

Understanding the Signature of Controversial Wikipedia Articles through Motifs in Editor Revision Networks

James R. Ashford

AshfordJR@cardiff.ac.uk
School of Computer Science and
Informatics, Cardiff University,
UK

Liam D. Turner

TurnerL9@cardiff.ac.uk
School of Computer Science and
Informatics, Cardiff University,
UK

Roger M. Whitaker

WhitakerRM@cardiff.ac.uk
School of Computer Science and
Informatics, Cardiff University,
UK

Alun Preece

PreeceAD@cardiff.ac.uk
Crime and Security Research Institute,
Cardiff University,
UK

Diane Felmlee

Dhf12@psu.edu
Dept. of Sociology Criminology,
Pennsylvania State University,
USA

Don Towsley

Towsley@cs.umass.edu
College of Information and Computer
Sciences, University of Massachusetts
Amherst, USA

ABSTRACT

Wikipedia serves as a good example of how editors collaborate to form and maintain an article. The relationship between editors, derived from their sequence of editing activity, results in a directed network structure called the revision network, that potentially holds valuable insights into editing activity. In this paper we create revision networks to assess differences between controversial and non-controversial articles, as labelled by Wikipedia. Originating from complex networks, we apply motif analysis, which determines the under or over-representation of induced sub-structures, in this case triads of editors. We analyse 21,631 Wikipedia articles in this way, and use principal component analysis to consider the relationship between their motif subgraph ratio profiles. Results show that a small number of induced triads play an important role in characterising relationships between editors, with controversial articles having a tendency to cluster. This provides useful insight into editing behaviour and interaction capturing counter-narratives, without recourse to semantic analysis. It also provides a potentially useful feature for future prediction of controversial Wikipedia articles.

KEYWORDS

Network Motifs, Wikipedia, Controversy, Complex Networks

This paper is published under the Creative Commons Attribution 4.0 International (CC-BY 4.0) license. Authors reserve their rights to disseminate the work on their personal and corporate Web sites with the appropriate attribution.

WWW '19 Companion, May 13–17, 2019, San Francisco, CA, USA

© 2019 IW3C2 (International World Wide Web Conference Committee), published under Creative Commons CC-BY 4.0 License.

ACM ISBN 978-1-4503-6675-5/19/05.

<https://doi.org/10.1145/3308560.3316754>

ACM Reference Format:

James R. Ashford, Liam D. Turner, Roger M. Whitaker, Alun Preece, Diane Felmlee, and Don Towsley. 2019. Understanding the Signature of Controversial Wikipedia Articles through Motifs in Editor Revision Networks. In *Companion Proceedings of the 2019 World Wide Web Conference (WWW '19 Companion)*, May 13–17, 2019, San Francisco, CA, USA. ACM, New York, NY, USA, 8 pages. <https://doi.org/10.1145/3308560.3316754>

1 INTRODUCTION

Wikipedia has become a tremendous platform for crowdsourcing knowledge, representing a cornerstone of the World Wide Web [7]. It allows the "wisdom of the crowd" [24] to potentially emerge, providing intelligence on a vast range of topics [4]. However, complex dynamics support the emergence of content, since the formation of Wikipedia articles involves both human cooperation and human conflict, based on the extent of convergent and divergent views. Narrative and counter-narrative frequently jostle for presence in articles, representing a source of friction that is seen through editor interaction [22] and in the semantics of article content [20]. Wikipedia conveniently provides a list of controversial content that are labelled by the Wikipedia community themselves¹.

