Mind Your POV: Convergence of Articles and Editors Towards Wikipedia's Neutrality Norm

UMASHANTHI PAVALANATHAN, Georgia Institute of Technology, USA

XIAOCHUANG HAN, Georgia Institute of Technology, USA

JACOB EISENSTEIN, Georgia Institute of Technology, USA

Wikipedia has a strong norm of writing in a “neutral point of view” (NPOV). Articles that violate this norm are tagged, and editors are encouraged to make corrections. But the impact of this tagging system has not been quantitatively measured. Does NPOV tagging help articles to converge to the desired style? Do NPOV corrections encourage editors to adopt this style? We study these questions using a corpus of NPOV-tagged articles and a set of lexicons associated with biased language. An interrupted time series analysis shows that after an article is tagged for NPOV, there is a significant decrease in biased language in the article, as measured by several lexicons. However, for individual editors, NPOV corrections and talk page discussions yield no significant change in the usage of words in most of these lexicons, including Wikipedia's own list of “words to watch.” This suggests that NPOV tagging and discussion does improve content, but has less success enculturating editors to the site's linguistic norms.

CCS Concepts: • **Human-centered computing** → *Empirical studies in collaborative and social computing*;

Additional Key Words and Phrases: Wikipedia, online communities, community norms, norm enforcement, objective language, writing style.

ACM Reference Format:

Umashanthi Pavalanathan, Xiaochuang Han, and Jacob Eisenstein. 2018. Mind Your POV: Convergence of Articles and Editors Towards Wikipedia's Neutrality Norm. *Proc. ACM Hum.-Comput. Interact.* 2, CSCW, Article 137 (November 2018), 23 pages. <https://doi.org/10.1145/3274406>

1 INTRODUCTION

Wikipedia is the largest online collaborative content creation community. The English Wikipedia contains nearly 5.6 million articles and nearly 140 thousand actively contributing accounts [78]. Similar to the norms or rules of other online communities [12], Wikipedia also has a set of norms regarding content and member participation [80]. The three core content policies of Wikipedia are: “neutral point of view (NPOV)”, “verifiability”, and “no original research” [77]. Wikipedia provides strict guidelines about NPOV, and non-complying articles are marked with an NPOV tag by editors. Such NPOV-tagged articles are listed under the “NPOV disputed” Wikipedia category. While the NPOV tagging system has been in place for over a decade, its impact on Wikipedia has not been quantitatively measured. In this paper, we study the effectiveness of NPOV tagging as a language norm enforcement strategy in Wikipedia, measuring its impact on articles as well as individual editors. Specifically, we ask whether tagging helps articles and editors converge towards

Authors' addresses: Umashanthi Pavalanathan, Georgia Institute of Technology, School of Interactive Computing, USA; Xiaochuang Han, Georgia Institute of Technology, College of Computing, USA; Jacob Eisenstein, Georgia Institute of Technology, School of Interactive Computing, USA.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2018 Copyright held by the owner/author(s). Publication rights licensed to ACM.

2573-0142/2018/11-ART137 \$15.00

<https://doi.org/10.1145/3274406>

the community's prescribed writing style, as a step toward the larger goal of understanding how collaborative communities converge on language standards.

1.1 Norm Enforcement in Wikipedia

Wikipedia provides detailed guidelines of the preferred writing styles through a manual of style [81]. The manual of style contains guidelines not only for formatting such as capitalization, punctuations, and abbreviations, but also for grammar and vocabulary. The *Neutral Point of View (NPOV)* guideline states that “*All encyclopedic content on Wikipedia must be written from a neutral point of view (NPOV), which means representing fairly, proportionately, and, as far as possible, without editorial bias, all of the significant views that have been published by reliable sources on a topic.*” The guideline further states that “*This policy is non-negotiable, and the principles upon which it is based cannot be superseded by other policies or guidelines, nor by editor consensus.*” These editorial policies are enforced by the community of editors to maintain high quality of the articles.

