

Cultural Diversity of Quality of Information on Wikipedias

Dariusz Jemielniak

New Research on Digital Societies (NeRDS) group, Kozminski University, Jagiellonska, 59, 03-301, Warszawa, Poland. E-mail: darekj@kozminski.edu.pl

Maciej Wilamowski

Faculty of Economic Sciences, University of Warsaw, Warwaw, Poland. E-mail: mwilamowski@wne.uw.edu.pl

This article explores the relationship between linguistic culture and the preferred standards of presenting information based on article representation in major Wikipedias. Using primary research analysis of the number of images, references, internal links, external links, words, and characters, as well as their proportions in Good and Featured articles on the eight largest Wikipedias, we discover a high diversity of approaches and format preferences, correlating with culture. We demonstrate that high-quality standards in information presentation are not globally shared and that in many aspects, the language culture's influence determines what is perceived to be proper, desirable, and exemplary for encyclopedic entries. As a result, we demonstrate that standards for encyclopedic knowledge are not globally agreed-upon and "objective" but local and very subjective.

Introduction

As of 2016, there were nearly 2.2 million Wikipedians (as contributors to Wikipedia are often called) worldwide who had edited at least 10 times¹ (the number of recently created accounts is significantly higher, as the English Wikipedia alone has nearly 29 million registered users).² This fact makes Wikipedia, or rather Wikipedias, easily the largest social movement of humankind—never before have so many people collaborated to create a common outcome (Jemielniak, 2014).

These Wikipedians want to create any librarian's dream: the systematized, free, and high-quality sum of all human knowledge. However, it is not entirely clear what is meant by high quality in an encyclopedia. Apart from simple accuracy

of information, are more references always good? How many photos should illustrate a given topic? What is the optimal length of an article? Should external links be used? These and other questions are highly relevant to a larger discussion on the quality of knowledge and information. It is also not clear as to what extent this presented information is culturally influenced. Is it possible, as we naturally tend to assume, that universal knowledge has some global standard of quality or that perhaps some languages and related cultures have a preference for visual representation or differing referencing standards? We attempt to uncover this information in this article.

Problem Statement

Based on data from major Wikipedias from different linguistic cultures, we want to discover whether the quality of information (that is, structured, contextualized, categorized data) has a universally standardized character. Addressing this problem is important, as that quality of information is often assumed to be neutral to culture and globally shared. However, we suspect that different language cultures may have different understandings of what constitutes the best quality of encyclopedic information. In particular, we speculate that the commonly expected reliance on external and internal links, the use of images, the proportion of words and characters to references, and other variables may differ between different language cultures, independent of the maturity and size of the project. If this is so, different language cultures have different standards in terms of what constitutes a good, informative encyclopedic entry. Consequences of this phenomenon would be substantial: it would mean that, for example, simply translating encyclopedias or textbooks perceived as exemplary (and of the highest quality) may not result in what is considered the highest quality in the language and culture into which the given text has been translated.

It is important to observe that, since we study different language versions of Wikipedias without a possibility of

¹See: <https://stats.wikimedia.org/EN/TablesWikipediaZZ.htm>

²See: <https://en.wikipedia.org/wiki/Special:Statistics>

Received July 29, 2016; revised January 19, 2017; accepted April 19, 2017

© 2017 ASIS&T • Published online 7 August 2017 in Wiley Online Library (wileyonlinelibrary.com). DOI: 10.1002/asi.23901

recognizing the actual country of origin of contributors, our conclusions are limited to what is commonly shared by specific language cultures, and not national cultures (in the cases when more than one nation speaks a given language).

Literature Review

The Wikipedia movement is largely self-organized and relies on an ad-hoc “division of labor, because people gravitate to work they enjoy, but little hierarchy” (Ayers, Matthews, & Yates, 2008, p. 217) and is generally described as nonhierarchical (Bruns, 2008). Unlike many other F/L/OSS (Free Libre and Open Source Software) projects (Lakhani & Wolf, 2003; Von Hippel & Von Krogh, 2003), participation in Wikipedia does not help anyone improve their professional status or career; writing an encyclopedia is not an experience highly valued for any particular profession, and contributions to Wikipedia are also often too compartmentalized to allow for an easy evaluation of one’s contributions as a whole.

However, the results of this largely uncoordinated and nonprofessional cooperation are amazing not only in terms of quantity but also quality. In a famous 2005 study published by *Nature*, Wikipedia was already going head-to-head with the Britannica encyclopedias and lexicons (Giles, 2005) in terms of the number of errors. Although at that time the results were not only surprising but also controversial, currently the situation has become much more clear. Wikipedia has grown to be ~50 times larger than Britannica in terms of word count,³ and Britannica went out of print in 2012. English Wikipedia’s overall linguistic readability receives mixed evaluations (Lucassen, Dijkstra, & Schraagen, 2012; Yasserli & Kertész, 2012), but the quality of information, even in areas as sophisticated as, for example, mental health, is consistently higher than other professional encyclopedic sources (Reavley et al., 2012). It is also much better referenced (Rivington, 2007), even if its perceived credibility is still problematic (Flanagin & Metzger, 2011).

Considering the previously mentioned unique character of Wikipedias, as well as its tremendous success, it is not surprising that Wikipedia has been a topic of many academic studies (Jemielniak & Aibar, 2016; Konieczny, 2016; Lih, 2009; Reagle, 2010). However, most of these studies thus far have focused on the English Wikipedia (Hara, Shachaf, & Hew, 2010), with some notable exceptions (Sundin, 2011). Most of the studies addressing the diversity of projects have also focused on the organizational and sociological aspects, rather than actual article output differences between languages (Konieczny, 2010; Morell, 2011; Tkacz, 2010).

The dominance of the English Wikipedia is so prevalent that it is sometimes even misleadingly referred to as “the Wikipedia” as if it was the only one out there. However, different communities on different Wikipedia projects (and other Wikimedia) differ significantly. Moreover, due to the highly nonhierarchical character of community governance, they create their own rules independently, both in terms of

procedures for decision-making and decisions about article standards, notability, etc. As a result, many observations made only on English Wikipedia are not as global as they are portrayed. It is quite clear, though, that there are significant cultural differences between different Wikipedias (Pfeil, Zaphiris, & Ang, 2006) and that they are related not only to language issues, for example, the ways in which editors interact with each other (Hara et al., 2010), but also to more general issues, for example, accepted information quality, collaborative innovation networks, etc. (Nemoto & Gloor, 2011; Stvilia, Al-Faraj, & Yi, 2009).