In an age of misinformation [1, 17, 26], understanding characteristics of controversial articles has increased in importance. Because of the controversial nature of some topics, the narrative in a Wikipedia article may contain misleading information that stops a neutral consensus emerging. Prior work in this area has established insights such as the predictability of controversy from editor behaviour [20], such as deletions, reversions, and statistics from the collaboration network, prediction of article quality taking insights

¹https://en.wikipedia.org/wiki/Wikipedia:List_of_controversial_issues

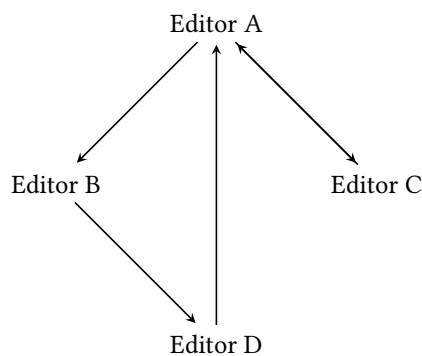


Figure 1: A network generated using the editor order A, B, D, A, C, A , from newest edit to oldest. Each editor is characterised by a letter A to D , and each occurrence in the list marks a single revision of the article by the corresponding author. An edge is formed between the current editor and its adjacent neighbour in the sequence, forming the directed edges $(A, B), (B, D), (D, A), (A, C), (C, A)$.

from multiple models [29], and interactions between users, bots, admin and pages [9]. There has also been a number of different types of network developed to assess Wikipedia articles, including collaboration networks [6] that capture the positive or negative relationship between editors, edit networks that capture “undoing” of edits by a third party [12] and affiliation networks [10].

Our focus in this paper is to further understand the relationship between small groups of editors, as induced by their editing sequences, by using a *revision network*. This does not require information on the nature of the editing undertaken - it simply captures the ordering in which editing occurs and is therefore a simple metric to infer. Editors are represented by nodes and a directed edge from node A to B indicates that “*Editor A edits the article after Editor B*” (see Figure 1). From this, we seek to determine the extent to which controversial articles exhibit a distinctive signature relative to those that are deemed non-controversial.

There has already been some consideration of revision networks in the literature [8, 11], where more recently the emphasis [27] has been to combine them with other network representations. However, given the fundamental nature of revision networks, it is interesting to question the extent to which they hold sufficient information to characterise controversial Wikipedia articles. Currently this is not well-understood, and motivates our work.

Hypotheses

We hypothesise that *interaction differences between small sub-groups of Wikipedia editors is sufficient to distinguish between controversial articles and non-controversial articles*. To address this we consider the extent to which revision networks

of Wikipedia articles have different local induced substructures based on their having controversial classifications. Our approach is based on techniques from complex networks [15, 16], that have been successful in classifying diverse and complex biological networks based on their latent induced subgraphs.

To achieve this at scale, and in contrast to previous literature [9, 20, 22], we assess a relatively large sample of Wikipedia articles, involving over 21,000 Wikipedia articles, by determining their subgraph ratio profiles. Each such profile represents the under and over representation of induced triads in the revision network of a Wikipedia article using 13 dimensions of connected triads, while also normalising for differences in network size.

Understanding the relationship between Wikipedia articles is important, as it allows us to determine the extent to which they may cluster. The sample we are considering represents 21,631 articles each represented in 13 dimensions. Therefore we perform dimensionality reduction, and project the subgraph ratio profiles into lower dimensional spaces. This allows us to examine the relationship both within and between controversial and non-controversial articles.

The results draw attention to distinctive clustering patterns concerning controversial Wikipedia articles. We further assess the results by examining the correlation with a range of variables, allowing to understand the role of substructures. The findings reaffirm that the sequence of editing provides an important mechanism to understand Wikipedia articles, independent of an article’s topic and without recourse to semantic analysis.

2 RELATED WORK

Understanding the content of crowd-sourced platforms such as Wikipedia and the behaviour of their contributors is of wide research interest [24]. Wikipedia represents a dynamic network of articles with a structure resembling that of the World Wide Web [30], whereby dominant articles act as connectivity hubs. Dynamics also exist within the formation and maintenance of individual Wikipedia articles, through open and collaborative editing.

Interactions between editors range from positive to negative, where debates and arguments lead to different patterns of revision (e.g., [14, 20]), capturing behaviours such as vandalism [19] and the propagation of disinformation [13]. Characterising articles and contributors through revisions provides a means for Wikipedia to manage and review its content. This is potentially labour intensive and has led to interest in creating and exploiting methods to detect issues (e.g., [2, 23]).

