

AUTHOR COPY
July 2015

RESEARCH ARTICLE

Conceptualising the geographic world: The dimensions of negotiation in crowdsourced cartography

Andrea Ballatore^{a*} and Peter Mooney^b

^a*Center for Spatial Studies, University of California, Santa Barbara, CA*

^b*Department of Computer Science, Maynooth University, Ireland*

(Received 00 Month 200x; final version received 00 Month 200x)

In crowdsourced cartographic projects, mappers coordinate their efforts through online tools to produce digital geospatial artefacts, such as maps and gazetteers, which were once the exclusive territory of professional surveyors and cartographers. In order to produce meaningful and coherent data, contributors need to negotiate a shared conceptualisation that defines the domain concepts, such as road, building, train station, forest, and lake, enabling the communication of geographic knowledge. Considering the OpenStreetMap Wiki website as a case study, this article investigates the nature of this negotiation, driven by a small group of mappers in a context of high contribution inequality. Despite the apparent consensus on the conceptualisation, the negotiation keeps unfolding in a tension between alternative representations, which are often incommensurable, i.e., hard to integrate and reconcile. In this study, we identify six complementary dimensions of incommensurability that recur in the negotiation: (i) ontology, (ii) cartography, (iii) culture and language, (iv) lexical definitions, (v) granularity, and (vi) semantic overload and duplication.

Keywords: Crowdsourced cartography; semantic negotiation; ontology engineering; volunteered geographic information; OpenStreetMap

*Corresponding author. Email: aballatore@spatial.ucsb.edu

Introduction

Over the last decade, Web 2.0 technologies have enabled a major change in the landscape of production and consumption of geographic information. The combination of inexpensive personal computers, GPS sensors, smartphones, high-speed Internet connections constitutes an ideal platform for amateur mappers to practise forms of crowdsourced cartography (Dodge and Kitchin 2013). Following the unexpected success of Wikipedia, wiki tools are being applied to the mass collection of geographic data in OpenStreetMap and WikiMapia. Relying on volunteered work, these projects aim at generating reusable maps of the world, free from the organisational, methodological, and epistemological constraints adopted by traditional mapping agencies (Gerlach 2014). What Goodchild (2007) named ‘volunteered geographic information,’ citizen scientists produce and share data about their surroundings, and can be conceived as millions of sophisticated ‘citizen sensors,’ whose contributions can sustain considerable scientific and commercial applications .

A crucial and understudied aspect of these forms of peer-production of geographic information lies in their rich, open, and fragmented conceptualisations. By conceptualisation, we mean a set of inter-related concepts that are used to describe aspects of reality. For example, emergency mappers create knowledge about an area using concepts such as building, damage, access, shelter, and presence of water, while cyclist mappers might see the same portion of reality in terms of cycle lane, cycle track, elevation, traffic, and scenic routes. The success of such crowdsourced cartography largely depends on the ability to coordinate around shared practices, channelling efforts constructively around digital artefacts (Ballatore 2014).

In fact, these groups of volunteered mappers can be framed as information communities of semantic agents, i.e., agents who are able to communicate by encoding and decoding meaning into and from symbols, such as words, maps, sounds, etc. These different conceptualisations can be defined at varying degrees of formalisation, ranging from implicit (e.g., commonsensical meaning intuitively shared among agents) to explicit (e.g., strictly defined meaning that requires conscious effort from the agents) (Ballatore *et al.* 2013). Conceptualisations can be encoded in information using a range of knowledge representation tools, such as simple vocabularies, taxonomies, and formal ontologies.

Crowdsourced cartography has important ontological aspects in two senses. In the philosophical sense, these new mapping practices attract attention in human geography and critical GIS for their collective, transient, and mutable nature (Schuurman 2006, Warf and Sui 2010, Elwood *et al.* 2012, Leszczynski and Wilson 2013, Gerlach 2014, Perkins 2014). On the other hand, geographic information science (GIScience) has been engaging in formal ontology in the tradition of analytic philosophy and knowledge engineering, aiming at clarifying formally the logical constructs used to encode geographic information (Smith and Mark 2001, 2003, Kuhn 2003, Janowicz *et al.* 2013). As these two approaches are grounded in divergent intellectual frameworks and goals, in this study we adopt the GIScientific approach, without engaging directly in the rich debates in geography.

Regardless of the specific technical infrastructure, the development and negotiation of a conceptualisation involves two dimensions. First, these processes revolve around cartographic ontology, particularly the decision of what entities are mappable, how they are structured internally, and how they connect to one another. Second, crowdsourced cartography encounters the age-old problems of semantics. Implicitly grounded in a correspondence theory of truth (Kuhn 2009), contributors try to constrain the semantic linkage between symbols and the underlying concepts, i.e., the psychological constructs residing in the human mind, pointing to objects in the real world, traditionally called

referents. In different information communities, the same term often refers to incompatible concepts, resulting in more or less commensurable semantic spaces, in need of either translation or alignment. For example, the meaning of terms such as ‘mountain’ or ‘river’ can vary widely depending on the agent’s local context (Janowicz 2012). When working on collaborative world maps, a stable vocabulary has to be negotiated among a heterogeneous population to act as a central component of the mechanism.

In their practices, volunteered mappers engage in forms of ontology engineering, defining ‘explicit specifications of conceptualizations’ that enable them to share information with other semantic agents (Guarino *et al.* 2009, p. 1). As semantic ambiguity is an intrinsic and unavoidable feature of human communication, building shared meanings around terms and ontological constructs is unsurprisingly complex. Given that absolute agreement is impossible to reach even in close, small, homogeneous groups of semantic agents with precise purposes, as observable in academic and industrial contexts, it is not trivial to understand how conceptualisations emerge from the interactions of large numbers of individuals in a media ecology sustained by asynchronous, Web-based forums and mailing lists.

While the spatial dynamics behind the generation of crowdsourced maps have received academic attention (e.g. Haklay 2010, Mooney and Corcoran 2014), the ontological engineering that underpins volunteered mapping is as yet largely unknown. To illuminate this central aspect of crowdsourced cartography, this article investigates the semantic negotiation in OpenStreetMap (OSM), the leading endeavour in the field. After surveying the existing work on semantic negotiation, we describe the OSM semantic ecosystem, in which meanings are generated and recorded. Methodologically, we operated as follows. Taking the OSM Wiki website as a case study, we carried out a qualitative analysis, investigating the website’s growth and contribution inequality over time. Subsequently, we devised a measure of the negotiation to identify the terms that, overall, attracted most negotiation in the website. On the top 25 terms, we performed a content analysis, extracting the types of incommensurability that characterise this collaborative form of ontology engineering, providing salient exemplars.