While there are many studies focusing on different types of topics related to the Wikipedia community and technology (Jemielniak, 2016; Müller-Birn, Dobusch, & Herbsleb, 2013; Osman, 2013; Yuan, Cosley, Welser, Xia, & Gay, 2009), including its social construction, demographics, and governance (Collier & Bear, 2012; Halfaker, Geiger, Morgan, & Riedl, 2013; Hill & Shaw, 2013; Jemielniak, 2015), as well as the content itself (Halavais & Lackaff, 2008; Rad & Barbosa, 2012; Rad, Makazhanov, Rafiei, & Barbosa, 2012; Reavley et al., 2012) and its internal rules and conflicts (Chen, 2010; de Laat, 2012; Geiger, Halfaker, Pinchuk, & Walling, 2012; Konieczny, 2009), most of them do not rely on a cross-project analysis of different Wikipedias.

Only recently have some authors turned to multicultural and multilingual analyses of the Wikipedia phenomena (Jemielniak, 2013; Yasserli, Spoerri, Graham, & Kertész, 2014). Yet there is still a significant gap in social research on Wikipedia communities that needs to be addressed, as studying Wikipedia projects as separate entities, while informative, does not allow the drawing of more general conclusions on national or linguistic preferences and styles of editors and readers. This is important because Wikipedia is arguably the best embodiment of what people of different language cultures perceive as good, well-presented information. By studying what different Wikipedian communities consider to be the paragon of encyclopedic entries, we delve into the cultural differences in understanding information quality in general. We already know that there are cultural differences in the styles of contributions on Wikipedia, correlating with Hofstede’s dimensions (Pfeil et al., 2006). It is also clear that the standards of quality are likely enacted through collaborative communities of practice and newcomer enculturation (Østerlund, Mugar, Jackson, & Crowston, 2016). However, the cultural differences in what is considered a well-articulated knowledge presentation are still undiscovered.

In our article, we attempt to partially address this issue by analyzing data in articles that are considered among the best in the eight largest Wikipedias (relying on Latin alphabet).

Materials and Methods

All analyzed Wikipedias have “Good” and “Featured” articles. Both categories are entirely community-driven: Wikipedians vote as to whether a given article is sufficiently well developed to deserve a “Good” status, and then another vote can decide about an even bigger promotion to “Featured.” These statuses are not written in stone and an

³https://en.Wikipedia.org/wiki/WP:Size_comparisons

TABLE 1. List of analyzed Wikipedias with number of Featured and Good articles.

Wiki	Number of articles		Ratio
	Featured	Good	
de	2,336	3,742	0,62
en	4,214	19,895	0,21
es	1,060	2,953	0,36
fr	1,411	2,158	0,65
it	574	303	1,89
pl	620	1,789	0,35
pt	630	640	0,98

article can be demoted as well, such as when the community decides that it no longer matches the standards of other distinguished articles. Thus, it could be assumed that the Good and Featured articles represent the cultural preferences of the studied communities quite well. However, data about these articles have not been collected in publicly available statistical analyses or in any accessible data sets, and primary research is required to collect them.

To study the proposed problem in detail, we designed a program that allowed us to analyze all Featured and Good articles from the eight largest Wikipedias (Table 1) and gathered the characteristics that statistically describe them, particularly the number of words, characters, images, references, and internal and external links. To the best of our knowledge, these are all the measures that are quantifiable and possibly related to the quality of the articles.

The number of chosen Wikipedias is deliberate: at the time of data collection, only these eight had a number of Featured and Good articles significant enough for our analysis to be possible.

The ratio of the number of Featured to Good articles is an interesting characteristic. In all analyzed projects, except Italian Wikipedia, this ratio is below 100%. It suggests that Featured articles are indeed more prestigious. A greater number of Good articles indicate that Wikipedia editors have lots of articles to choose from and are at liberty to apply rigorous selection criteria. It is also interesting to note that half of the analyzed Wikipedias have a ratio of the number of Featured to All articles between 0.08 and 0.1%. Additionally Wikipedias with the lowest numbers of articles also have lower values of this metric. It shows that the Featured articles are developed in a conservative manner—younger less-developed projects do not rush to promote articles to inflate their numbers.⁴

We used a crawler bot written in PHP to iterate over every single article page, Featured or Good, in all of the Wikipedias listed in Table 1 and gather relevant information

⁴It is worth mentioning that bot edits have played an increasingly important role in these measures and may distort the whole picture in the future. For instance, Swedish Wikipedia doubled between September 2012 and June 2013. As a result, when compared to most other Wikipedias with a similar number of articles, articles on Swedish Wikipedia have on average roughly half of the content of the French, German, Italian, Russian, and Spanish editions.

TABLE 2. Mean and median wordcount for the studied Wikipedias.

	Mean word count				Median word count			
	All articles		Top 300 articles		All articles		Top 300 articles	
	Featured	Good	Featured	Good	Featured	Good	Featured	Good
De	5,790	4,002	12,997	11,771	4,851	3,108	11,774	10,750
En	4,418	2,504	10,410	11,291	3,855	1,864	9,966	10,675
Es	7,862	3,342	13,756	9,570	6,733	2,643	12,408	8,515
Fr	8,785	4,959	16,714	12,453	7,732	3,908	15,295	11,106
It	6,224	4,275	8,456	4,311	5,479	3,533	7,753	3,546
Pl	4,840	2,446	7,197	6,065	3,991	1,870	6,193	4,879
Pt	6,545	2,652	9,636	4,075	5,712	2,100	8,652	3,580
Sv	4,617	3,499	5,059	4,232	3,921	2,897	4,292	3,643

on April 13, 2014. To obtain the number of words as accurately as possible, we counted all the words that were within the main body of the article in the paragraph HTML tags (<p>).⁵ To obtain the number of references, we counted the number of occurrences of the following fragment “<li id=cite” in the article source code. The numbers of images and internal and external links were gathered using the wiki API to obtain the most accurate data. The API provided by Wikipedia allows access to an article’s descriptive statistics for any article in the database (see: <https://en.wikipedia.org/w/api.php?action=help&modules=parse>). The resulting data set contains descriptive statistics for every Featured or Good article for the eight analyzed countries (Table 2). All the measures (means, medians, ratios, etc.) were calculated in R (<https://www.r-project.org/>).

Graphical analysis of the articles in our data set (Figure 1) shows that the number of words per article is distributed significantly differently for Good and Featured articles. All distributions are, as expected, strongly asymmetrical, with a positive skew. We do observe a large difference in kurtosis, which is higher for Good than Featured articles.

The differences in number of words between Good and Featured articles are expected, as the first category is a preliminary stage for the latter one. At the same time, we observe a positive skew of the distribution of the number of words for Good articles. Hence, there are many articles with a high number of words that are not promoted to the Featured category from Good despite their volume. This observation suggests that an analysis of the articles with the highest number of words within each group may yield interesting findings. Therefore, we decided to also calculate all the characteristics for the 300 articles with the highest number of word within each group (hence, Top 300).