Controversial articles have become an increasing point of focus, and characterised as such by Wikipedia. Automated methods for classifying articles have received much attention




Figure 2: 13 possible combinations of connected triads in directed networks.

(e.g., [21, 27, 29]). The associated revision log for Wikipedia articles has been shown to provide a basis to examine potential controversy through examining the collaborative behaviour of individual editors within an article [23] or across multiple articles [28]. An article’s revision log identifies the structure underlying temporal interactions [27], and provides insight into how articles and contributors’ habits may evolve over time [9]. Features from the aggregation of this, such as number of edits, revision, and previous version restorations have been shown to correlate (e.g., [23]).

Treating the revision log as a network between editors [22] has been shown to provide additional useful features using graph theory and social network analysis techniques (e.g., [21]). This has ranged from global features such as the degree distribution (e.g., [23]), through to analysis of local sub-structures concerning the articles with which editors interact (e.g., [28]).

However, there has been little investigation of controversial articles based on the under or over representation of local-substructures. Called network motif analysis, this approach originated from biology [5, 15], and has been used to good effect in characterising other complex networks, including technology (e.g., [18]). In terms of Wikipedia, motif analysis has been used to determine how articles point to each other [30] and in assessing interactions between editors and different Wikipedia articles [28]. Our contribution is two-fold: firstly to use motif analysis to understand the fundamental revision networks, and secondly to consider the role of motifs in differentiating between controversial and non-controversial articles.

3 METHODS

Dataset

We collected revision history logs and article meta-data of a sample of Wikipedia articles ($N = 21,631$) through Wikipedia’s web API². The revision logs contain time-series events and meta-data attributing the revision to a particular user at a given time. Within this set of articles, a subset ($N = 2,661$) are considered to be ‘controversial’ as they are listed in Wikipedia’s “List of controversial issues”¹. The remaining articles ($N = 18,970$) are random articles that are not contained in the controversial issue list to serve as a basis for comparison. These were taken from an original

sample of 23,000 articles (20,000 non-controversial and 3,000 controversial), from which articles were removed if they did not contain sufficient information for motif analysis.

Network Construction


For each article, we construct a revision network (in the same manner as [11, 27]) where nodes represent unique editors and directed edges show that an editor added a revision after another editor. We traverse the revision log list to build a network that spans the article’s lifetime, adding nodes and edges as they appear in each event. Specifically let the revision network of a Wikipedia article be defined by $G = (V, E)$, where each editor is represented by a node $v \in V$. An edge $(v_i, v_j) \in E$ indicates that editor v_i edits the article after editor v_j . This excludes self-loops and editor v_i editing after editor v_j multiple times does not result in a multi-edge.

A simple example can be found in Figure 1, which describes how the network is constructed. Two extreme examples from the dataset shown in Figures 3a and 3b. This form of representation is potentially useful as large articles do not typically follow a linear or incremental structure. For example, it is highly likely that users will restore work back to an earlier revision should a revision become vandalised or irrelevant. Editors are likely to refer back to previous editor’s work. These behaviours result in complex sequential patterns that are captured through revision networks.


Network Motif Analysis

Network motif analysis focuses on determining the under or over representation of induced subgraphs [5, 16], as compared to an alternative sample of graphs (i.e., a null-model that acts a relevant baseline for comparison). Dyads, triads and tetrads are often considered as induced subgraphs such over and under-represented induced subgraphs are called network *motifs* or *anti-motifs*. In this analysis, we examine each article’s revision network using triads, representing how all possible triples of editors may sequentially interact. Triads are sufficiently large enough to capture both direct and indirect reciprocity between editors, while not being of a scale that is impeded by combinatorial complexity - there are 13 possible connected triads, shown in Figure 2. The coded names listed in Figure 2 are provided as part of the convention used in the triad census algorithm [3].

²<https://en.wikipedia.org/w/api.php>



(a) A revision network taken from the non-controversial *The Web Conference* Wikipedia article.



(b) A revision network taken from the controversial *Brexit* Wikipedia article. The large volume of editors produces a highly dense cluster of interactions.