Semantic negotiation in crowdsourced cartography

For our study, useful conceptual tools are found in the philosophical area of geographic ontology, and that of ontology engineering. As Smith and Mark (2001) put it, the traditional philosophical area of ontology ‘seeks to study in a rational, neutral way all of the various types of entities and to establish how they hang together to form a single whole (‘reality’)’ (p. 592). In this sense, the age-old effort of ontology can be seen as a systematic work of analysis and clarification of the internal structure of the entities that tacitly populate human cognition and language.

From an applied perspective, the areas of conceptual modelling and ontology engineering aim at mitigating recurring problems in the authoring of conceptualisations with artificial languages, such as database schemas, object-oriented models, or first-order and description logics. As many diverse semantic agents can be involved in the definition of an ontology, managing semantic disagreements is key to creating any complex, knowledge-intensive system. In an analogous way, collaborative ontology design methods aim to involve a number of participants in the process who hold diverging viewpoints about the same domain (Karapiperis and Apostolou 2006). Such methods emphasise that semantic agreements are essentially social, inter-subjective phenomena (Kuhn 2009), and that an information community needs constant, intense social interaction to function. The

success of a conceptualisation depends therefore on the commitment and motivations of the participants, on their epistemological biases, and on the social skills of the process' coordinator, who has to steer complex social interactions towards consensus.

As a conceptualisation emerges, a central preoccupation for volunteer cartographers is the classification of real-world entities according to it. Collaborative tagging has emerged from the Web 2.0 paradigm, allowing large numbers of non-expert users to attach arbitrary labels to items. These 'folksonomies' (i.e., taxonomies defined by ordinary people) tend to be more flexible, scalable, decentralised, and easier to maintain than more formal classification schemas and ontologies, but suffer from semantic incoherence, idiosyncratic usage of tags, as well as polysemy and synonymy, limiting the tags' informational value. (Rorissa 2010). All of these aspects converge in OSM, described in the next section.

The OpenStreetMap semantic ecosystem

In this study, we focus on the case of OpenStreetMap (OSM), observing the environment in which the project's ontological and semantic negotiation is carried out.¹ Since its inception, OSM has established itself as the most ambitious crowdsourced cartographic project. The project maintains a free and open world vector map, generated by the contributions of 2.2M users as of July 2015. Unlike traditional cartographic datasets, the conceptual modelling of OSM explicitly rejects a top-down, expert-driven ontology. The scope of OSM has no precise boundaries, and a huge number of geographic objects are included in the map, modelling heterogeneous sections of the natural and built world.

From a semantic viewpoint, OSM is a semi-structured folksonomy, which allows contributors to create any new term (a key-value pair called 'tag') to describe the objects that they find worth mapping. This conceptualisation is an unfinished, evolving product, permanently open to re-negotiation and modification. The project faces a tension between the technical need for a unified conceptualisation, and the vital necessity to enable contributors to express their local knowledge. OSM contributors create new spatial objects on the map, and describe their semantics with terms. For example, terms `highway=primary` is used to label primary roads, and `amenity=university` indicates universities. Unlike Wikipedia, where the vast majority of the editing occurs on the wiki website itself, OSM relies on the complex interplay between different media. The central media elements of this semantic ecology are the following:

- *Vector map.* This is the core artefact of the project, consisting of a large planetary vector map.
- *Mailing lists.* The bulk of communication and coordination occurs in the project mailing lists.
- *Forum.* The forum plays the same role as the mailing lists, but it is considerably less used.²
- *Wiki website.* The central component of this ecosystem is the OSM Wiki website,³ which hosts the core documentation of the project, including the definitions of terms and usage guidelines.
- *Meta-data monitors.* Web services such as *TagInfo* monitor and provide summary statistics of the actual tag usage in the vector map, and this information is used to make semantic decisions.

¹<http://www.openstreetmap.org>

²<http://forum.openstreetmap.org>

³<http://wiki.openstreetmap.org>

- *Map editors*. The vast majority of contributors edit the map through dedicated software editors, such as *iD* and *Potlatch 2*.⁴ These tools mediate and constrain the editing process, and each editor reflects ontological assumptions and commitments.
- *Map renderers*. These software tools, such as *Mapnik*,¹ visualise vector data following a stylesheet containing rules that determine the representation of map features based on their type.

In OSM, semantic negotiations occur at all times. New tags are proposed in the wiki, voted upon, utilised directly in the vector map, recorded by the meta-data services, and interpreted by the map renderers. The intended meaning of existing tags is constantly discussed on the mailing lists, and changes to the software editors' map features are heavily debated. As the vast majority of these interactions are public, archived, and freely accessible, the OSM ecosystem constitutes an ideal ground to study cartographic semantic negotiation, as we show in the next section.

Case study: the OSM Wiki website

As the full mapping of the interactions in the OSM semantic ecosystem is difficult to reconstruct and interpret, we selected the OSM Wiki website² as a case study. At the core, the OSM Wiki website is a set of pages dedicated to specific topics—for example, a page focuses on the mapping of rivers.³ Four types of users edit the wiki: normal users, administrators, bureaucrats and bots. Normal users, after creating an account, can read and edit any page of the wiki. Similarly to the Wikipedia's sysops, the administrators can block users, delete, and protect pages. Bureaucrats, on the other hand, form an elite, and are closely linked to the Data Working Group, which deals with important issues in copyright violation, disputes, and vandalism, and coordinate the development of broad project policies. Finally, bots are automated scripts that perform maintenance operations on the wiki.

First, we analysed the temporal evolution of the website, measuring its growth in terms of users, pages, and content. Subsequently, we focused on the patterns of participation inequality. To carry out this quantitative analysis, a complete dump of OSM Wiki website was downloaded.⁴ The dataset contains a tripartite graph of 808,503 revisions performed by 16,420 active contributors over 90,627 pages.⁵ Since its creation in April 2005, the OSM Wiki website has experienced rapid and tumultuous growth, summarised in Figure 1. Plot (a) shows the overall editing activity, including the number of active users, edited pages, revisions, and words added to the website. From the observation of the temporal development, it is possible to identify considerable regularities throughout this collaborative writing process. All these variables follow a similar curve, with an initial faltering start, followed by exponential growth between 2005 and 2008. The growth then slows down and reaches a peak between 2009 and 2011 (marked with red '+'), and finally slowly decreases. This trend can be compared with Wikipedia that, after a phase of exponential growth, reached a peak and started decreasing, possibly because of management overheads, coverage of 'easy' topics, and surging coordination costs (Suh *et al.* 2009).

⁴<http://wiki.openstreetmap.org/wiki/Editors>

¹<http://wiki.openstreetmap.org/wiki/Renderers>

²http://wiki.openstreetmap.org/wiki/Map_Features

³<http://wiki.openstreetmap.org/wiki/Tag:waterway=river>

⁴The dataset includes activity from April 2005 to November 2012.