Fig. 1. An NPOV tag displayed at the top of an article.¹

Wikipedia articles which do not follow the NPOV guidelines are marked by editors with tags such as {{POV}}, {{NPOV}}, and {{POV-section}} so that they can be included in the editing workflow (Figure 1). From the explicit tagging of the non-NPOV content in Wikipedia articles and the availability of the content before and after NPOV revisions, we can develop a better understanding of the acceptable writing style in Wikipedia. The example below from our dataset shows a portion of a sentence before and after it was corrected for NPOV. Usage of **dictated** in (a) suggests a point of view and it is replaced with **overseen** in (b) to make the language more neutral.

- (a) “BL Republic is **dictated** by UberConsul Lars and High Marshall Bill ...”
- (b) “BL Republic is **overseen** by UberConsul Lars and High Marshall Bill ...”

1.2 Research Questions and Findings

We study the effects of NPOV tagging and revision as norm enforcement strategies to make Wikipedia an objective source of knowledge at the article level and at the editor level.

RQ1: What are the article-level effects of NPOV tagging?

RQ2: What are the editor-level effects of NPOV correction?

RQ2a: What effect does NPOV correction have on the writing style of the editors?

RQ2b: What effect does NPOV correction have on the engagement of editors?

We address these questions using a corpus of nearly 7,500 articles in the *NPOV dispute* category, and we identify revisions in which a NPOV tag was added or removed, or in which content was corrected for NPOV. We then identify the first revision in which an article received a NPOV tag as the treatment point for the article. Similarly we identify editors whose content contributions were corrected for NPOV as treatment editors and identify the treatment revisions. As a measure of language, we use a set of lexicons associated with non-objective or biased language.² Considering

¹Image obtained from https://en.wikipedia.org/wiki/Natural_Gas_Choice on 30-August-2018.

²Lexical analysis is chosen because of the existence of well-validated lexicons, as described later. We leave for future work the analysis of bias in dimensions such as syntax, semantics, and pragmatics.

editor contribution and article revisions before and after the treatment points, we analyze the effect of treatment using an interrupted time series analysis [8].

We find that the NPOV tagging helps articles to converge to the desired writing style of Wikipedia. However, for individual editors, NPOV corrections yield no statistically significant change in most lexicons used to characterize biased language, including Wikipedia's own "words to watch" lexicon. This finding holds even when NPOV corrections are paired with one-to-one interactions in the form of talk page discussions, which trigger email prompts to the corrected editor. To summarize our results: NPOV norm enforcement helps to make the articles better, but does not lead individual editors to converge to the writing style norms of Wikipedia.

1.3 Terminology

Below, we provide a brief summary of different terminology used in this paper:

Edit refers to any changes made to a Wikipedia article.

Revision is a version of a Wikipedia article resulted from editing the previous revision.

Article contains chronologically ordered revisions.

Reverts in Wikipedia refers to edits that restored the current revision to a previous revision.

NPOV tagging refers to the addition of any NPOV templates to the articles.

NPOV correction refers to a set of edits that are made to correct for NPOV errors. These are detected from the comment metadata in a revision.

Treatment refers to an intervention made to the subject of interest (i.e., article, editor).

Pre-treatment refers to the period prior to the treatment.

Post-treatment refers to the period after the treatment.

Treatment to the articles refers to the addition of an NPOV tag.

Treatment to the editors refers to editors' contribution being corrected for NPOV.

Biased language refers to words/phrases which are not neutral and may introduce a point of view or attitude.

2 BACKGROUND AND RELATED WORK

Prior work falls into following main areas: online community norms, writing style norms, norm enforcement, effects of norm enforcement, and motivation for participation and editor roles in Wikipedia.