Figure 3: Revision network of two articles - one non controversial and one controversial.

For each article, we calculate the subgraph ratio profile (SRP) as defined by Milo et al [15]. This accounts for variations in network size. This is achieved by determining the relative abundance of each type of triad compared to random graphs generated by the null-model. For each type of triad i , we firstly calculate Δ_i :

$$\Delta_i = \frac{(Nreal_i - \langle Nrand_i \rangle)}{(Nreal_i + \langle Nrand_i \rangle) + \epsilon} \quad (1)$$

where $Nreal_i$ is the number of such triads observed in the graph under observation, and $Nrand_i$ is the average frequency of such triads seen across the sample of networks composed by the null model. In this case the null model uses 100 random graphs with the same number of nodes and edges as the graph under observation. The value of ϵ is set to 4 [15, 25] to prevent the result from being misleadingly large when a subgraph rarely appears. This process is repeated for each triad i and normalised across triads to form the *subgraph ratio profile* (SRP) for a given network. The i^{th} SRP, denoted SRP_i , denotes the extent of under or over representation of the triad i , and is defined by:

$$SRP_i = \Delta_i / (\sum \Delta_i^2)^{1/2}. \quad (2)$$

Article comparison and dimensionality reduction

The SRP composed for each article provides a 13-dimensional vector whose components indicate the extent of triad representation relative to networks in the null model. To assess these collectively, we use principal component analysis

(PCA) that allows the SRPs to be considered in a lower dimensional space. We apply three and two dimensional PCA. Finally, we make comparison with a number of external variables (number of editors/nodes, age of article, and revision rate) to understand potential correlations with motifs.

4 RESULTS

Subgraph Ratio Profiles

We first examine the SRPs that arise from both controversial and non-controversial articles in isolation (Figures 4 and 5 respectively). To find significant triads, we use a cutoff value of +0.3 and -0.3. We determine that controversial articles are strongly represented by triads 111D, 111U and 201, which attain average SRP scores of 0.382, 0.375, 0.372, with relatively low dispersion (SDs of 0.136, 0.149 and 0.124 respectively). Interestingly, together these represent a chain of three nodes, where one edge is reciprocated, with the other edge covering all possible directional types (i.e., reciprocated, directed in, directed out).

In contrast, the results for the non-controversial articles provide a different profile. Here 021D and 021U are significantly under represented (average SRP scores of -0.511, -0.485), albeit with higher standard deviations present (SDs of 0.192, 0.2). Interestingly these anti-motifs (021D and 021U) relate to a lack of subgraphs where directed edges either emanate from or are received by a single node in the triad. Such configurations relate to the role of a mediating editor




Figure 4: Subgraph ratio profiles of all controversial articles. Average displayed in red.




Figure 5: Subgraph ratio profiles of all non-controversial articles. Average displayed in red.




Figure 6: Subgraph ratio profiles of all articles, with controversial articles displayed in red and non-controversial in black.

that may be presented with or respond to the editing of others. In other words, such mediators have a reduced role in non-controversial articles.

The comparison between these subgraph ratio profiles is shown in Figure 6. These profiles are quite distinct. We also calculate the Pearson correlation coefficient for each distinct pair of articles in three groups - controversial articles, non-controversial articles and all articles. Controversial articles provide the greatest correlation to each other ($M=0.41193$, $SD=0.41262$). Non-controversial articles have a lower mean correlation ($M=0.37498$, $SD=0.36749$) which is similar to the result when considering all articles together ($M=0.37569$, $SD=0.36844$).

Principal Component Analysis

We analyse the 21,631 subgraph ratio profiles to determine the relationship in terms of relative clustering. We apply PCA in order to reduce the 13 dimensions of the SRPs down to a more manageable form. We initially project the SRPs into 3-dimensional space for clarity, as seen in Figure 7. This presents a distinctive region where controversial articles are dominant. This provides evidence for a distinction between the controversial and non-controversial articles, consistent with the variation in motifs identified in the previous section.

