily *osmChange* data file. Based on the referenced changeset ID it was possible to count how many create, modify and delete operations were performed for nodes, ways and relations for each changeset. Finally, based on changeset ID, proper *OsmCha* and OSM links were created in order to enable us to review manually particular changesets. The outputs of this process were the changeset XML data, change statistics, links and a file containing changeset discussion where available. We created a C# console application for parsing these textual files and storing obtained data into the PostgreSQL database. During the parsing process the application used the Language Detection API in order to retrieve the language used for the comments. The changeset bounding boxes were then visualized in Leaflet.

3.2 General changesets analysis results

We counted how many changesets have comments (responses). The number of uncommented changesets is

10,510. Taking into account that the expected way of reviewing the changesets (Neis, 2017) was through comments, it means that only 2.68% changesets display evidence of being reviewed through the actual comments provided by other contributors. This indicates only small evidence of reviews actually happening. The next step was to check the number of changes per changeset. The result was unexpected with many changesets having high numbers of changes. Ten changesets with highest number of changes are shown in Table 1:

Table 1: Top 10 changesets with highest number of changes.

Changeset Id	No. Changes	Comments Count	HOT-OSM
51581404	7118	0	No
51587971	2855	0	Yes
51646404	2548	0	Yes
51468822	2322	0	Yes
51592470	2305	1	Yes
51757887	2246	0	Yes
51679311	2192	0	Yes
51762813	2068	0	Yes
51588296	1984	0	Yes
51647890	1874	0	Yes

Obviously, such changesets cannot easily be reviewed, which is demonstrated by the absence of the comments, except for one of them. The content of the comment was: *""@TidoB I am quite sure that these round shaped objects are not buildings. Please check! Happy mapping*", which suggests that the original contributor made some geometrical errors during their mapping. The average number of changes per changeset was 100.75. However, changesets which had been commented upon had approximately 39.2 changes. This reduces to 31.33 when the Table 1 changesets are removed from analysis. The uncommented changesets' average was 102.45, due to the much higher number of uncommented changesets.

It is evident that a significant number of changesets contained the hashtag in the form of *#hotosm-project-xxxx*. Of the 10,799 analysed changesets, there were total of 6070 changesets belonging to some HOT-OSM (Humanitarian OSM) project. Of these, only 24 were actually commented. This means that HOT-OSM projects get commented approximately 0.39% of the times compared to 6.11% for none HOT-OSM related changesets. In addition, HOT-OSM project changesets had approximately 164 changes compared to only 8.4 changes for the regular changesets. This clearly shows that HOT-OSM related changesets are skewing the data.

From the Table 2, it is noticeable that in general, the highest number of changes is performed on the nodes, specifically on nodes creation. The second highest number of changes is for ways creation. Modifications are around 5% of the changes while deletions occur very rarely. We speculate that new OSM contributors usually tend to map new data, thus creating it. They do not change existing data and may not even try deletions. It is also noticeable that relation changes are very rarely part of the changesets which requires review. We think it is because of its more complex nature, causing novice contributors to avoid interaction with them. If only changesets with existing comments are analysed we gain a little more insight into what type of changesets are attracting more attention and having more changes to be reviewed.

Table 2: Average number of changes per object type and change type for each changeset.

Feature type	Change type	Average Number of changes
Node	Create	78.25
	Modify	4.97
	Delete	2.18
Ways	Create	12.51
	Modify	1.78
	Delete	0.28
Relation	Create	0.030
	Modify	0.050
	Delete	< 0.001

One might expect that changesets with fewer changes are more favourable for reviewing. Table 3 shows the average number of changes per object type and change type for changesets with comments.

Table 3: Average number of changes per object type and change type for each changeset that has comments (that went through the review process).

Feature type	Change type	Average Number of changes
Node	Create	26.51
	Modify	3.05
	Delete	2.83
Ways	Create	3.75
	Modify	2.63
	Delete	0.21
Relation	Create	0.10
	Modify	0.14
	Delete	0

When compared to Table 2, it is noticeable in Table 3 that reviewed objects have less changes for all the three change types when feature types are nodes and ways. This gives some evidence to support the assumption that more reviews happen on changesets with less changes. For relations however, it is not the case and average number of changes is higher for the reviewed changesets than overall. Nevertheless, this number of changes is still much lower than the number of changes for other two feature types.

Table 3 also shows that changesets with modifications and deletions attract the attention of reviewer(s) (the average numbers are higher compared to average numbers for all the changesets). This is most likely because such modifications are more likely to happen in already well mapped regions with lots of experienced users who will notice such changes in data and then react to them. In areas where new mapping occurs, local knowledge of the mappers builds up and experienced contributors will be able only to review data compliance to the general recommendations but less likely positional or thematic accuracy of created data.