⁵Bots were removed from the dataset to focus on manual contributions.

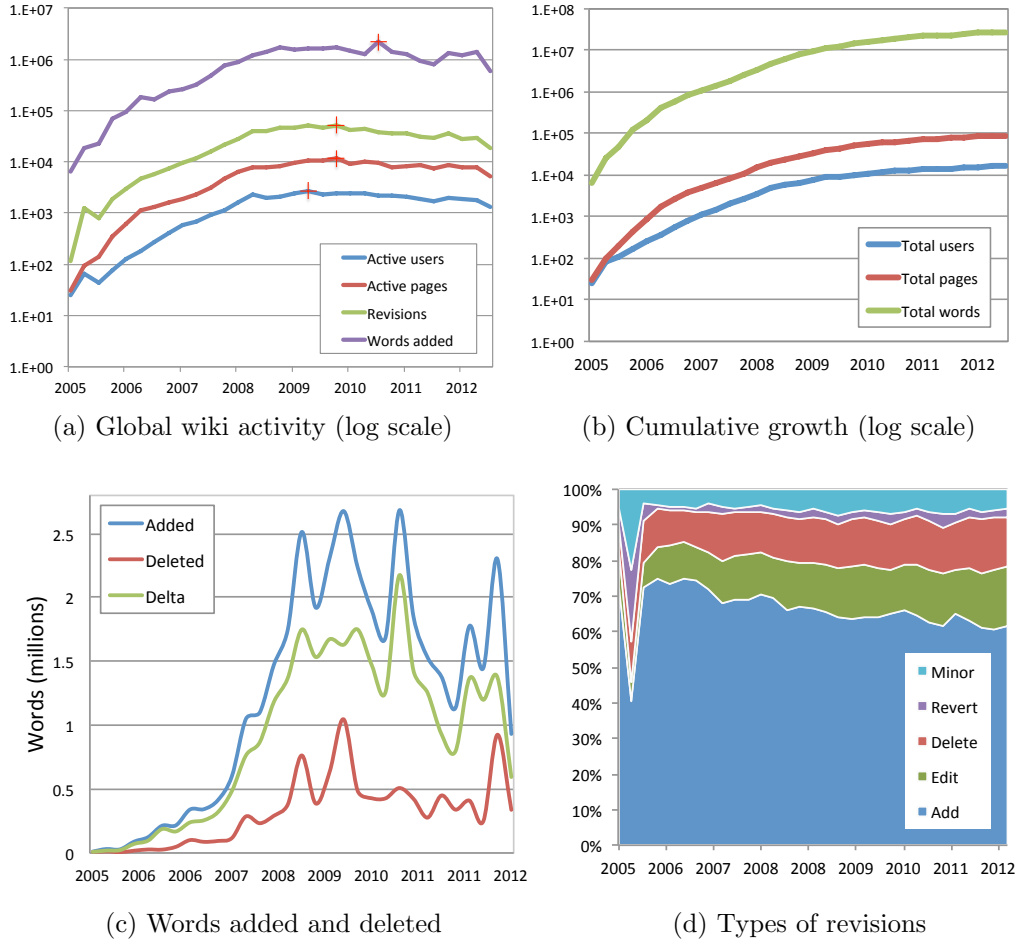


Figure 1.: Temporal development of the OSM Wiki website

Relation	Top 1%	Top 10%	Top 20%	Gini coeff.
Revisions per contributor	39.1	81.5	91.1	.87
Added words per contributor	48.6	87.9	95.3	.92
Deleted words per contributor	51.0	90.3	96.8	.97
Revisions per page	30.2	71.0	82.8	.77
Added words per page	39.5	81.4	92.9	.89
Deleted words per page	71.5	98.7	99.9	.98

Table 1.: Contribution inequality in the OSM Wiki website. Each number is the percentage of the global contribution.

The production and maintenance of OSM Wiki website is driven by a very active minority of users. To observe the weight of the top contributors, we computed the cumulative distributions at typical cutoff points (1%,10%,20%) along several dimensions (see Table 1). These results indicate radical contribution inequality in every aspect of the production process, with the top 1% of users generating a staggering 48.6% of the wiki's words. Overall, the top 20% of users, pages, and revisions consistently generated more than 82% of the content.

This analysis indicates that the OSM Wiki website is created, edited, and maintained by the top 20% of contributors, which includes approximately 3,200 people, with a core group of about 15 major contributors, whose composition changes over time. As the least

Page p variables	symbol	median	mean	max	skewness
# unique contributors	c_p	2	4.0	393	14.2
# revisions	r_p	3	9.9	2,574	38.7
# reverts	e_p	0	.3	61	20.7
# deleted/added word ratio	δ_p	0	0.0	1	3.3
# active days	d_p	2	5.1	735	22.7

Table 2.: Indicators of negotiation

active 80% contributors generated 4.7% of the words added to the wiki, the impact of the ‘long tail’ of occasional contributors appears negligible. To observe the change of contribution inequality over time, we computed the Gini coefficient for every quarter, with respect to the contributors and size of the contributions. The Gini coefficient is a commonly used measure of income inequality based on statistical dispersion, originally conceived in economics, and then applied to other domains. After an initial phase of instability, the coefficient tends to grow over time from an overall mean .66 in 2005, to an overall mean of .82 in 2012. These results are highly consistent with those reported about Wikipedia by Ortega *et al.* (2008).

The dimensions of incommensurability

As the OSM Wiki website is a central element of the project’s semantic ecosystem, negotiation leaves traces in the pages’ edit history for each term of the conceptualisation, and their analysis can illuminate the nature, scope, and recurring patterns in the incommensurability encountered by contributors. While edit wars in Wikipedia tend to reflect known controversies in the domains that fall within the encyclopaedia’s scope, such as politics, ethics, history, and religion (Rad and Barbosa 2012), the specific areas of intense negotiation in OSM need to be identified. For this purpose, we carried out a qualitative analysis in three steps. First, we devised a measure N_p of the negotiation intensity, and we computed it on the entire dataset. Second, we selected the terms that generated most negotiation, resulting in a corpus of edits, and related emails and forum discussions. Third, we performed a full content analysis of this corpus, identifying recurring dimensions of incommensurability encountered by contributors.

To define a measure of negotiation, we assume that controversial pages attract many contributors, generate a high number of revisions and reverts, have a high deleted/added word ratio (i.e., a lot of content is deleted as well as created), and that such negotiation is also sustained over time. Several indicators can be used to quantify the intensity of the negotiation in a page, including the number of unique contributors (c), revisions (r), reverts (e), the deleted/added word ratio (δ), and active days (d) (see Table 2).