The PCA coefficients (Table 1) that define the three dimensions reveals that principal component one (x axis in




Figure 7: PCA scatter plot representation of the 13-point feature vector in 3D clustering space.

Figure 7) primarily depends on triads 111D and 111U. Principal component two (y axis in Figure 7) primarily depends on triad 021C. The third principal component primarily depends on triads 021C and 030C. However, we also note that when represented in the three dimensional space, the revision networks have limited dispersion in the third dimension (i.e., vertical dimension as plotted).

Calculating the percentage of explained variance by principal component confirms that the first principal component produces 53.7% of the shared variance, the second produces 22.3% and the third produces the least with 6.7%. This confirms that the third principal component provides a limited contribution to representation of the total variance across the significant ratio profiles. This supports representation through two principal components, as plotted in Figure 8, with the relative composition of each principal component being near identical to PC-1 and PC-2 in Table 1. As anticipated, this is similar dependency on the first and second primary components in the three dimensional representation.

Representation in two dimensions further clarifies the distinction between controversial and non-controversial revision networks. In particular, from Figure 8 we note that both classes of article exhibit a similar maximum and minimum range against principal component two, which is primarily defined by the linear path between three nodes (021C). However, it is the variation in the first principal component, dominated by 111D and 111U, which represent linear paths with reciprocation on one edge, that differentiate the non-controversial from controversial. High values in principal component one correlate with controversial articles - in other

Table 1: PCA coefficients displaying the strongest triads

	021D	021U	021C	111D	111U	030T	030C	201	120D	120U	120C	210	300
PC-1	0.332	0.321	-0.126	0.509	0.522	0.145	0.205	0.318	0.078	0.079	0.204	0.117	0.065
PC-2	0.427	0.438	0.643	-0.278	-0.242	0.207	0.104	-0.052	0.059	0.06	0.047	0.074	0.06
PC-3	0.074	0.127	-0.609	-0.398	-0.332	0.136	0.441	0.115	0.131	0.130	0.19	0.167	0.112

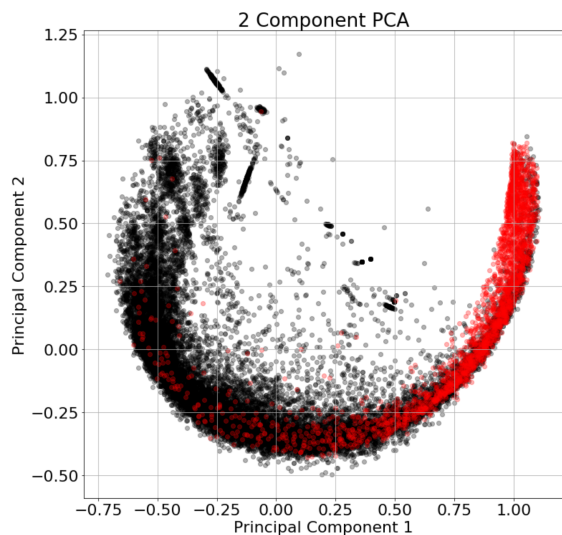


Figure 8: PCA scatter plot repression of the 13-point feature vectors in 2D clustering space.

words, controversial articles exhibit more reciprocation on top of possible linear paths.

Additional Metrics

We examine the relationship between revision networks and primary external variables (number of editors/nodes, age of article, and revision rate) using motifs. Specifically, using the dimensions of two-dimensional PCA analysis, we examine the correlation with the external variables, and how this differs between controversial and non-controversial articles. The results are shown in Figures 9, 10 and 11.

The greatest differences between controversial and non-controversial articles occur with respect to article age (Figure 10). Here, controversial articles with high age cluster against high values of principal component one, and to some extent this occurs for principal component two. This contrasts against the clustering seen for non-controversial articles.

Figure 10 also shows that controversial articles have a tendency to be older. The high values in principal component one and two which align with dense clustering of controversial articles show that while such articles accumulate the linear revision path between authors (021C which dominates




Figure 9: Scatter plot of each principal component combined with node count for both non-controversial and controversial articles.