Figure 1: Changeset bounding boxes: commented ones in blue and uncommented ones in red

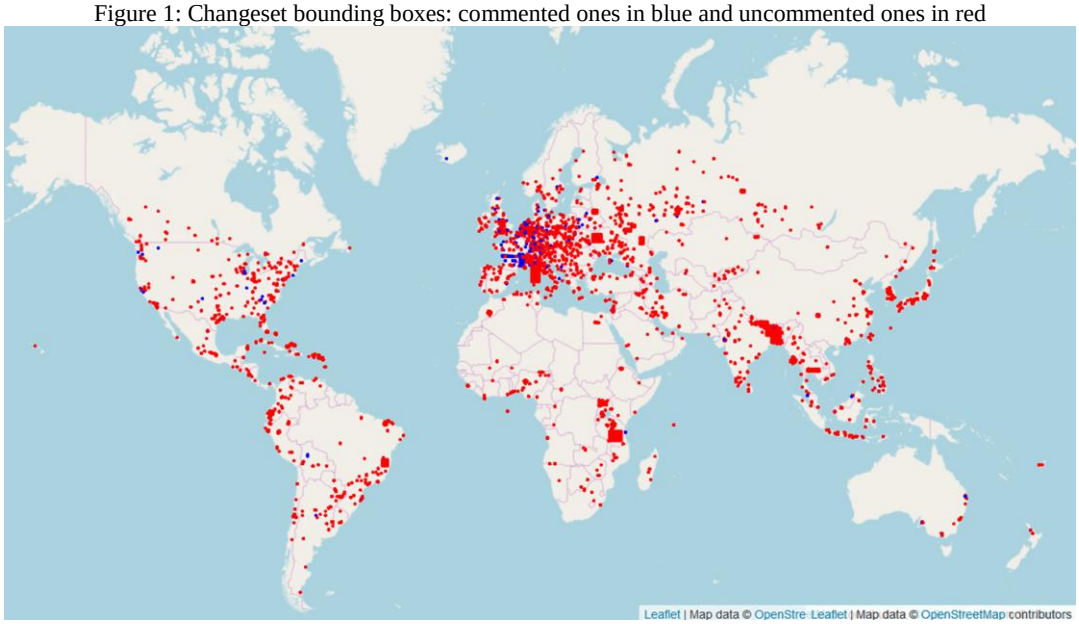

3.3 Spatial distribution

Figure 1 shows the locations of the bounding boxes of the changesets that were analysed. From Figure 1, it can be noticed that the highest density of changesets with *review required* were created in Europe and an even the higher percentage of them was commented. This supports our previous findings that modifications and deletions were attracting higher percentages of reviews, since such are happening in more mapped areas with experienced contributors.

Figure 2 shows the additional spatial distribution of the changesets created in the scope of some HOT-OSM projects. It is noticeable that, even though there is a higher number of changesets belonging to some HOT-OSM project their spatial dispersion is much smaller than the other changesets. HOT-OSM changesets are tightly coupled around the regions that were the subject of HOT-OSM activity. It is also noticeable that most of the changesets are grouped in the regions where OSM is already very well developed. This may suggest that new mapping efforts are less concentrated on less well mapped areas/regions. Where mapping is occurring in these areas few contributors are using the new reviewing option.