As is possible to notice, the variables express widely different units, and are heavily skewed positively, making the usage of parametric statistics problematic. Hence, we adopt a ranking function rk as a non-parametric approach to normalise the variables, and we define a negotiation intensity measure N_p , where p is a page and p_t its corresponding ‘talk page’ (a discussion page, if present). $v(i)$ represents the negotiation intensity for a page i , and α is a weighting factor controlling the importance of the negotiation in the page itself and its talk page:

$$v(i) = \frac{rk(c_i) + rk(r_i) + rk(e_i) + rk(\delta_i) + rk(d_i)}{5}$$

$$N_p = \alpha v(p) + (1 - \alpha) v(p_t) \quad \alpha \in [0, 1] \quad (1)$$

N_p	Key/tag page	Page p					Talk page p_t				
		c	r	e	δ	d	c	r	e	δ	d
1	highway	103	235	10	.99	129	53	128	2	.13	60
2	boundary	97	246	2	.99	130	65	219	1	.14	100
3	building	43	100	3	.97	65	27	121	1	.28	47
4	barrier	38	112	8	.94	57	39	135	1	0	72
5	access	82	238	22	.43	136	86	296	1	.39	141
6	traffic calming	59	119	5	.65	69	22	32	2	.77	19
7	amenity=place of worship	33	71	4	.72	44	49	127	2	.05	66
8	waterway=riverbank	55	139	8	.69	74	10	44	4	0	23
9	service	39	107	2	.89	44	17	45	3	.01	21
10	wheelchair	36	114	3	.61	60	13	51	1	.09	20
11	natural=tree	48	109	2	.48	60	25	65	1	.17	34
12	shop	29	50	2	.94	37	31	68	1	.09	41
13	bridge	43	105	5	.34	60	38	78	1	0	43
14	highway=turning circle	45	70	4	.76	31	17	43	1	0	12
15	sport=shooting	40	85	2	.74	44	12	20	2	.52	11
16	area	40	107	2	.91	51	8	37	1	0	15
17	smoothness	17	75	18	.60	30	9	58	11	.27	15
18	natural=coastline	25	73	2	.56	54	16	39	1	.19	25
19	railway=subway entrance	34	67	2	.70	42	13	24	2	.01	16
20	power	27	74	1	.98	45	22	38	1	.47	27
21	mountain pass	32	98	2	1	38	7	9	3	.02	7
22	is in	30	54	3	.33	37	17	41	2	0	22
23	highway=unclassified	23	93	4	.34	31	14	38	1	0	16
24	railway=level crossing	31	49	1	.92	30	18	32	1	.65	19
25	power=generator	21	44	3	.29	30	17	34	2	.05	18

Table 3.: Top key/tag pages w.r.t. negotiation measure N_p . (c) contributors; (r) revisions; (e) reverts; (δ) del./added word ratio; (d) active days.

The formula ranks the pages along five ranked variables, as a proxy to negotiation intensity. In this sense, all the five variables contribute in equal measure to the final rank of a page. The measure N_p was then computed on the OSM Wiki website, aiming at the identification of the pages that attracted most negotiation. As generic hub pages such as *Map Features* and *Proposed Features* tend to generate intense but thematically unfocused activity, we considered only *Key:* and *Tag:* pages in English, for a total of 3,185 pages. Parameter α was set to .5 to keep a balance between the pages and their associated talk pages. Out of all pages, we selected the top 25. Unlike pages with lower N_p , these pages represent the core semantic negotiation that occurred in the project. Past this group, the variables rapidly decrease, showing a relatively low level of negotiation. Table 3 displays the intensity of negotiation in the top pages, showing the variables in detail, highlighting the wide differences between terms. As is possible to observe, the top of the group include terms related to the road network, such as *highway*, *barrier*, and *access*, the conceptualisation of boundaries, buildings, and vague natural features such as riverbanks.

To identify recurring themes in the negotiation, we considered these 25 pages and the related talk pages, as well as threads in the mailing lists, collecting a corpus of about 162,000 words. We performed a full content analysis of the corpus, focusing on the negotiation-related issues explicitly articulated in the discussions. We then identified a number of categories that capture the dimensions of incommensurability. This analysis uncovered that, when expert conceptualisations exist, for example for land cover categories, OSM contributors tend to adopt them, integrating them into the project’s context. This conceptual integration appears to be successful and relatively uncontroversial in most instances. By contrast, when no pre-existing source is available, the negotiation occurs along the following dimensions, described in the following sections: (i) ontology,

(ii) cartography, (iii) culture and language, (iv) lexical definitions, (v) granularity, and (vi) semantic overload and duplication. These dimensions are complementary, and often co-occur in the same conceptual area. For example, many ontological conflicts have a cultural dimension, and definitional conflicts are often due to deeper ontological ambiguities. Ontological and definitional negotiation plays an important role in most pages, while semantic overload only affects the negotiation of very generic terms.

Integration with expert conceptualisations

Expert-defined official conceptualisations are not necessarily less contingent, controversial and contested, but they provide useful templates, already deployed within information communities in real contexts. The lengthy negotiations and difficulties encountered in the Infrastructure for Spatial Information in the European Community (INSPIRE) offer a striking illustration of the challenges in achieving actual semantic interoperability between heterogeneous information communities.¹ OSM contributors who are working on broadly explored domains re-use existing conceptualisations, rather than designing new ones. Nevertheless, when contributors re-use a conceptualisation in the project, a set of recurring issues emerges.

Existing expert conceptualisations are often too detailed for the non-specialist scope of OSM, and contain culturally-specific, obscure technical jargon. The main strategy adopted in these cases is that of selection and simplification: contributors choose an appropriate subset of the conceptualisation, and translate it into the OSM context. This was the case with land use classification. CORINE Land Cover is a European programme launched in 1985 to enable semantic interoperability on land cover between European countries. In some European countries, such as France, the CORINE nomenclature has been imported into OSM. A set of semantic transformation rules were developed to transform CORINE land cover nomenclature into OSM, while many complex classes were not included.

In some well-defined and circumscribed domains, international standards are sufficiently clear and established and appear to be imported without much friction. Notably, the conceptualisation of administrative boundaries in page *Key:boundary* benefited from the presence of accepted international standards. European contributors started to adopt the Nomenclature of Territorial Units for Statistics (NUTS) defined by Eurostat. Another conceptualisation adopted by OSM is the codes for the representation of names of countries and their subdivisions, defined in the ISO 3166 standard. When official or scientific conceptualisations are either unavailable or not publicly accessible, contributors resort to crowdsourced and unorthodox sources. Also because of the project's similar ethos, Wikipedia and its related projects are a major source of inspiration for OSM contributors, which often rely on it for definitions and classifications. A railroad modelling manual (Mallery 1992) is used as a source for the OSM conceptualisation of bridges. In general, the lack of established standards forces contributors to generate a new conceptualisation.