2.1 Online Community Norms

Norms are habitual behaviors that characterize a social group and differentiate it from other social groups [41]. While norms are explicitly defined through detailed guidelines (e.g., Reddit) and FAQs (e.g., Usenet) in some communities, in other communities norms are not formally codified, but emerge socially through the interactions of the members [12]. Online communities exhibit norms [14] related to member behavior such as content appropriateness [7], writing style [81], and adherence to community policies [52]. Community norms evolve as its members negotiate norms [49] and as newcomers arrive and old members become less active [21]. While community norms are studied in terms of deviant behavior such as abusive language [73], trolling [17], and vandalism [34], norms also encourage participation [68] and help communities achieve their goals [11, 15]. As a collaborative content production community, Wikipedia has several norms [80] related to neutral point of view [77], supportive communication [66], vandalism [63], and member participation [11]. In this work, we focus on Wikipedia's norms related to NPOV writing style.

2.2 Writing Style Norms

Communicative competence is the ability to use language appropriately, in accord with social situation and associated norms [44]. This notion of communicative competence can be used to explain linguistic norms of communities—the language used by community members should not only be correct, but should also be appropriate. A related concept is *legitimate language*, which is defined in any community as the language produced by a subset of speakers/writers with symbolic authority, and is often codified into explicit standards [10]. In the context of collaborative online writing, legitimate language can be considered to be the language produced by the high-status members of the community, and is governed by community policies.

Among online communities, Wikipedia is well known for its focus on the quality of the content produced by volunteer contributors [11]. Wikipedia provides detailed guidelines of the preferred writing styles through a manual of style [81]. A large body of research focused on Wikipedia's writing style norms including the prediction of article quality [55], biased content [1], quality flaws [2, 3], and vandalism detection [39]. Focusing specifically on bias in Wikipedia, Recasens et al. [67] built a dataset of Wikipedia phrases that are corrected for NPOV, and then trained a classifier using linguistic features from hand-crafted lexicons of factive verbs, hedges, and subjective intensifiers. We build on this prior work by using these lexicons to characterize biased language, but note that while Recasens et al. [67] aimed to identify the linguistic aspects of norms, our focus is to understand the effects of NPOV tagging and correction. Specifically, we are interested in the change in the rate of biased language use when an article has an active NPOV tag, and when an editor is corrected for NPOV.

2.3 Enforcing Norms

Deviant behavior is an ongoing challenge to online communities, requiring persistent regulatory effort [48]. In addition to providing detailed guidelines, online communities also take active moderation actions against the violations of community guidelines [26], including both technical and social approaches [36]. Technical moderation includes banning posts with any keywords from predefined word lists (e.g., Yik Yak [25]) and banning posts based on the IP addresses from which posts originate (e.g., Hacker News [59], Yelp [65]). Social moderation includes both distributed approaches [53] where the community members vote on content (e.g., Yik Yak, Slashdot) [52] and centralized approaches where a small number of users called moderators maintain community practices by removing posts they find inappropriate (e.g., Reddit) [36].

In Wikipedia, a rich set of policies and guidelines articulate strategies for seeking consensus, principles of encyclopedic content, and appropriate user behavior. The policy environment in Wikipedia encodes and explains norms, but the policies are not imposed top down; rather policies are created and managed by the Wikipedia community itself [9]. Norms are enforced by policy citations [9], where template messages [24] are added to articles and discussion talk pages [45]. The policy pages are subject to the same editing processes of Wikipedia articles and contributors are given a participatory role when creating or editing policy pages. While several quality control mechanisms to enforce objective language in Wikipedia articles have been in place for a long time, the impact of these mechanisms has not been quantitatively measured. This gap in the literature motivates our broad research question in this study: whether NPOV tagging and corrections help Wikipedia articles and editors to converge to desired writing styles norms.

2.4 Effects of Norm Enforcement

Most prior work on the effects of norm enforcement has focused on moderation techniques related to the design of online communities. Users report significantly higher intent to participate in a

community that is moderated, compared to an unmoderated community [82]. A study on Slashdot – a popular website with a good amount of moderation – found that users come to a consensus about the community’s moderation policies and such moderation practices enable large scale civil participation [53]. Reddit’s decision to ban hate communities resulted in a reduction of abusive language [16]. An analysis of different moderation styles in online health support communities found that positive and rewarding moderation styles are more effective than negative and punishing styles [58]. In online discussion forums related to education, peer moderation is found to be more encouraging active participation than moderation by superiors [71].