Figure 2: Changeset bounding boxes: HOT-OSM ones in blue and others in green

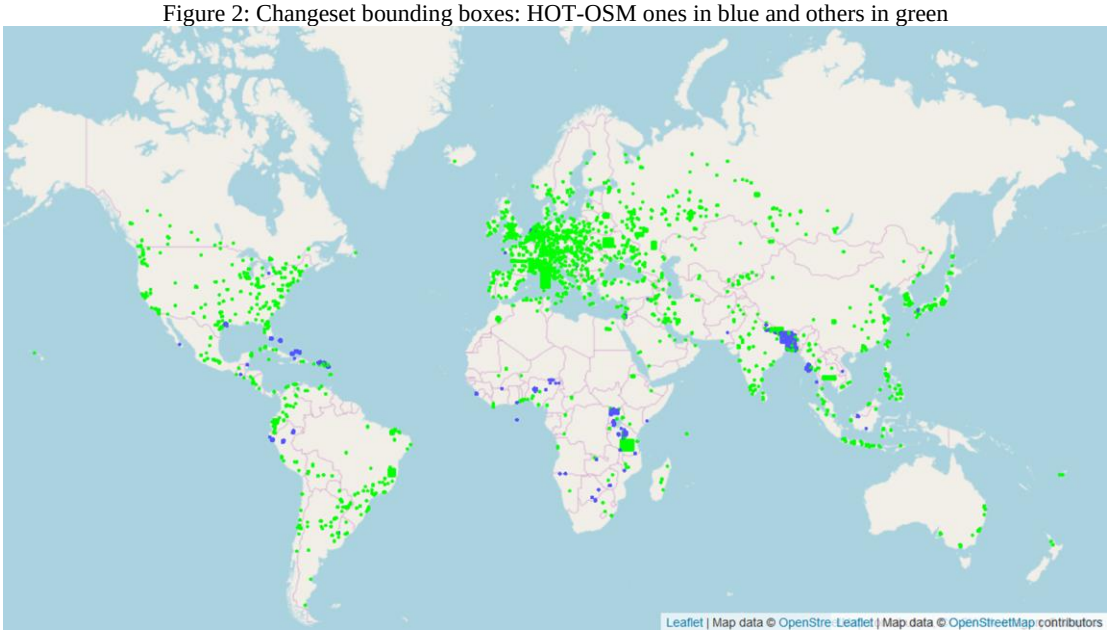

3.4 Discussions language

Table 4 shows the number of time different languages were used in cases where changesets were reviewed.

Table 4: Languages used for review comments

Language	Number of usages
English	91
German	51
Dutch	39
French	35
Polish	29
Russian	24
Italian	6
Spanish	4
Danish	3
Swedish	2
Portuguese	2
Czech	1
Icelandic	1
Ukrainian	1

Looking at the language analysis from the Table 4, we were expecting higher usage of English language than it actually is. Even though reviews should mainly rely on the contributor's local knowledge this does not exclude the language of the contributors providing reviews. In cases where comments are written in local languages they may not be understandable for other OSM contributors.

4 Conclusions and Future Work

This analysis has allowed us to draw a number of conclusions which are summarized in the following items:

- Surprisingly high number of changesets with *review_request* actually have a very low number of discussions/comments.
- Some changesets with huge number of changes also have *review_request*.
- Many changesets belong to HOT-OSM projects, which could potentially be more efficiently reviewed by senior project managers and leaders within the HOT-OSM organization.
- At this early stage of adoption it seems that the current review structure may not be fit for purpose due to too many reviews being requested and too few active contributors perform reviews (evidenced by comments).

Our analysis allows us to make the following suggestions for managing the *review_requested* process in the future:

- Review requests should be made only by relatively novice contributors. Given that current average of created changesets per contributor who requested review is 91, it seems reasonable to enable this feature only for user with significantly less created changesets. Such a constraint could reduce the overall number of changesets requiring review while also lowering the average number of changes per changeset.
- Motivational approaches could be used to facilitate reviews for unmapped or remote areas where too few local experienced contributors exist.

Future work will include a more in-depth analysis of the changesets to investigate if actual data editing took place after reviewing. We noticed some evidence of such activity during this analysis. Now that the *review_requested* process is several months old this will provide us with a useful indication if this structure of peer-review within OSM has potential to be successful and sustainable in the longer term.

References

Antoniou, V. User Generated Spatial Content: An Analysis of the Phenomenon and Its Challenges for Mapping Agencies. Ph.D. Thesis, University College London (UCL), London, UK, 2011

Ludwig, I.; Voss, A.; Krause-Traudes, M.A. Comparison of the street networks of Navteq and OSM in Germany. In *Advancing Geoinformation Science for a Changing World*; Geertman, S., Reinhardt, W., Toppen, F., Eds.; Springer: Berlin, Germany, 2011; pp. 65–84.

Graser, A.; Straub, M.; Dragaschnig, M. Towards an open source analysis toolbox for street network comparison: Indicators, tools and results of a comparison of OSM and the official Austrian reference graph. *Trans. GIS* 2014. [CrossRef]

Neis, P.; Zielstra, D. Recent developments and future trends in volunteered geographic information research: The case of OpenStreetMap. Future Int. 2014, 6, 76.

Neis, P, Review requests of OpenStreetMap contributors – How you can assist!, http://neis-one.org/2017/09/review-requests-osm/, 2017, last accessed 14/02/2018

Planet OSM website, https://planet.openstreetmap.org/ , last accessed 14/02/2018

Davidovic, N., Mooney, P., Stoimenov, L. and Minghini, M., 2016. Tagging in Volunteered Geographic Information: An analysis of tagging practices for cities and urban regions in OpenStreetMap. *ISPRS International Journal of Geo-Information*, *5*(12), p.232.

Minghini, M., Antoniou, V., Fonte, C.C., Estima, J., Olteanu-Raimond, A.M., Minghini, M., Antoniou, V., Fonte, C.C., Estima, J., Olteanu-Raimond, A.M. and See, L., 2017. The relevance of protocols for VGI collection. *Mapping and the Citizen Sensor*, pp.223-247.

Senaratne, H., Mobasheri, A., Ali, A.L., Capineri, C. and Haklay, M., 2017. A review of volunteered geographic information quality assessment methods. *International Journal of Geographical Information Science*, *31*(1), pp.139-167.