Ontological negotiation

The most central dimension of the negotiation falls in the realm of ontology engineering, i.e., the attempt to extract a formal conceptualisation from tacit knowledge, but in an informal, computer-mediated setting. In the context of OSM, contributors encounter topological, mereological, and geometrical issues. Indeed, the purpose of their efforts is not the definition of a sound geographic ontology per se, but aims towards a simple,

¹<http://inspire.jrc.ec.europa.eu>

shared, consensual conceptualisation of the entities, so that they can be easily rendered in a cartographic form. To represent geographic entities, a *topology* is needed. A topology aims at defining formally how entities can be connected or contiguous, and relies on a theory of boundaries, connectedness, interiority/exteriority, and separation. A *mereology*, by contrast, is a theory that specifies how the constituent parts result in whole entities (e.g., a tree is a part of a forest).

The conceptualisation of boundaries plays a prominent role in geographic ontology engineering. Many geographic terms are semantically vague, and their demarcation is bound to cause disagreement (Agarwal 2005). While *bona fide* boundaries and objects emerge from discontinuities in the world (e.g., between islands and seas), *fiat* boundaries and objects are the result of an arbitrary human choice, either social, political, cultural, or linguistic. Once these mereotopological details have been decided, conflicts can still arise around *geometry*, as the same real-world entities can be expressed using different geometrical representations, such as points, polylines, polygons, multi-polygons, or relations between existing geometries. Consistently with the tenets of naïve geography (Smith and Mark 2003), fields are not represented in OSM, and crisp, discrete boundaries are preferred to vague ones.

In the OSM Wiki website, ontological conflicts are present in most discussions, particularly along the boundaries between heterogeneous entities. The negotiation of *Key:boundary* highlights the topological difficulties in defining crisp boundaries for natural entities, and in representing enclaves and exclaves, joints and separation between adjacent entities. More specific mereological issues are discussed in relation to *Key:building*. Complex buildings are composed of different parts, and alternative mereologies can be used to represent them. The intricacies of the built world are particularly visible in the discussion around *Key:wheelchair*, affected by the variety of barriers preventing wheelchair access. As the mapping of roads is the very foundation upon which OSM was developed, the conceptualisation of *Key:highway* causes major geometrical conflicts on how roads should be modelled and connected for routing purposes. For example, *highway=turning circle* conversations in the OSM Wiki website focus on how to geometrically represent a turning circle.

The modelling of large rivers in page *Tag:waterway=riverbank* constitutes another striking example of mereotopological and geometrical issues. Contributors believe that large rivers cannot be represented with a simple polyline, and suggest using multi-polygons. Additional ontological complexity is added by the necessity to represent river islands, and the joints with tributaries and distributaries, and the unclear boundaries between large rivers and seas. Analogous difficulties are encountered in the modelling of coastlines (*Tag:natural=coastline*). Although coastlines are *bona fide* boundaries, they are affected by seasonal phenomena, and they connect to rivers and lakes through complex and vague joints. Ontological difficulties occur also with bridges (*Key:bridge*), as they intersect with rivers, seas, and man-made infrastructures.

Cartographic negotiation

To achieve its goals, OSM applies the paradigm and ethos of free and open source software (FOSS) to the cartographic domain (Crampton 2009). After an initial co-existence of several competing and loosely connected map renderers, *Mapnik* has recently emerged as the dominant one.¹ Map renderers can interpret stylesheets, formal cartographic specifications defining rules about which entities are represented and how, at different scales

¹<http://mapnik.org>

and in different geographic areas. As stylesheets are complex and need high consistency, they are mainly manipulated by a small number of professional cartographers.² A tight interplay exists between stylesheets and the adoption of terms in the map, driven by the primacy of visual feedback returned by the renderers. Community guidelines invite contributors not to ‘tag for the renderer’ by entering incoherent data to obtain a certain visual style. A coordination problem exists between map renderers, map editors, and the OSM Wiki website, as contributors frequently lament disconnects and inconsistencies between them.

Renderers, because of their rule-based nature, often struggle with atypical instances that contributors encounter in their localities. Complex intersections in the road network can cause artefacts and garbled graphics on the rendered maps. Because of the bi-dimensionality of the data and the variation in urban layouts, the rendering order of buildings and roads is often debated, highlighting that universal rules—although needed by the technical infrastructure—clash with the complexity of geographic realities. To solve such cartographic negotiation, the actual tag usage statistics are also mentioned to support a particular solution among alternatives. Furthermore, as the stylesheets are mainly used for temperate regions in Europe and North America, contributors started to develop alternative visual styles for Polar and other regions.¹

Cultural and linguistic negotiation

Geographic conceptualisations are inevitably dependent on the information communities in which they originate. In general, *bona fide* entities are more likely to show cross-cultural invariance, while *fiat* entities, as they emerge from social processes, are likely show a higher degree of cultural dependence. The global, universalistic scope of OSM clashes with the heterogeneity of its contributors and objects of interest (Ballatore *et al.* 2013). Since the beginning, the project opted for a centralised conceptualisation in British English, to be adopted across the board in the world map. Unsurprisingly, this choice results in a fundamental tension. On the one hand, thanks to the hegemony of the language on the web and in software engineering, English is indeed an effective choice for a *lingua franca*. However, this choice forces contributors to perform complex semantic negotiations to adapt British English terms and concepts to the virtually infinite variety of local conceptualisations. Administrative and other *fiat* boundaries discussed in *Key:boundary* are strictly dependent on specific national legislation and political arrangements which are highlighted in cross-border mapping (Witschas 2010). As a result, ambiguities, difficulties in expressing accurate local knowledge, misinterpretation of terms, and other semantic failures are pervasive in the OSM Wiki website. In a recurring pattern, contributors initially attempt a universalistic conceptualisation, and then progressively fragment it into inter-operable regional or national schemas.

Initially, friction occurred between the two largest communities of German and British contributors to find a conceptualisation in English that can accommodate the differences in their administrative systems. Consensus was slowly built around a numeric code (*admin_level*), which is then associated with different terms in different countries. For example, level 6 corresponds to a *county* in the UK, a *provincia* in Italy, and a *district* in Japan. Similarly, *Key:highway* has seen much conflict since the inception of the project, when contributors attempt to express their local road classification with British English terms, even within the English-speaking world and its *freeways* and *motorways*. Ultimately, an intricate translation schema was established, illustrating the complexity

²<http://github.com/gravitystorm/openstreetmap-carto>

¹http://wiki.openstreetmap.org/wiki/Polar_Regions_Rendering_Issues

of this cross-cultural road classification.² Another major area conflict revolved around *Key:access*, which is a core part of OSM, particularly for routing, and is dominated by often incompatible national legal codes. Even in a narrow sub-domain such as the access to bus lanes, contributors visibly struggled with the complexity of the national road legislation of each country, finally leading to the definition of national schemas.