We decided to use 300 as a basis for our top word-count analysis for the following reasons. The number

⁵We also tested the method with the total word count in the main body of articles and it did not affect our findings. However, we chose to use the word count of the main text, as it seems to be more representative for the quality of the article than the total number of words that can be inflated, for example, due to large tables.

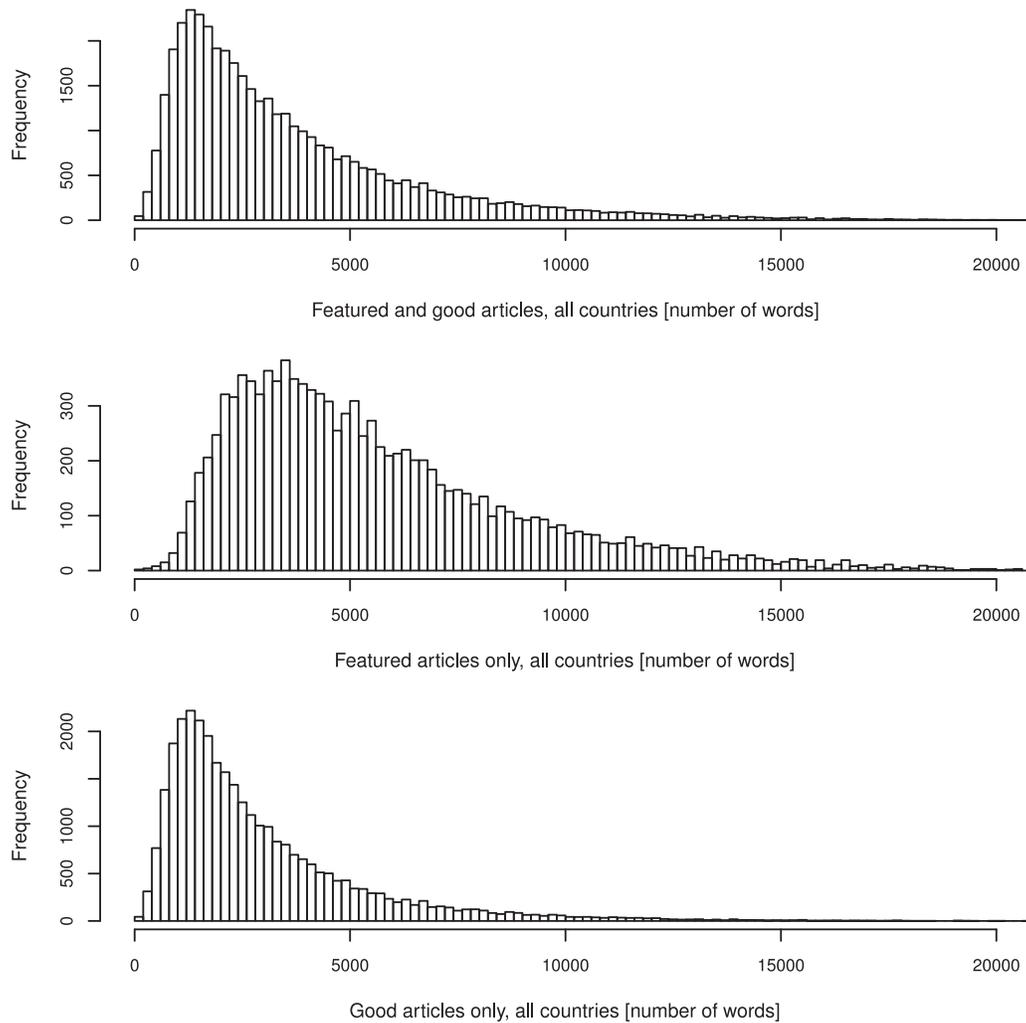


FIG. 1. Distribution of word count for good and featured articles collected from multiple Wikipedias.

should be as high as possible to capture the average characteristics of the articles with highest number of words and not the size of several outliers. On the other hand, as we wanted to focus on the right end of the distribution of the number of words, the number of analyzed articles should be smaller than half of the total number of articles. The differences in the number of articles between countries (Table 1) make meeting those two criteria simultaneously next to impossible, as there are almost 20,000 Good articles in English Wikipedia and only 300 in Italian. The number 300 is therefore a middle ground that best meets the set criteria in most cases.

The apparent smoothness of distributions, computed from our data set that contains values for articles from eight different Wikipedias, could falsely suggest that the distribution of the word count is similar across data sources. However, a closer look at the data (Table 2) clearly shows that this is not the case. The numbers of words differ between Wikipedias with respect to the values of mean and median as well as their difference and ratio. A more detailed analysis is needed and will be presented in the latter part of our article.

Findings

Result 1: Wikipedias Differ Substantially in Standards

Our first important conclusion is that Wikipedias are highly diverse in terms of editing patterns and content format. The largest Wikipedia communities vary significantly in what they perceive as the desirable format of an encyclopedic article. Although we have not been able to analyze the language style used across projects, for obvious reasons, we have observed even more fundamental differences on the level of proportions between visual and textual content, as well as in terms of article length. Irrespective of whether we observe Featured or Good articles or whether we analyze the top 300 articles or all of them, there is clearly substantial diversity between the projects in terms of the average count of words, characters, and images.

For example, when we analyze all Featured articles, French Wikipedia has about twice as many words, characters, or images per article compared to Polish or English Wikipedia (Figure 2). It is a clear leader within analyzed projects in both Featured and Good articles in all three measurements, with only one exception: for Good articles, it is



FIG. 2. Word, character, and image count in Featured and Good articles for both all articles and a subset of the top 300, where results are a proportion to the leader (value 1). [Color figure can be viewed at wileyonlinelibrary.com]

second to Italian Wikipedia in terms of the average number of images per article. The differences between Wikipedias are also visible in the absolute number of references and internal and external links (Figure 3). For all three characteristics, we can see that the numbers for the leading Wikipedias are approximately twice as large as for other projects (usually the Polish, Swedish, and German Wikipedias). We discuss the matter of references and links in more detail in result 3 of our article.

Our findings regarding diversity among Wikipedias are sustained even for the articles in the top 300 category. Interestingly, the results are not significantly different for the top 300. French Wikipedia is the clear leader, second only to German Wikipedia in one category: image count in Good articles. The only major difference in relative position is a better standing for English Wikipedia, most likely due to the long tail both in Good and Featured articles and the high refinement of its top 300. The disparities are nevertheless substantial. In fact, in the top 300 Featured articles, French Wikipedia has 50% more words than five of the eight analyzed Wikipedias.⁶

⁶It should be noted that the mere length of an article does not necessarily have to directly translate into its quality—after all, sometimes shortening an article may improve its quality and readability—but we are not analyzing these factors here.