The effects of moderation on the quality of collaboratively created content and long term behavior of individual community members have not been investigated much. In the context of Wikipedia, Halfaker et al. [38] found that the action of “reverting” edits has the effect of reducing motivation and quantity of work, particularly for new editors. However, they also found that reverts result in higher quality contributions. Halfaker et al. [37] found that enforcing quality control mechanisms affects the retention of high-quality newcomers. Personalized warning messages, as opposed to pre-defined template messages, are found to be effective in retaining newcomers [33]. Motivated by this prior literature, which studies the effectiveness of norm enforcement actions both at the platform-level (i.e., articles) and individual member-level (i.e., editors), we seek to answer our research question about the effectiveness of NPOV norm enforcement both at the article level (RQ1) and at the editor level (RQ2).

2.5 Motivation for Participation and Editor Roles in Wikipedia

The scale and success of Wikipedia as a volunteer-run collaborative content producing community have attracted a large body of research about various aspects of Wikipedia including motivations to contribute [32, 60, 70], formation of community [6, 19, 74], content quality [18, 54], editorial and authorship attributes [28], sharing and collaborative knowledge building [54], learning [32], coordinations of complex tasks [47], and editor functional roles [4, 31]. Prior studies on the motivations for voluntary contribution to Wikipedia provide several explanations including the ideology for contributing to Wikipedia as a variant of open-source application [60], internal self-concept motivation [84], cognitive (e.g., learning new things or intellectual challenge) and affective reasons (e.g., pleasure) [64], and the sense of individual efficacy [11].

Another line of work focused on different editor roles that emerge in Wikipedia [4, 11, 51, 61, 76]. Bryant et al. [11] find that while novice editors contribute edits to articles related to their domain of expertise, expert editors contribute towards improving the quality of Wikipedia itself. Arazy et al. [4] identified functional roles in Wikipedia such as technical administrator, border patrol, quality assurance technicians, administrators, and directors, and find that editor roles determine activity patterns across variety of editing tasks. Yang et al. [83] automatically identify various editor roles based on low-level edit types and find that different editor roles contribute differently in terms of edit categories and articles in different quality stages require different types of editors. In this work, we focus on editor roles that make contributions to improve the quality of Wikipedia articles, particularly articles in the NPOV dispute category.

3 DATASETS

In this section we describe the datasets and extraction of treatment articles and editors. Figure 2 summarizes our data pipeline and methods.

3.1 Wikipedia NPOV Corpus

We begin with the NPOV corpus from Recasens et al. [67], which they used to build a classifier to detect bias inducing terms in phrases. The NPOV corpus contains 7,464 articles from the English

Fig. 2. Flowchart showing the data and method setup for each research question.

Wikipedia with a total of 2.7 million revisions. It was constructed by retrieving all the articles that were in the *NPOV dispute* category³ in early 2013, together with their full revision history. Wikipedia editors are encouraged to identify and revise biased content to achieve neutral tone, and several NPOV tags are used to mark biased content.⁴ Articles that are marked with any of the NPOV tags will be listed in Wikipedia's category of *NPOV disputes*. Each version of the article is considered as a *revision*, and an article is a set of chronologically ordered revisions. Each revision of the article contains the version of the article content along with metadata such as revision ID, timestamp, contributor, revision comment, and an SHA-1 hash key.

3.2 Identifying NPOV Tagging and Removal

Several NPOV tags are used to mark biased content and there is no standardized template or metadata available to directly detect tag addition and removal. Therefore, to identify in which revision an article was tagged as violating NPOV and when that tag was removed, we used a set of NPOV-tag patterns based on a preliminary inspection of a small subset of NPOV tagged articles.⁵ In order to detect the addition or removal of an NPOV tag in the absence of any readily available metadata, we need to compare each consecutive revisions and then inspect the difference in the textual content for an addition or removal of an NPOV tag. To automatically extract the changes between two consecutive revisions in terms of added and removed content, we used the Diff Match and Patch library,⁶ which was also used by Recasens et al. [67]. If any of the NPOV tags was present in the added content, then we mark that revisions as an NPOV tag addition point for that article. Similarly, if any of the NPOV tags was present in the removed content, we mark that revision as an NPOV tag removal point for that article.