Overall, the issues encountered by contributors can be understood in the well-established framework of translation theories, particularly in the form of *problems of equivalence* between languages (Bassnett 2002, pp. 30-44). Hence, to obtain an overview of the problems of equivalence that cause friction in the process of negotiation, we considered a case study of a large community within OSM, that of Italian contributors. The Italian community maintains a page that hosts the translations of English terms into Italian,¹ which helped us identify the recurring semantic difficulties in the expression of local knowledge based on a conceptualisation in British English. The following problems of equivalence regularly spring up:

- *Local terms that do not have a precise translation in English.* The Italian *bar* can be roughly translated as ‘cafe’ in English, but corresponds to a different concept. The English term ‘bar’ usually defines what in Italian is called an ‘American bar’ or ‘lounge bar.’
- *Terms depending on practices, laws, and vocabularies of local institutions.* These include the road and land use classifications, postcodes, and addresses. Several terms of OSM are based on the UK Highway Code and, as Italy has a radically different code (‘Codice stradale’), this cultural translation is particularly problematic. *Courthouse* is a legal term specific to North America, and is used to map a ‘tribunale’, which has a roughly similar role in the Italian legal system.
- *Specific local terms that are merged into more general English terms.* ‘Trattoria’, ‘pizzeria’, and ‘ristorante’ refer to different types of restaurants, and they are all mapped with the term *amenity=restaurant*, losing local knowledge. In Italy, several police forces exist, with different roles and jurisdictions (‘Polizia,’ ‘Carabinieri,’ and ‘Polizia Municipale’), all mapped as *police*. Similarly, many specific types of accommodation² are classified in OSM as *guest houses*.
- *English terms referring to concepts that do not exist outside the English-speaking world.* Examples are *allotments*, *food court*, *commons*, *motel*, and *village green*. These terms are occasionally used incorrectly by non-Italian mappers to map entities located in Italy.

Definitional negotiation

Lexical definitions of terms are an essential normative tool to build a shared conceptualisation. Definitions constrain the intended usage of terms, specifying the necessary and sufficient conditions for their application, providing salient exemplars. The volunteered lexical definitions adopted in OSM are *intensional*, describing the internal structure of the concept, and *precising*, narrowing the usage of a common term for the specific context of the project. Lexical definitions can also be *stipulative*, defining novel meanings for an existing terms (Ballatore *et al.* 2013). As the OSM conceptualisation revolves around a set of terms, conflicts frequently arise about their lexical definitions.

Problems occur when definitions are *underspecified*, i.e., they lack salient details that

²http://wiki.openstreetmap.org/wiki/Highway:International_equivalence

¹http://wiki.openstreetmap.org/wiki/IT:Map_Features

²*Pensioni a conduzione familiare, Bed & Breakfast, agriturismo, affittacamere, locande, foresterie, etc.*

enable the correct interpretation of the terms being defined. By contrast, definitions can be *overspecified* when they include irrelevant or confusing intensional details. Many definitional conflicts spring from cultural and linguistic difficulties encountered when defining intrinsically vague geographic concepts, and from the polysemy of terms. These mostly intensional conflicts can result in *extensional* conflicts in the map, in which contributors disagree on whether individual objects fit a definition or not, called ‘tag wars’ (Mooney and Corcoran 2012).

The discussion on the conceptualisation of buildings in *Key:building* includes definitional conflicts. The definition of the term *tower* fosters a debate on its scope, in particular whether it should include sky-scrapers, communications towers, lighthouses, and turrets. This issue originates also from the unclear difference between towers as tall buildings or light structures, such as the Eiffel Tower. Contributors also struggled with the polysemy of the term *terrace*, whose definition might refer to a covered patio, a row of townhouses, a level paddy on a hillside, or to a terrain artificially graded to resemble rice fields. The conflict was solved with the precisising definition of ‘row of linked residential houses.’

Underspecified definitions of *church*, *chapel*, and *cathedral* caused confusion on how to distinguish between them, calling for differential definitions. Notably, *Tag:highway=unclassified* has provoked much prolonged discussion and conflict. This counter-intuitive term is used in the UK to classify minor public roads typically at the lowest level of the interconnecting grid network. Unclassified roads have lower importance in the road network than tertiary roads, and are not residential streets or agricultural tracks. As the term intuitively means roads that have not been classified, several contributors suggest that it is a very specific UK term, which cannot be easily understood by non-Britons. Because of the UK centrism of the OSM conceptualisation, discussions also emerged about *Key:shop* regarding the definitions for shops outside of the UK. For example, *chemist*, *pharmacy*, *drugstore*, and *medical-supplies*, all refer to businesses selling medical goods, but their usage for non-British contexts remains problematic.

Granularity negotiation and infinite knowledge

The collaborative geographic conceptualisation in OSM involves the categorisation of entities at different conceptual levels. Contributors often disagree on the conceptual *granularity*, which is the level of detail to be included in the conceptualisation. To understand this aspect, a useful distinction can be drawn from cognitive psychology applied to folk taxonomies, following Rorissa (2010). According to a seminal categorisation theory by Rosch (1999), abstract concepts belong to the *superordinate level* (e.g., furniture), while concrete concepts are located at the *basic level* (e.g., chair). More specific concepts, rarely used in day-to-day language, belong to the *subordinate level* (e.g., swivel chair). The identification of geographical objects is more challenging than the discrete table-top objects traditionally studied by cognitive psychologists. Notably, geographical objects are not just located in space, but they tend to be part of the Earth’s surface, and display great mereotopological complexity (Casati and Varzi 1999).

Along this vertical dimension from general to specific geographical concepts, OSM contributors encounter a number of problems. First, as the project has no imposed scope and scale, conflicts arise from divergent views of what parts of the geographical reality should be included in the map. This problem can be called *infinite knowledge*: given an object, a potentially infinite amount of information about it can be elicited. For example, given a school building, it is possible to start to describe it by its location and architectural structure in terms of main walls and roof. Then a huge number of details can be added about its materiality (specific model of windows, construction materials, etc.), and so

on *ad libitum*. The choice of what details should be included is completely arbitrary, and determined with respect to the desired application. Second, when a superordinate category is too generic, its usage is not constrained enough and different conceptualisations are likely to emerge. Subordinate categories can also cause problems, as they normally involve technical terms whose meaning is obscure or confusing outside—and often within—their domain of origin. Very specific subordinate categories are often little used and abandoned after a period of low usage. By contrast, categories at the basic level tend to cause less confusion, although they can indeed be the object of cultural and linguistic negotiation.