These observations have led us to notice that countries also differ substantially in terms of their Good and Featured articles' saturation in images, references, internal links, and external links per word count (Figure 4). After all, for a reader, the sheer number of images (and other items) is less important in terms of article reception than their proportion to the text. We found that the differentiation of such saturation is also very high, and a standard deviation of these values amounts to 18–30% of the average value of a given variable (both for all and for the top 300 Featured and Good articles, as we have tested for top 300 previously). This means that Wikipedias clearly have different practices of desired proportions of these four elements. We are going to discuss this result in more depth and present the most significant of the other five differences observed in the studies material.

Result 2: Romance-Speaking Countries Like to Watch

Our preliminary analysis indicated that the differences between Wikipedias are particularly strong for images. The Wikipedias that lead in the absolute number of images stay at the top of the rankings even when we analyze image saturation despite the large differences in word counts.

Although we do not have a strong explanation for this phenomenon, we have observed an interesting correlation:

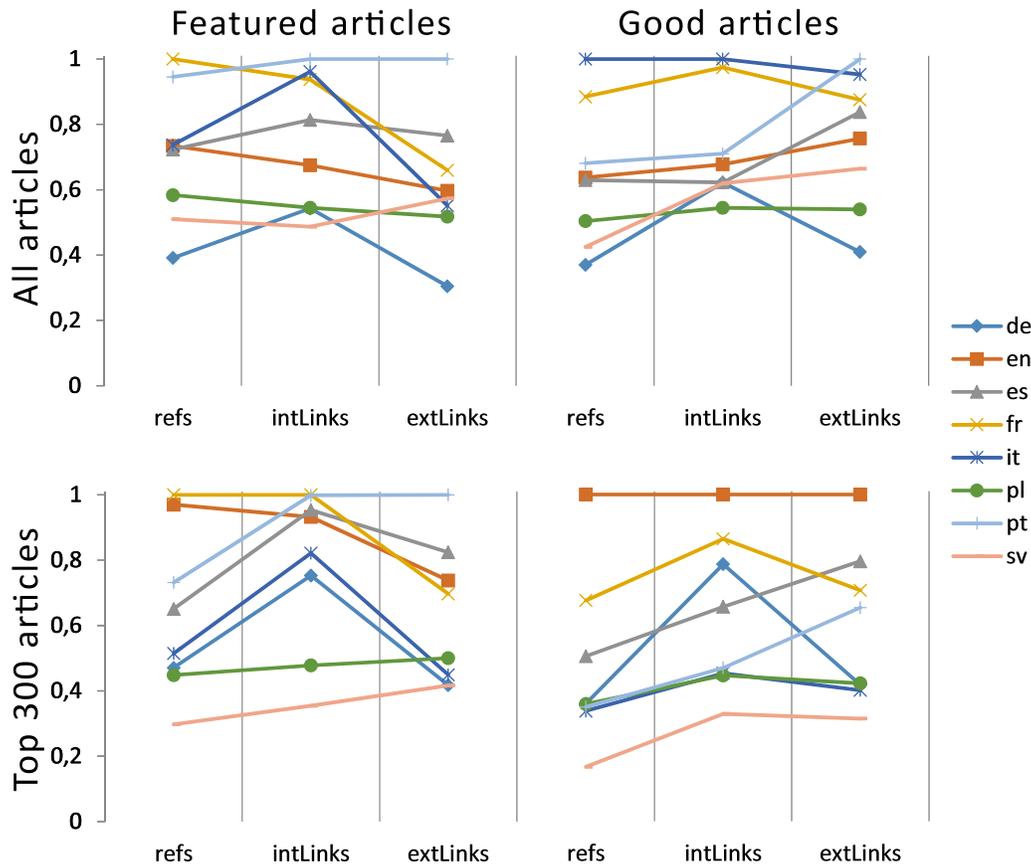


FIG. 3. Number of references, internal and external links in Featured and Good articles for both all articles and a subset of the top 300, where results are a proportion to the leader (value 1). [Color figure can be viewed at wileyonlinelibrary.com]

in terms of the total number of images per an average top 300 Featured article, Romance-speaking countries occupy the first, second, fourth, and fifth places. For the top 300 Good articles, they have the first and third places. When all Featured articles are considered, they take the top four positions. They also have the first two positions when all Good articles are considered. Romance-speaking countries also occupy the first two positions in terms of image saturation for all four categories.

More specifically, our data suggest that, for example, Italians apparently have a strong affinity for images, while the French tend to describe all phenomena in detail.

Because both the absolute average number of images and article saturation in images are clearly higher, we can speculate that the encyclopedic culture of these countries relies to a larger extent on a visual than a textual layer of communication.

Result 3: Different Approaches to References and Internal and External Links

As mentioned before, both the saturation and the absolute numbers of references and internal and external links are highly diverse (Figures 3 and 4). For references, there is no one big grouping of projects in any category: not in absolute numbers and, what may be more important, not in saturations.

The situation is fairly similar for external links, although we can see one grouping in saturation for the top 300 Good articles. For the internal links, however, the projects are much closer together, with very clear groupings in saturation. We discuss these differences in more detail below.

French Wikipedia is the leading project in terms of absolute number of references. On average, it has twice as many as the least advanced Wikipedias in this criterion. It is the leading project in this category for Featured articles, and it takes the second place for Good articles (Figure 3). In absolute numbers, it dominates six out of seven other projects in all four categories. However, the picture is very different when we look at saturations, where French Wikipedia is always very close to the mean saturation of references (Figure 4).

Interestingly, English Wikipedia, highly regarded for its references, places as low as fourth in the ranking for absolute number of references, when we look at all Featured and Good articles. It is, however, leading for Good articles and it is close to the leader for Featured in the 300 group (Figure 3). We explain this difference by the fact that English Wikipedia is more mature than other projects, and especially its top 300 Good articles are basically more developed than the top 300 Good articles in other projects. Moreover, it is leading in all four categories in terms of saturation (Figure 4). These results suggest that the reputation for references of English Wikipedia is indeed deserved, despite the fact that