Accurate identification of valid tag additions and removal is challenging due to the presence of vandalism and reverts.⁷ From our initial inspections of a small set of NPOV tagged articles, we

³https://en.wikipedia.org/wiki/Category:All_NPOV_disputes

⁴e.g., `{{POV}}`, `{{POV-check}}`, `{{POV-section}}`, etc. When such tag is added, a template such as the following is displayed in the article page: "The neutrality of this article is disputed. Relevant discussion may be found on the talk page. Please do not remove this message until the dispute is resolved."

⁵`{{{.*(POV|Pov|PoV|Point Of View|pov|NPOV|Npov|npov|NEUTRALITY|Neutrality|neutrality|NEUTRAL|Neutral|neutral).*}}}`.

⁶<http://code.google.com/p/google-diff-match-patch>

⁷Vandalism and edit wars are primary challenges for Wikipedia editors. There is a large body of work regarding this. For example, see Geiger and Ribes [34], Potthast et al. [62], Shachaf and Hara [72].

identified several cases of vandalism (e.g., NPOV tags were removed intentionally or accidentally, and then added back within the next few revisions; revisions containing NPOV tags were reverted to a prior revision without a tag). To minimize the effect of vandalism in the accurate identification of NPOV tag addition and removal, we used two heuristics. First, we used the SHA-1 hash key present in the revision metadata to track reverts and disregarded reverted revisions. Second, if an NPOV tag was added immediately after a prior tag removal, we consider the last tag removal as unreliable. If such changes appear within five revisions after a tag removal, we merge the consecutive tag addition-removal pairs and consider only the latest revision as a valid tag removal revision. In this way, we identified 6,512 articles that had at least one NPOV tag addition and 1,095 articles with multiple NPOV tag additions.

3.3 Identifying Treatment Editors

For our editor-level analysis, we define treatment as editing part(s) of an article's content to correct for NPOV. Treatment editors are editors who originally contributed the portion of the text that was corrected for NPOV. To identify whether a revision was corrected for NPOV, we first check if the comment metadata for that revision contains "NPOV" or "POV" or any case variations. If one of these strings appear in the comment and if that revision was not identified as an NPOV tag addition or removal revision, then we consider it to be a *NPOV correcting* revision. While Recasens et al. [67] checked for only the presence of NPOV strings in the comment to detect NPOV corrections, in our initial inspection we found that sometimes comments of NPOV tagging and removal revisions also contain these NPOV strings. Therefore, we additionally checked if the revision had a tag addition or removal as we do not consider those as NPOV correcting revisions.⁸

Once we have identified NPOV correcting revisions, we trace back previous revisions to identify the treatment editors—editors who originally contributed the portion of text that was corrected for NPOV. As a preprocessing step, we removed revisions which are reverts of previous revisions in order to identify the editor who originally contributed a span of text. Otherwise, the editor who made the revert could have been attributed to a portion of text. To trace the treatment editors, we first built a detailed record of which part of the article was contributed by which editor at the character granularity level from the beginning of the article revision history. We then checked all of the NPOV correcting revisions, extracted their first *diff* compared to the previous revision, and looked into the record to identify which editor originally contributed to that *diff-ed* language.⁹

Automated scripts or bots are used in Wikipedia to perform repetitive and mundane tasks such as minor fix for syntax and adding dates to NPOV tags. Because our interest is in studying the writing style of human editors, we removed bots by checking for the presence of "bot" in the username and presence for the "bot" editor group attributes.¹⁰ In total, we identified 5,645 treatment editors after removing 20 bot accounts. Note that each treatment editor could be NPOV corrected more than once. On average, a treatment editor is corrected for NPOV 1.9 times, and 32.4% of the treatment editors are corrected for NPOV more than once.