OSM contributors struggle with superordinate categories that include large numbers of basic level and subordinate level categories. The gradual inclusion of details in the *Key:access* page greatly complicated its conceptualisation. A taxonomy of categories of access rules was ultimately developed to provide an overview from the most generic category (*access*) to the most specific (e.g., *taxi* and *bus*). When needed, specificity can be greatly intensified. As OSM developed primarily to map stable objects in developed countries in peacetime, its basic geographic conceptualisation is not sufficient to support applications of OSM such as humanitarian responses and disaster management, often indicated as promising areas of application of geographic crowdsourcing (Haklay *et al.* 2014). Since the first Humanitarian OSM Team (HOT) initiative after the Haitian Earthquake in 2010, contributors produced an extended conceptualisation geared towards the needs of field workers, called the *Humanitarian Data Model*, which includes terms such as *practicability* for roads and *earthquake:damage* for buildings.¹ Similarly, in the case of the conceptualisation of wheelchair access (*Key:wheelchair*), contributors involved in the German project Wheelmap² increased the level of detail to provide more informational value for wheelchair users.

Semantic overload and duplication

Because developing a geographic conceptualisation is a remarkably complex task, several conflicts in OSM are caused by cases where the same terms are used to represent ontologically distinct aspects of an entity. Using the ontological terminology adopted by Gangemi *et al.* (2002), terms in OSM can represent *endurants* (entities wholly present in time, such as a bridge), *perdurants* (events and states in which an entity can be in), and *qualities* (properties of entities, such as size and colour). In this context, *semantic overload* occurs when the same term is used to refer to incompatible ontological dimensions, such as perdurants and qualities. *Duplication* is the opposite phenomenon, when the same ontological dimension is spread across different terms, reducing the clarity of the conceptualisation.

As buildings are the results of fast-paced social processes, conflicts arose about the term *building*, when contributors mixed the physical enduring aspects of building parts (e.g., wall, arch, hall, etc.), and the many perdurant social functions of buildings (e.g., school, shop, or house). The same term is also utilised to specify qualities of buildings (e.g., public, private, medical, or industrial). Contributors started out stating that the primary classification of buildings should be based on a type (e.g., hotel, house, or school). However, it is often the case that buildings are converted for a different use, although they preserve the original architectural structure. For example, many residential houses in London are converted to hotels, and contributors classify a building based either on its original usage embedded in the building's architecture, or on its current usage.

¹http://wiki.openstreetmap.org/wiki/Humanitarian_OSM_Tags

²<http://wheelmap.org>

Further confusion is added by usages that refer to states, without specifying structure or usage (abandoned, demolished, derelict, or vacant), and in the cases of multi-functional aggregates of buildings. Facing these issues, contributors reviewed the usages of the term, and attempted to split the dimensions into different terms, performing an ontological cleanup.

Semantic overload is also experienced by *Key:service*, which allows contributors to provide additional information about parts of highways, railways, and waterways. The term is currently used for diverse structures, including service yard entrances for trains, sidings, storage areas for rail cars, passing zones for trains in a railway yard, parking aisles, emergency accesses, and spur tracks. It also includes open man-made waterways, such as canals, used for transportation, irrigation, or the generation of electrical power. Non-native English speakers complained that these terms are too technical and hard to understand, while others point out the difficulties encountered when using the term in combination with other terms. The problem appears to stem from *Key:smoothness* being confused with *Key:surface*, which describes the physical surface of roads and footpaths. Some contributors claim that ‘the tag as proposed is irreparably broken.’ The next section concludes this investigation of the dimensions of incommensurability in OSM and points out directions for future research.

Discussion and conclusion

Taking the OSM Wiki website as a case study, we have examined the semantic negotiation that underpins crowdsourced cartography. In this process, through a quantitative and qualitative analysis, we have identified a pattern of rapid growth coupled with high contribution inequality, and we have shown that intense negotiation occurs to settle disputes about what should be represented in the map and how. The incommensurability in different ways of slicing the geographic world occurs along complementary dimensions. Our study shows that the negotiation unfolding in the OSM Wiki website is driven by a small core of highly motivated and productive individuals, with 1% of the most active users producing almost 50% of the content—a degree of inequality considerably higher than that found in comparable projects, such as Wikipedia (Ortega *et al.* 2008). Moreover, crucial components of the infrastructure, such as the map renderers and the stylesheets, are controlled by a handful of professional cartographers, who carry out centrally-planned crucial tasks.

This research indicates a word of caution against sweeping generalizations, acknowledging the tensions and contradictions playing out in crowdsourced cartography. OSM mappers are continuing a well-established tradition of universalizing cartography, innovating how the data is collected and encoded, rather than proposing deeper innovation in map-making. From our investigation, it emerged that consensus about the conceptualisation seems to reach an equilibrium, particularly in domains already described and formalised by existing international standards. However, such an equilibrium is an optical illusion, rather than a reality. Our inspection of talk pages, mailing lists, and discussion boards, rather than a clear recipe to solve semantic and ontological fractures, reveals the co-existence of several loosely coordinated processes at play, in which meanings are negotiated and articulated in a shifting network of human actors and tools.

By analysing the Italian conceptualisation, we have highlighted OSM’s problematic adoption of a conceptualisation in British English, which forces mappers outside the UK to express their local geographic realities through the eyes of an alien culture. OSM experiences a tension between its rhetoric of decentralisation, de-bureaucratisation, and

empowerment, and its technologically enforced Anglo-centrism, deeply built-in in its geographic conceptualisation. Nevertheless, OSM's negotiation of a conceptualisation can be considered as a case of successful coordination against the inevitable fragmentation of human language and cognition. Recurrently, contributors set off to find a universalistic conceptualisation and, after encountering insurmountable problems, resorted to more contingent and localized approaches. Notably, the community initially struggled to merge all national and regional road classifications into one single global conceptualisation, to then realise that such an approach would never succeed. Similarly, mappers adopted the conceptualisation defined by international standards, such as the CORINE Landcover nomenclature, simplifying it for their purposes.

Several promising directions might be taken in future work. The mappers who drive the development of the conceptualisation in OSM form a small group, who could be investigated from a social-scientific perspective, for example in terms of nationality, demography, and socio-economic status. From a theoretical perspective, while this study explicitly took a GIScientific approach, further explorations in the geographic conceptualisations can supply empirical evidence and insights to inform the debates in human geography and critical cartography. More pragmatically, a better understanding of the semantic and conceptual issues in crowdsourced cartography will support the assessment of data quality, which is a crucial and unsatisfied need of many information consumers (Ballatore and Zipf 2015).