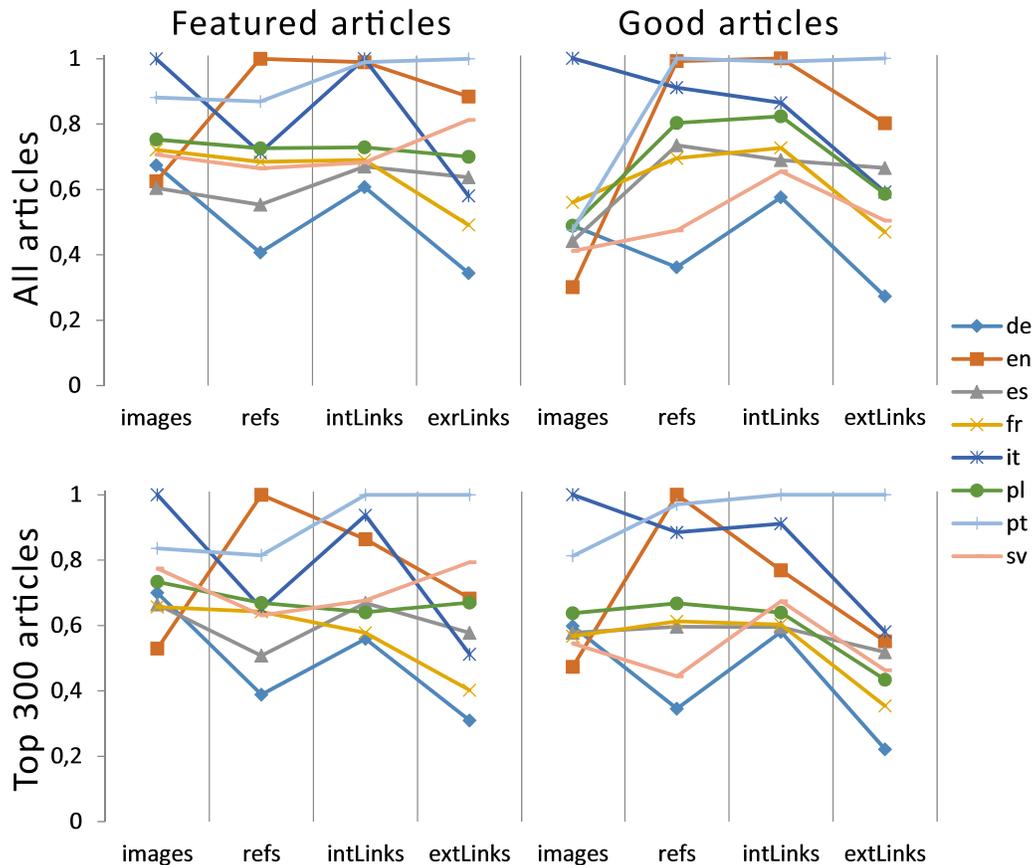


FIG. 4. Saturation in images, references, internal links, and external links per word count in Featured and Good articles for both all articles and a subset of the top 300, where results are a proportion to the leader (value 1). [Color figure can be viewed at wileyonlinelibrary.com]

not all English Featured and Good articles are extensive in volume. This result may indicate that English Wikipedia's Featured articles are well-developed and mature (as mentioned, sometimes a shorter article is more difficult to develop than a lengthy one).

Swedish Wikipedia is a visible outlier when we look at the references. It is in the last position or close to the bottom in absolute count in all four categories, and in saturation for Good articles. It is closer to the mean saturation of references only for Featured articles. The situation is fairly similar for external and internal links as well. This observation indicates that while the editors of Swedish Wikipedia may have expanded it in terms of sheer article count (which is a measure referenced by Wikipedians and a source of common pride), they have not been as successful in matching the referencing standards of other major projects. This result also signals that our choice of eight top Wikipedias makes sense: the smallest one studied (Swedish) is already an outlier, as it has made a stretch to match the leading ones in terms of article count.

In the analyzed group German Wikipedia is another project worth mentioning. Despite the fact that it is the second largest in terms of the number of All, Featured, and Good articles, and that it can be considered a well-developed and mature project, it is dead last in saturation of references in

all four categories. It also takes the last place for Featured and Good articles for all articles. Additionally, this project is also the last in saturation of both external and internal links in all four categories. We do not have a good explanation of this phenomenon, besides a cultural one: perhaps the editors of German Wikipedia basically put less emphasis on referencing.

When we look at external and internal links, the Portuguese Wikipedia is somehow peculiar. When we look at the saturation, Portuguese Wikipedia is at the top or very near to it in all four categories for both external and internal links. Moreover, for external links it is leading by a large margin. It would seem that its editors value links in general, and perhaps treat them as a useful extension of the article more than on other Wikipedias.

The rankings of other countries for external and internal positions are less distinctive. One minor, although interesting, observation is that internal links seem to be easier and more common to generate than external ones (the difference from the leader is smaller than in the case of references and external links), or standards in their creation are much more uniform across projects, possibly due to common suspicion of external links as possible tools of promotion. This result is also confirmed in Figure 4 when the saturation per number of words is considered (the distance to the leader is the

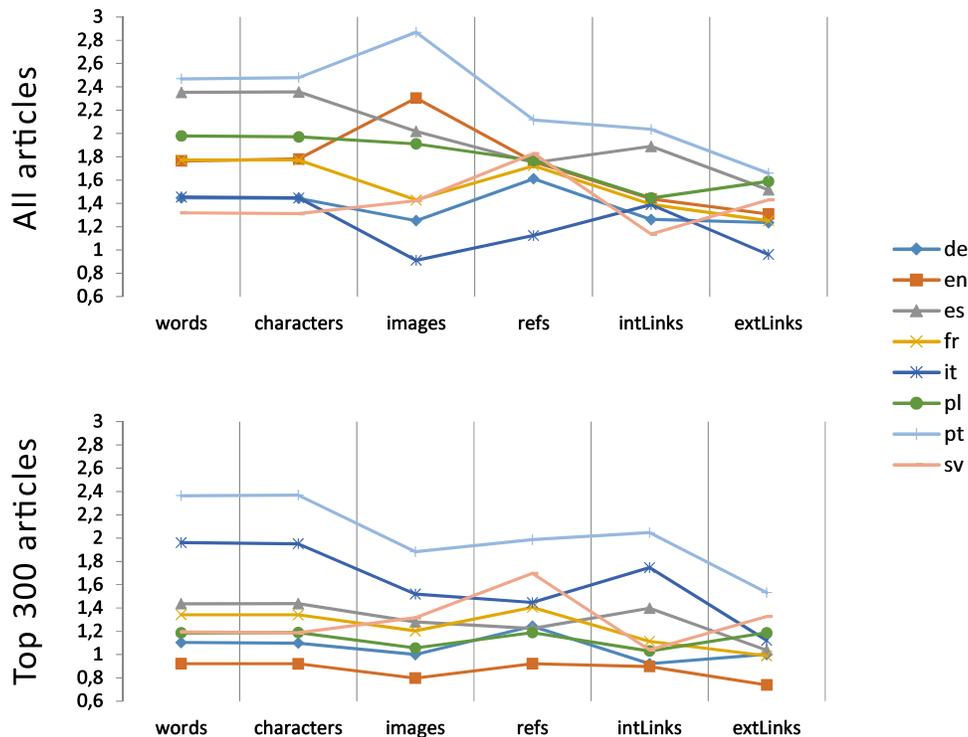


FIG. 5. Proportion of the number of words, characters, and images in Featured to Good articles. [Color figure can be viewed at wileyonlinelibrary.com]

smallest in internal links, and all results cluster around two values). The top six Wikipedias, when top 300 Featured articles are compared, have a relatively similar number of internal links (Figure 3). Only Polish and Swedish Wikipedias are clearly behind. This observation indicates that the best-developed articles stabilize across projects, with a relatively similar number of internal links.