⁸Note that an NPOV removal revision could include corrections for NPOV and removal of NPOV tag. However, in our initial inspections, we found that NPOV issues are often solved in revisions prior to the NPOV tag removal revision. In this way, few NPOV correcting instances could have been missed, but we aimed for precision rather than recall.

⁹Note that recent work such as Flöck and Acosta [30] proposed other sophisticated algorithms to attribute authorship of revisioned content. Future work could consider employing such algorithms.

¹⁰Editor attributes are obtained from the Media Wiki API <https://www.mediawiki.org/wiki/API:Query>. Examples of editor group attributes include "reviewer", "administrator", and "bot".

3.4 Dataset for Article-Level Analysis

In RQ1 we study the effect of NPOV tagging on the linguistic style of the articles. For this analysis we use the revisions of articles in the NPOV corpus (§ 3.1) and characterize the linguistic style of each of the revisions. To observe the biased language trend over time, we limited this dataset to include NPOV articles which have at least 40 revisions after the treatment (i.e., addition of an NPOV tag) and considered up to 40 revisions before and after the treatment for the regression analysis.¹¹

3.5 Dataset for Editor-Level Analysis

To study the editor-level effects of NPOV correction, we collected each treatment editor's textual contribution to the Wikipedia article namespace¹² by querying the Media Wiki API.¹³ In RQ2a we study the effect of NPOV correction on the writing style of the treatment editors. To characterize treatment user writing style before and after treatment, we first obtained up to 150 revisions to which the editors contribute (in all Wikipedia articles, not restricting to the initial NPOV corpus) before and after their first treatment point. To control for any platform level changes such as revisions to NPOV policies, we restricted the pre and post treatment revisions to be within an year from the treatment. We then queried for the parent revision of each of the editors' revisions and extracted the text added by the treatment editors using the *diff* library. Since our interest is in the language use, we considered only the revisions with any content addition.¹⁴ In RQ2b we study the changes in treatment editor engagement before and after NPOV correction. To characterize treatment editor engagement we obtained all the revisions (in all Wikipedia articles, not restricting to the initial NPOV corpus) they have contributed to during two months before and after the treatment point.

Attributing Editor Content Contribution. Our approach to identify the editor who originally contributed NPOV corrected content uses the *diff* of two consecutive revisions as described in § 3.3. One potential issue with this approach is content mis-attribution due to vandalism [23]. To minimize this issue, we identify and remove reverted revisions. When characterizing editor language (§ 3.5), we compare each of the editor's revision to its parent revision and compute the *diff* of both revisions. When making revisions, editors could re-introduce content that was previously contributed by another editor. Even though the re-introduced content may not be originally contributed by the editor who makes the revision, we assume that the editor is at least partially responsible for the language used in that content as they include it in the revision they make. While our approach is simple in terms of implementation, other sophisticated approaches (e.g., Flöck and Acosta [30]) exist and such approaches are specifically designed for the task of authorship attribution for revisioned content. We acknowledge that such methods are superior to our approach, and future work could use an implementation of such algorithms.

¹¹We also considered the number of revision threshold window W of $[10 \leq W < 20]$ and $[20 \leq W < 40]$ and the results are similar qualitatively. The number of articles included in the analysis for each of these thresholds are 1054, 1095, and 1763 for the revision windows $[10 \leq W < 20]$, $[20 \leq W < 40]$, and $[W \geq 40]$ respectively.

¹²Wikipedia includes multiple namespaces such as article page, article talk page, user page, user talk page, etc. Since our interest is to characterize editor writing style in Wikipedia articles, we restricted our dataset to the article namespace. Prior work, such as Elia [27], has found significant differences between the linguistic style of article text and talk page text due to the conversational nature of talk pages.

¹³<https://www.mediawiki.org/wiki/API:Query>

¹⁴Examples of other revision contributions include addition of links, templates, and citations.

4 METHODS

In this section we describe methods related to preprocessing of Wikipedia data for our analyses, identifying biased language, and measuring effect of treatment.