Figure 10: Scatter plot of each principal component combined with article age for both non-controversial and controversial articles.

principal component two), controversial articles also accumulate instances of linear paths where one edge is reciprocated (i.e., 111D and 111U which dominate principal component one).

5 DISCUSSION

Motif analysis of revision networks gives insight into how the temporal editing relationship between small groups of

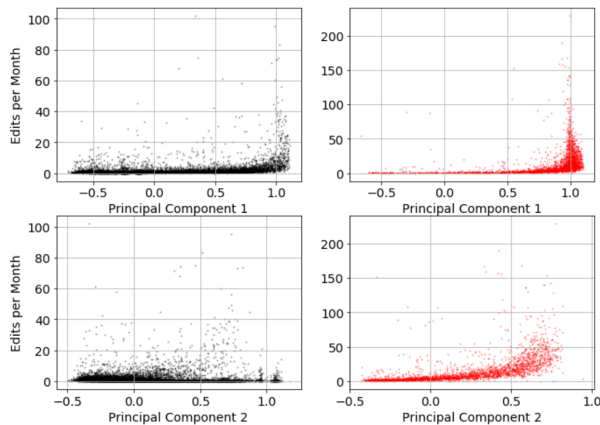


Figure 11: Scatter plot of each principal component combined with edit rate (mean number of edits per month) for both non-controversial and controversial articles.

Wikipedia authors create signatures that allow controversial articles to be distinguished. In contrast to previous work, we have investigated this using a relatively large sample of Wikipedia articles, where distinct patterns emerge. This provides strong support for our hypothesis, and reaffirms the importance of the revision network as a simple but fundamental element in editing Wikipedia.

Through motif analysis, we identify that reciprocation on linear paths among triads of editors in the revision network is over represented in controversial Wikipedia articles. These motifs are defined by the triads 111D, 111U and 201. In contrast, the revision networks from non-controversial articles exhibit two anti-motifs, involving the under representation of triads involving two directed edges either arriving at or emanating from a mediating node (triads 021D and 021U). These motifs and the underlying subgraph ratio profiles represent an unusual and distinctive profile that we believe represent distinctive "super-families" beyond those seen in other technologically related networks, such as the world wide web [30].

Performing dimensionality reduction upon the subgraph ratio profiles from each revision network allows us to further understand the relationships between Wikipedia articles. Our analysis shows that the structure of the data is amenable to reduction to two dimensions, where the principal components are dominated by triads 111D and 111U in the first component, and mainly 021C in the second component, but with lesser contributions from triads 021D and 021U.

The results from two-dimensional principal analysis are illuminating - the dominant triads in both components, as listed above, correspond to linear paths, i.e., open triads

which represent sequences of editing without indirect reciprocity. The extent and format of reciprocated (i.e., bidirectional) edges on these open triads is sufficient to define the two principal components. The dominant triads in the first principal component each involve reciprocation on one edge, where as interestingly, in the second principal component, the dominant triads are open triads with no reciprocated edges. From this we deduce that short paths, rather than short loops of editing that represent indirect reciprocity, are important features in characterising Wikipedia revision networks.

We also observe through two-dimensional principal component analysis that it is the first principal component that strongly distinguishes between the revision networks of controversial and non-controversial articles. The dominant triads defining this capture the extent of direct reciprocation being present. Finally, through consideration of the principal components against additional external variables, we find in particular that article age plays a role in distinguishing the controversial articles. High values in both principal components aligns with strongest clustering of controversial articles, which is not the case for non-controversial articles.

6 CONCLUSION

The analysis has given insights into the structure underlying revision networks from Wikipedia articles, and has shown that a relatively small number of features, in terms of substructures in revision networks, characterise controversial Wikipedia articles. The results have identified key clusters of editorial interactions to this effect, in support of the hypothesis. These are distinctive and indicate that the revision networks for controversial and non-controversial Wikipedia articles have differentiated subgraph ratio profiles. Our study gives understanding as to how prediction or classification of articles can be enhanced using the latent structures relating to editor behaviour. This also reaffirms the importance of the revision network as a simple but useful representation for assessment of Wikipedia articles.