The previous statement does not hold true for references and external links however, as in these areas projects are much more diversified. This may be somewhat surprising, especially in terms of references. After all, the rules of Wikipedias require proper sourcing of all information. Apparently, however, what is “proper” and sufficient differs across countries. To confirm these differences as project-, and possibly language culture-dependent, we compared the results for the top 300 Featured articles with all Featured articles (Figures 3 and 4). Although the ranking of Wikipedias occasionally varies, the shapes of the curves remain similar. We clearly see, for instance, that French Wikipedia is comparatively much more reliant on references than external links, quite unlike Portuguese Wikipedia (whose reliance on external links is even more visible when all Good articles are considered). Interestingly, the analyzed articles from German Wikipedia are the least saturated in references in all categories.

All in all, a relative position of Wikipedias is quite stable and there are only minor changes in rankings when ordered by average number of references and internal and external links and when compared in ranking by total number of words, characters, and images in all Featured articles

(Figures 2 and 3). French Wikipedia is the clear leader in four categories and is very close to the leader in one more. English Wikipedia’s lowest position in terms of an average word, character, and image count for all Featured articles may come as a surprise. Apparently, their Featured articles are typically two times smaller than the leader in the category.

Result 4: Featured and Good Articles Differ Substantially but Not in the Top 300

The substantial differences in the development of Wikipedias can also be observed between Featured and Good articles within projects. When all Featured and Good articles are considered, the Featured ones are significantly larger by 30–140% (Figure 5). This result suggests that Featured articles are much better developed on average than Good ones.

However, when just the top 300 Featured and Good articles are considered, the differences decrease strongly (with an exception of Italian and Portuguese Wikipedias). This suggests that different countries have different standards for Good and Featured articles: Italians and the Portuguese seem to grant less developed articles a Good status such that even the most developed in this category are still significantly less developed than the best Featured ones. For other Wikipedias, the proportions are less surprising. Again, English Wikipedia stands out: the top 300 Good articles on English Wikipedia are larger in terms of words, characters, and the number of images on average than the top 300 Featured articles (Figure 5). This could signify that the

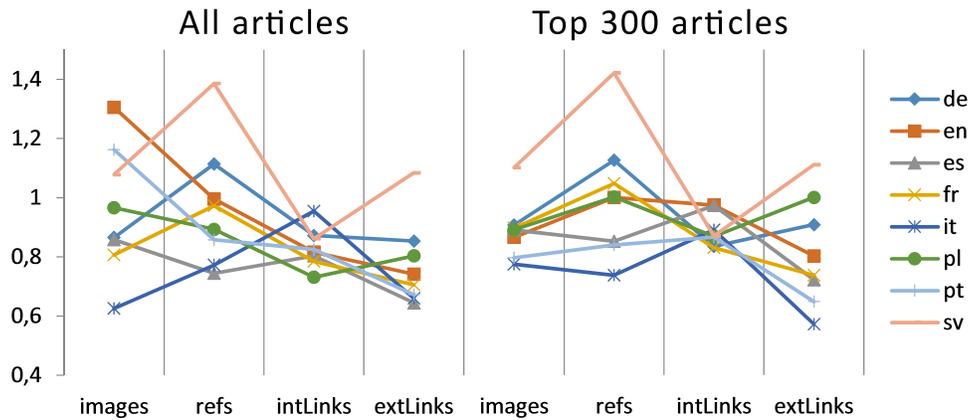


FIG. 6. Proportion of saturations of the number of images, references, internal and external links in Featured to Good articles. [Color figure can be viewed at wileyonlinelibrary.com]

transition of an article from the Good to the Featured category on English Wikipedia is not easy, and also that it is not enough to be well-developed to become Featured. As observed before, English Wikipedia appears to be the most mature and has higher saturation in many of the studied indicators, even if it has a smaller article size.

Result 5: Saturation Is Similar in Good and Featured Articles

Quite interestingly, we can observe significant differences between absolute length and the development of Featured and Good articles within projects. For instance, on average, a Featured article on Portuguese Wikipedia has 2.5 times more words than a Good one; on Polish Wikipedia, twice as many; etc. Nevertheless, when proportions of saturation between Featured and Good articles are considered, they are close to 1 (although typically below 1, as Figure 6 shows, which indicates that the Good articles still need to catch up to the Featured ones in terms of text development). Featured articles may be twice as long as Good ones, but the proportions of images, references, and internal and external links remain stable.

Additionally, Wikipedias are consistent in terms of their style: saturation in images, references, and internal and external links is relatively stable when Featured and Good articles are compared within each of the projects. When we analyze the saturation in the number of images, references, and internal and external links in Featured articles as a proportion of Good ones, that proportion is generally close to 1 (Figure 6). This result clearly indicates that there may be national (or rather, language culture-dependent) styles of Wikipedias; preferred proportions of images, references, and internal and external links to the word count of an article are project-dependent.

When the top 300 Featured and Good articles for each Wikipedia are considered, the proportions of saturation are also stable and even closer to 1. This suggests that the differences in saturations between Wikipedias are not a result of the stage of development or hours of work applied to the

articles, but rather come from the differences in language culture among the editors. Language culture styles of Wikipedias are quite consistent within and between Featured and Good, although there are significant differences between them.

Result 6: Saturation Is Stable and So Is the Ranking Order

The differences in the average sizes of Good and Featured articles between Wikipedias are, as described above, are substantial (Figures 2 and 3). However, it is interesting to note the degree of these differences. Despite large variations in the saturation of images, links, and references between projects (Figure 4), the ranking order of Wikipedias in the absolute numbers does not change to a great extent. The ranking of a project almost never changes by more than 1 within each category in terms of the number of words, characters, or images. Or to put it differently, the lines in Figure 2 do not often cross. The situation is similar for references and links, although the ranks are slightly more variable (Figure 3). For example, Italian Wikipedia has the highest saturation of images for top 300 articles by far but ranks only 4th and 5th in the total number of images for Featured and Good articles, respectively (in the top 300). What is more, the leaders in top 300 Good articles (Germany and France) have ~70% more images than Italian Wikipedia on average.