The conceptualisation of the world in OSM is destined to remain a multi-authored, unfinished and transient product, suggesting, as Kitchin and Dodge (2007) argued, that maps 'have no ontological security, they are of-the-moment; transitory, fleeting, contingent, relational and context-dependent' (p. 340). Our analysis of the geographic conceptualisation constructed by volunteer mappers confirms the presence in the negotiation of recurring problems that are unlikely to find technological fixes, but require intense social and cultural negotiation to handle the inevitable fragmentation in how information communities view their worlds.

Acknowledgements

The authors thank Derek O'Callaghan (University College Dublin) for his valuable input, and the E.U. COST Action TD1202 for providing a platform for discussion. This research received no grant from any funding agency in the public, commercial or not-for-profit sectors.

References

- Agarwal, P., 2005. Ontological considerations in GIScience. *International Journal of Geographical Information Science*, 19 (5), 501–536.
- Ballatore, A., Bertolotto, M., and Wilson, D., 2013. Computing the Semantic Similarity of Geographic Terms Using Volunteered Lexical Definitions. *International Journal of Geographical Information Science*, 27 (10), 2099–2118.
- Ballatore, A., 2014. Defacing the map: Cartographic vandalism in the digital commons. *The Cartographic Journal*, 51 (3), 214–224.
- Ballatore, A. and Zipf, A., 2015. A Conceptual Quality Framework for Volunteered Ge-

- ographic Information. *Conference on Spatial Information Theory (COSIT)*. In press, 1–20.
- Bassnett, S., 2002. *Translation studies (3rd ed.)*. Oxon, UK: Routledge.
- Casati, R. and Varzi, A.C., 1999. *Parts and Places: The Structures of Spatial Representation*. Cambridge, MA: MIT Press.
- Crampton, J.W., 2009. Cartography: maps 2.0. *Progress in Human Geography*, 33 (1), 91–100.
- Dodge, M. and Kitchin, R., 2013. Mapping experience: Crowdsourced cartography. *Environment and Planning A*, 45 (1), 19–36.
- Elwood, S., Goodchild, M., and Sui, D., 2012. Researching Volunteered Geographic Information: Spatial Data, Geographic Research, and New Social Practice. *Annals of the Association of American Geographers*, 102 (3), 571–590.
- Gangemi, A., et al., 2002. Sweetening ontologies with DOLCE. In: A. Gomez-Perez and V.R. Benjamins, eds. *Knowledge Engineering and Knowledge Management: Ontologies and the Semantic Web*, Vol. 2473 of LNCS Springer, 166–181.
- Gerlach, J., 2014. Lines, contours and legends: Coordinates for vernacular mapping. *Progress in Human Geography*, 38 (1), 22–39.
- Goodchild, M., 2007. Citizens as Sensors: the world of volunteered geography. *GeoJournal*, 69 (4), 211–221.
- Guarino, N., Oberle, D., and Staab, S., 2009. What is an Ontology?. In: S. Staab and R. Studer, eds. *Handbook on Ontologies*. second ed. Springer, 1–17.
- Haklay, M., 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), 682–703.
- Haklay, M., et al., 2014. *Crowdsourced geographic information use in government, Report to GFDRR (World Bank)*. Technical report, London.
- Janowicz, K., 2012. Observation-Driven Geo-Ontology Engineering. *Transactions in GIS*, 16 (3), 351–374.
- Janowicz, K., Scheider, S., and Adams, B., 2013. A Geo-semantics Flyby. In: S. Rudolph, G. Gottlob, I. Horrocks and F. Harmelen, eds. *Reasoning Web. Semantic Technologies for Intelligent Data Access*, Vol. 8067 of LNCS Berlin: Springer, 230–250.
- Karapiperis, S. and Apostolou, D., 2006. Consensus building in collaborative ontology engineering processes. *Journal of Universal Knowledge Management*, 1 (3), 199–216.
- Kitchin, R. and Dodge, M., 2007. Rethinking maps. *Progress in Human Geography*, 31 (3), 331–344.
- Kuhn, W., 2009. Semantic engineering. LNGC, In: G. Navratil, ed. *Research Trends in Geographic Information Science*. Springer, 63–76.
- Kuhn, W., 2003. Semantic reference systems. *International Journal of Geographical Information Science*, 17 (5), 405–409.
- Leszczynski, A. and Wilson, M.W., 2013. Guest Editorial: Theorizing the GeoWeb. *GeoJournal*, 78 (6), 915–919.
- Mallery, P., 1992. *Bridge and Trestle Handbook (4th Ed.)*. Newton, NJ: Carstens Publications.
- Mooney, P. and Corcoran, P., 2012. Characteristics of heavily edited objects in OpenStreetMap. *Future Internet*, 4 (1), 285–305.
- Mooney, P. and Corcoran, P., 2014. Analysis of Interaction and Co-editing Patterns amongst OpenStreetMap Contributors. *Transactions in GIS*, 18 (5), 633–659.
- Ortega, F., Gonzalez-Barahona, J.M., and Robles, G., 2008. On the inequality of contributions to Wikipedia. In: *Hawaii International Conference on System Sciences, Pro-*

- ceedings of the 41st Annual*, 304–304.
- Perkins, C., 2014. Plotting practices and politics: (im)mutable narratives in OpenStreetMap. *Transactions of the Institute of British Geographers*, 39 (2), 304–317.
- Rad, H.S. and Barbosa, D., 2012. Identifying controversial articles in Wikipedia: A comparative study. *In: Proceedings of the Eighth Annual International Symposium on Wikis and Open Collaboration* ACM.
- Rorissa, A., 2010. A comparative study of Flickr tags and index terms in a general image collection. *Journal of the American Society for Information Science and Technology*, 61 (11), 2230–2242.
- Rosch, E., 1999. Principles of Categorization. *In: E. Margolis and S. Laurence, eds. Concepts: Core Readings*. Cambridge, MA: MIT Press, 189–206.
- Schuurman, N., 2006. Formalization Matters: Critical GIS and Ontology Research. *Annals of the Association of American Geographers*, 96, 726–739.
- Smith, B. and Mark, D.M., 2001. Geographical categories: An ontological investigation. *International Journal of Geographical Information Science*, 15 (7), 591–612.
- Smith, B. and Mark, D.M., 2003. Do mountains exist? Towards an ontology of landforms. *Environment and Planning B*, 30 (3), 411–428.
- Suh, B., *et al.*, 2009. The singularity is not near: Slowing growth of Wikipedia. *In: Proceedings of the 5th International Symposium on Wikis and Open Collaboration*.
- Warf, B. and Sui, D., 2010. From GIS to neogeography: ontological implications and theories of truth. *Annals of GIS*, 16 (4), 197–209.
- Witschas, S., 2010. Cross-Border Mapping: Geodata, Geonames, Multilinguality and More. Lecture Notes in Geoinformation and Cartography, *In: G. Gartner and F. Ortog, eds. Cartography in Central and Eastern Europe*. Berlin: Springer, 163–180.