ACKNOWLEDGEMENTS

This research was sponsored by the U.S. Army Research Laboratory and the U.K. Ministry of Defence under Agreement Number W911NF-16-3-0001. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the U.S. Army Research Laboratory, the U.S. Government, the U.K. Ministry of Defence or the U.K. Government. The U.S. and U.K. Governments are authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation hereon.

REFERENCES

- [1] Daron Acemoglu, Asuman Ozdaglar, and Ali ParandehGheibi. 2010. Spread of (mis)information in social networks. *Games and Economic Behavior* 70, 2 (2010), 194–227.
- [2] B Thomas Adler, Luca De Alfaro, Santiago M Mola-Velasco, Paolo Rosso, and Andrew G West. 2011. Wikipedia vandalism detection: Combining natural language, metadata, and reputation features. In *International Conference on Intelligent Text Processing and Computational Linguistics*. Springer, 277–288.
- [3] Vladimir Batagelj and Andrej Mrvar. 2001. A subquadratic triad census algorithm for large sparse networks with small maximum degree. *Social Networks* 23, 3 (2001), 237 – 243. [https://doi.org/10.1016/S0378-8733\(01\)00035-1](https://doi.org/10.1016/S0378-8733(01)00035-1)
- [4] Daren C Brabham. 2008. Crowdsourcing as a Model for Problem Solving: An Introduction and Cases. *Convergence* 14, 1 (2008), 75–90.
- [5] Dave Braines, Diane Felmlee, Don Towsley, Kun Tu, Roger M Whitaker, and Liam D Turner. 2018. The role of motifs in understanding behavior in social and engineered networks. In *Next-Generation Analyst VI*, Vol. 10653. International Society for Optics and Photonics, 106530W.
- [6] Ulrik Brandes, Patrick Kenis, Jürgen Lerner, and Denise Van Raaij. 2009. Network analysis of collaboration structure in Wikipedia. In *Proceedings of the 18th international conference on World wide web*. ACM, 731–740.
- [7] Anhai Doan, Raghu Ramakrishnan, and Alon Y. Halevy. 2011. Crowdsourcing Systems on the World-Wide Web. *Commun. ACM* 54, 4 (April 2011), 86–96.
- [8] Takashi Iba, Keiichi Nemoto, Bernd Peters, and Peter A Gloor. 2010. Analyzing the creative editing behavior of Wikipedia editors: Through dynamic social network analysis. *Procedia-Social and Behavioral Sciences* 2, 4 (2010), 6441–6456.
- [9] David Jurgens and Tsai-Ching Lu. 2012. Temporal Motifs Reveal the Dynamics of Editor Interactions in Wikipedia.
- [10] Gerald C Kane and Sam Ransbotham. 2016. Research note: Content and collaboration: an affiliation network approach to information quality in online peer production communities. *Information Systems Research* 27, 2 (2016), 424–439.
- [11] Brian Keegan, Darren Gergle, and Noshir Contractor. 2012. Staying in the loop: structure and dynamics of Wikipedia’s breaking news collaborations. In *Proceedings of the Eighth Annual International Symposium on Wikis and Open Collaboration*. ACM, 1.
- [12] Aniket Kittur, Bongwon Suh, Bryan A Pendleton, and Ed H Chi. 2007. He says, she says: conflict and coordination in Wikipedia. In *Proceedings of the SIGCHI conference on Human factors in computing systems*. ACM, 453–462.
- [13] Srijan Kumar, Robert West, and Jure Leskovec. 2016. Disinformation on the web: Impact, characteristics, and detection of wikipedia hoaxes. In *Proceedings of the 25th international conference on World Wide Web*. International World Wide Web Conferences Steering Committee, 591–602.