This result confirms and strengthens our observation that there are clearly different cultural norms between Wikipedias in terms of what is considered a paragon article and that the saturation in images, references, and links, as well as average word count, relies, to large extent, on local community perceptions of what is proper and right in this respect.

Discussion

In our study we have shown that the articles differ significantly between respective projects, not only in terms of word and character count but also in terms of relative saturation of images, references, and internal and external links. This is an important discovery, as knowledge is often

perceived as relatively easily translatable and universal.⁷ However, it seems that at least the standards of a high quality of knowledge vary substantially across language cultures. This phenomenon is perfectly exemplified through a closer look at Wikipedia. Although Wikipedias in different languages are often considered similar, especially if they are of comparative size, they still display substantial differences (Sundin, 2011). Some of these differences, such as the average top article length (in terms of its number of words or characters) can possibly be attributed to the stage of development of each of the projects. Similarly, there will be obvious differences in terms of areas covered and what is considered important and notable. For instance, most countries will cover their national heroes in greater detail.

However, some differences, such as the aforementioned saturation in images, references, or links (internal and external) per length of an article, which should theoretically probably be standardized and global, seem to be clearly linked to language culture preferences. These results indicate that quality standards, in terms of what is perceived as exemplary for an encyclopedic article, vary substantially depending on language and its culture. These findings are meaningful, as typically the quality of information is not perceived as something local and language culture-dependent. The results of our study may suggest that not only, as Barbara Czarniawska and Guje Sevón (Czarniawska-Joerges & Sevón, 2005) note, are ideas translated rather than simply diffused (Latour, 1993; Rogers, 1983), but also the very process of direct translation of knowledge is difficult (Estabrooks, Thompson, Lovely, & Hofmeyer, 2006) due to different cultural expectations.

More research is needed to more precisely confirm and verify the preferences of different language cultures for different information formats and standards. Our study is also obviously limited to the languages we have chosen, and the very fact that we rely on Wikipedia may already carry a certain bias (for instance, Wikipedias in all languages typically rely on more references than “professional” encyclopedias). Additionally, because our study relies only on a quantitative analysis, it is limited in its possible explanations of the studied phenomena and a follow-up qualitative study, focusing on the reasons for the described preferences, would definitely be enlightening, as would be approaching this fascinating topic from the fields of psychology and anthropology (Atran & Medin, 2008).

Conclusions

What makes a good encyclopedic article? What is a good representation of knowledge? Apparently, there is no one size of structured, contextualized, and categorized data (that is, knowledge) that fits all. Not only are there cognitive

⁷This common misconception has often been refuted by the researchers from Science and Technology Studies (STS), who have shown how seemingly universal knowledge is, in fact, very much culture-dependent and subjective (Latour, 1999; Latour & Woolgar, 1979).

differences between language cultures and clear preferences for abstract or concrete phrasing of information (Rau, Choong, & Salvendy, 2004), but the nature of understanding what makes a good description of a given phenomenon is culture-dependent.

As we have shown, different languages and cultures require different representations of knowledge, which go beyond just the issues of adequate language translation and are related more to the formatting, proportions, and the style of article construction. The very expectations of what justifies and grounds the presented sentences well enough to consider them solid and proven vary significantly. As a result, what one culture may perceive as meeting its high quality standards, another may not. The importance of this fact for computer-mediated communication and for library and information sciences is high.

Acknowledgments

Writing this article was possible thanks to a research grant from the Polish National Science Center (no. UMO-2012/05/E/HS4/01498).

References

- Atran, S., & Medin, D.L. (2008). *The native mind and the cultural construction of nature*. Cambridge, MA: MIT Press.
- Ayers, P., Matthews, C., & Yates, B. (2008). *How Wikipedia works: And how you can be a part of it*. San Francisco: No Starch Press.
- Bruns, A. (2008). *Blogs, Wikipedia, second life, and beyond: From production to produsage*. New York: Peter Lang.
- Chen, S. (2010). *Wikipedia: A republic of science democratized*. Albany Law Journal of Science and Technology, 20, 247–325.
- Collier, B., & Bear, J. (2012). Conflict, criticism, or confidence: An empirical examination of the gender gap in wikipedia contributions. In *Computer Supported Cooperative Work Conference Proceedings* (pp. 383–392).
- Czarniawska-Joerges, B., & Sevón, G. (2005). *Global ideas: How ideas, objects and practices travel in a global economy*. Malmö, Sweden: Liber & Copenhagen Business School Press.
- de Laat, P.B. (2012). Coercion or empowerment? Moderation of content in Wikipedia as ‘essentially contested’ bureaucratic rules. *Ethics and Information Technology*, 14, 123–135.
- Estabrooks, C.A., Thompson, D.S., Lovely, J.J.E., & Hofmeyer, A. (2006). A guide to knowledge translation theory. *Journal of Continuing Education in the Health Professions*, 26, 25–36.
- Flanagin, A.J., & Metzger, M.J. (2011). From Encyclopædia Britannica to Wikipedia. *Information, Communication & Society*, 14, 355–374.
- Geiger, R.S., Halfaker, A., Pinchuk, M., & Walling, S. (2012). Defense mechanism or socialization tactic? Improving Wikipedia’s notification to rejected contributors. In *Proceedings of the Sixth International AAAI Conference on Weblogs and Social Media* (pp. 122–129).
- Giles, J. (2005). Internet encyclopaedias go head to head. *Nature*, 438, 900–901.
- Halavais, A., & Lackaff, D. (2008). An analysis of topical coverage of Wikipedia. *Journal of Computer-Mediated Communication*, 13, 429–440.
- Halfaker, A., Geiger, R.S., Morgan, J., & Riedl, J. (2013). The rise and decline of an open collaboration community. How Wikipedia’s reaction to sudden popularity is causing its decline. *American Behavioral Scientist* (forthcoming). Retrieved from <http://www-users.cs.umn.edu/~halfak/summaries/The%20Rise%20and%20Decline/>.