- [14] David Laniado, Riccardo Tasso, Yana Volkovich, and Andreas Kaltenbrunner. 2011. When the wikipedians talk: Network and tree structure of wikipedia discussion pages.. In *Sixth International AAAI Conference on Weblogs and Social Media*. 177–184.
- [15] Ron Milo, Shalev Itzkovitz, Nadav Kashtan, Reuven Levitt, Shai Shen-Orr, Inbal Ayzenshtat, Michal Sheffer, and Uri Alon. 2004. Superfamilies of evolved and designed networks. *Science* 303, 5663 (2004), 1538–1542.
- [16] Ron Milo, Shai Shen-Orr, Shalev Itzkovitz, Nadav Kashtan, Dmitri Chklovskii, and Uri Alon. 2002. Network motifs: simple building blocks of complex networks. *Science* 298, 5594 (2002), 824–827.
- [17] Delia Mocanu, Luca Rossi, Qian Zhang, Marton Karsai, and Walter Quattrociocchi. 2015. Collective attention in the age of (mis)information. *Computers in Human Behavior* 51 (2015), 1198–1204.
- [18] Derek O’Callaghan, Martin Harrigan, Joe Carthy, and Pádraig Cunningham. 2012. Network Analysis of Recurring YouTube Spam Campaigns. In *Sixth International AAAI Conference on Weblogs and Social Media*.
- [19] Martin Potthast, Benno Stein, and Robert Gerling. 2008. Automatic vandalism detection in Wikipedia. In *European conference on information retrieval*. Springer, 663–668.
- [20] Hoda Sepehri Rad and Denilson Barbosa. 2012. Identifying Controversial Articles in Wikipedia: A Comparative Study. In *Proceedings of the Eighth Annual International Symposium on Wikis and Open Collaboration (WikiSym ’12)*. ACM, New York, NY, USA, Article 7, 10 pages.
- [21] Hoda Sepehri Rad and Denilson Barbosa. 2012. Identifying controversial articles in Wikipedia: A comparative study. In *Proceedings of the eighth annual international symposium on wikis and open collaboration*. ACM, 7.
- [22] Hoda Sepehri-Rad and Denilson Barbosa. 2015. Identifying Controversial Wikipedia Articles Using Editor Collaboration Networks. *ACM Trans. Intell. Syst. Technol.* 6, 1, Article 5 (March 2015), 24 pages.
- [23] Hoda Sepehri-Rad and Denilson Barbosa. 2015. Identifying controversial Wikipedia articles using editor collaboration networks. *ACM Transactions on Intelligent Systems and Technology (TIST)* 6, 1 (2015), 5.
- [24] James Surowiecki and James. 2005. *The wisdom of crowds*. Anchor Books. 306 pages. <https://dl.acm.org/citation.cfm?id=1095645>
- [25] Kun Tu, Jian Li, Don Towsley, Dave Braines, and Liam D. Turner. 2018. Network classification in temporal networks using motifs. In *Proceedings of the 3rd ECML/PKDD Workshop on Advanced Analytics and Learning on Temporal Data*.
- [26] M Del Vicario, A Bessi, F Zollo, F Petroni, A Scala, and G Caldarelli. 2016. The spreading of misinformation online. *Proc National Academy of Sciences* 113, 3 (2016).
- [27] Guangyu Wu and Pádraig Cunningham. 2015. Integration of multiple network views in Wikipedia. *Knowledge and Information Systems* 45, 2 (2015), 473–490.
- [28] Guangyu Wu, Martin Harrigan, and Pádraig Cunningham. 2011. Characterizing Wikipedia pages using edit network motif profiles. In *Proceedings of the 3rd international workshop on Search and mining user-generated contents*. ACM, 45–52.
- [29] Guangyu Wu, Martin Harrigan, and Pádraig Cunningham. 2012. Classifying Wikipedia articles using network motif counts and ratios. In *Proceedings of the Eighth Annual International Symposium on Wikis and Open Collaboration - WikiSym ’12*.
- [30] Vinko Zlatić, Miran Božičević, Hrvoje Štefančić, and Mladen Domazet. 2006. Wikipedias: Collaborative web-based encyclopedias as complex networks. *Physical Review E* 74, 1 (2006), 016115.