- Hara, N., Shachaf, P., & Hew, K.F. (2010). Cross-cultural analysis of the Wikipedia community. *Journal of the American Society for Information Science and Technology*, 61, 2097–2108.
- Hill, B.M., & Shaw, A. (2013). The Wikipedia gender gap revisited: Characterizing survey response bias with propensity score estimation. *PLoS One*, 8, e65782.
- Jemielniak, D. (2013). Netnografia, czyli etnografia wirtualna – nowa forma badań etnograficznych. *Prakseologia*, 154, 97–116.
- Jemielniak, D. (2014). *Common knowledge? An ethnography of Wikipedia*. Stanford, CA: Stanford University Press.
- Jemielniak, D. (2015). Naturally emerging regulation and the danger of delegitimizing conventional leadership: Drawing on the example of Wikipedia. In H. Bradbury (Ed.), *The SAGE handbook of action research*. London, Thousand Oaks, CA: Sage.
- Jemielniak, D. (2016). Wikimedia movement governance: The limits of a-hierarchical organization. *Journal of Organizational Change Management*, 29, 361–378.
- Jemielniak, D., & Aibar, E. (2016). Bridging the gap between Wikipedia and academia. *Journal of the Association for Information Science and Technology*, 67, 1773–1776. <http://onlinelibrary.wiley.com/doi/10.1002/asi.23691/abstract>.
- Konieczny, P. (2009). Governance, organization, and democracy on the Internet: The Iron Law and the evolution of Wikipedia. *Sociological Forum*, 24, 162–192.
- Konieczny, P. (2010). Adhocratic governance in the Internet age: A case of Wikipedia. *Journal of Information Technology & Politics*, 7, 263–283.
- Konieczny, P. (2016). Teaching with Wikipedia in a 21st-century classroom: Perceptions of Wikipedia and its educational benefits. *Journal of the Association for Information Science and Technology*, 67, 1523–1534. <http://onlinelibrary.wiley.com/doi/10.1002/asi.23616/abstract>.
- Lakhani, K., & Wolf, R. (2003). Why hackers do what they do: Understanding motivation and effort in free/open source software projects. In J. Feller, B. Fitzgerald, S. A. Hissam & K. R. Lakhani (Eds.), *Perspectives on free and open source software*. Cambridge, MA: MIT Press.
- Latour, B. (1993). *We have never been modern*. Cambridge, MA: Harvard University Press.
- Latour, B. (1999). *Pandora's hope: Essays on the reality of science studies*. Cambridge, MA: Harvard University Press.
- Latour, B., & Woolgar, S. (1979). *Laboratory life: The social construction of scientific facts*. Beverly Hills, CA: Sage Publications.
- Lih, A. (2009). *The Wikipedia revolution: How a bunch of nobodies created the world's greatest encyclopedia*. New York: Hyperion.
- Lucassen, T., Dijkstra, R., & Schraagen, J.M. (2012). Readability of Wikipedia. *First Monday*, 17(9). Retrieved from <http://www.uic.edu/htbin/cgiwrap/bin/ojs/index.php/fm/article/view/3916/3297>.
- Morell, M.F. (2011). The Wikimedia foundation and the governance of Wikipedia's infrastructure historical trajectories and its hybrid character. In G. Lovink & N. Tkacz (Eds.), *Critical point of view: A Wikipedia reader*. Amsterdam: Institute of Network Cultures.
- Müller-Birn, C., Dobusch, L., & Herbsleb, J.D. (2013). Work-to-Rule: The emergence of algorithmic governance in Wikipedia *Communities and Technologies Conference* (pp. retrieved on 04/04/16 from: <http://herbsleb.org/web-pubs/pdfs/Muller-Birn-work-to-rule-2013.pdf>). Munich, Germany.
- Nemoto, K., & Gloor, P.A. (2011). Analyzing cultural differences in collaborative innovation networks by analyzing editing behavior in different-language Wikipedias. *Procedia - Social and Behavioral Sciences*, 26, 180–190.
- Osman, K. (2013). *The role of conflict in determining consensus on quality in Wikipedia articles*. Paper presented at the WikiSym conference in Hong Kong, China (5–7 August).
- Østerlund, C., Mugar, G., Jackson, C., & Crowston, K. (2016). Typologies of learning in open online collaborative communities. *International Reports on Socio-Informatics (IRSI)*, 13, 15–22.
- Pfeil, U., Zaphiris, P., & Ang, C.S. (2006). Cultural differences in collaborative authoring of Wikipedia. *Journal of Computer-Mediated Communication*, 12, 88–113.
- Rad, H.S., & Barbosa, D. (2012). *Identifying controversial articles in Wikipedia: A comparative study*. Paper presented at the WikiSym conference Linz, Austria (27–29 August).
- Rad, H.S., Makazhanov, A., Rafiei, D., & Barbosa, D. (2012). *Leveraging editor collaboration patterns in Wikipedia*. Paper presented at the Hypertext conference, Milwaukee (25–28 June).
- Rau, P.-L.P., Choong, Y.-Y., & Salvendy, G. (2004). A cross cultural study on knowledge representation and structure in human computer interfaces. *International Journal of Industrial Ergonomics*, 34, 117–129.
- Reagle, J.M. (2010). *Good faith collaboration: The culture of Wikipedia*. Cambridge, MA: MIT Press.
- Reavley, N., Mackinnon, A., Morgan, A., Alvarez-Jimenez, M., Hetrick, S., Killackey, E. . . Jorm, A.F. (2012). Quality of information sources about mental disorders: A comparison of Wikipedia with centrally controlled web and printed sources. *Psychological Medicine*, 48, 1753–1762.
- Rivington, J. (2007). Wikipedia guns for Britannica extermination. *TechRadar*. Retrieved from <http://www.techradar.com/news/internet/broadband/wikipedia-guns-for-britannica-extermination-132792>.
- Rogers, E.M. (1983). *Diffusion of innovations*. New York: Free Press.
- Stvilia, B., Al-Faraj, A., & Yi, Y.J. (2009). Issues of cross-contextual information quality evaluation—The case of Arabic, English, and Korean Wikipedias. *Library & Information Science Research*, 31, 232–239.
- Sundin, O. (2011). Janitors of knowledge: constructing knowledge in the everyday life of Wikipedia editors. *Journal of Documentation*, 67, 840–862.
- Tkacz, N. (2010). Wikipedia and the politics of mass collaboration. *Journal of Media and Communication*, 2, 41–53.
- Von Hippel, E., & Von Krogh, G. (2003). Open source software and the “private-collective” innovation model: Issues for organization science. *Organization Science*, 14, 209–223.
- Yasseri, T., & Kertész, J. (2012). Value production in a collaborative environment. *arXiv, preprint arXiv:1208.5130*. Retrieved from <http://arxiv.org/pdf/1208.5130v2011.pdf>.
- Yasseri, T., Spoerri, A., Graham, M., & Kertész, J. (2014). The most controversial topics in Wikipedia: A multilingual and geographical analysis. In P. Fichman & N. Hara (Eds.), *Global Wikipedia: International and cross-cultural issues in online collaboration*. Lanham, MD: Scarecrow Press.
- Yuan, Y.C., Cosley, D., Welser, H.T., Xia, L., & Gay, G. (2009). The diffusion of a task recommendation system to facilitate contributions to an online community. *Journal of Computer-Mediated Communication*, 15, 32–59.