4.1 Identifying Biased Language

Bias in Wikipedia can emerge from several factors including language style, editors' point of view, cited sources, coverage of topics, etc. Prior work focused on political bias [35], cultural bias [13], gender bias [75], topic bias [29], and authoritative bias [22]. Our goal in RQ1 and RQ2a is to study the effects of NPOV tagging in article-level and editor-level language. Particularly, we are interested in characterizing the bias in articles and editor contribution in terms of the linguistic style that may introduce bias. The Wikipedia Manual of Style [81] states that *"There are no forbidden words or expressions on Wikipedia, but certain expressions should be used with caution, because they may introduce bias. Strive to eliminate expressions that are flattering, disparaging, vague, or endorsing of a particular view point."* Prior work [e.g., 40, 67] has addressed the task of identifying biased language in Wikipedia at different levels. However, we are not aware of a reliable system that can detect bias in terms of linguistic style in Wikipedia articles. Therefore, we use a set of linguistic style lexicons as a proxy to characterize biased language in Wikipedia.¹⁵

The Wikipedia Manual of Style provides a list of *words to watch* in a prescriptive manner [79], indicating style words that may introduce bias. We compiled these styles words into a lexicon called *Words to Watch*. Prior work of Recasens et al. [67] introduced the task of detecting bias inducing terms in phrases from Wikipedia articles and used a set of pre-compiled style lexicons that are indicative of expressions of attitude or point of view. These lexicons include hedges, factive verbs, assertive verbs, positive words, and negative words. We use these nine lexicons in addition to the words from Wikipedia manual.¹⁶ A list of all the lexicons we used is shown in Table 1 with a description, sources, and example terms. To characterize the amount of biased language in a text, we compute the coverage of each lexicon words per token in the text.

4.2 Measuring the Effect of Treatment on Biased Language Usage

Our goal in RQ1 is to investigate whether NPOV tag addition has an effect on the level of biased language in the NPOV tagged articles. Similarly, our goal in RQ2a is to investigate whether NPOV correction has an effect on the level of biased language usage of treatment users. Both of these are causal questions. To answer them, we apply the Interrupted Time Series technique for causal inference [8]. This method has been used in recent studies on hate speech [16] and conspiratorial discussions[57]. Interrupted Time Series analysis is a quasi-experimental design that can be used to evaluate the longitudinal effects of an intervention (i.e., treatment), through segmented regression modeling. The term quasi-experimental refers to an absence of randomization. ITS is a tool for analyzing observational data where complete randomization, or case-control design, is not affordable or possible. A notable strength of ITS with respect to measuring treatment effects using observational data is that ITS controls for the effects of secular trends in a time series of outcome measure. For example, if a measure is increasing or decreasing historically over time, and if we only look at the averages pre- and post- treatment, we would not notice the change in slope or level.

¹⁵We limit our analysis only to lexicon, while syntax, semantics, and pragmatic measures could also be studied.

¹⁶Note that Recasens et al. [67] used a bias lexicon created from the NPOV articles. We do not include it for our analysis as this lexicon is not entirely indicative of linguistic style and included content words such as *communist*, *historian*, and *migration*.

Table 1. List of Lexicons Used to Characterize Bias Language.

Lexicon Name	Description	Example Terms
Words to Watch	Style words that may introduce bias (Wikipedia [79])	<i>fortunately, notable, often, speculate</i>
Hedges	Used to reduce one's commitment to the truth of a proposition (Hyland [43])	<i>apparent, seems, unclear, would</i>
Assertives	Complement clauses that assert a proposition (Hooper [42])	<i>allege, hypothesize, verify, claim</i>
Positive Words	Positive sentiment terms (Liu et al. [56])	<i>achieve, inspire, joyful, super</i>
Negative Words	Negative sentiment terms (Liu et al. [56])	<i>criticize, foolish, hectic, weak</i>
Factives	Terms that presuppose the truth of their complement clause (Kiparsky and Kiparsky [50])	<i>regret, amuse, strange, odd</i>
Implicatives	Imply the truth or untruth of their complement, depending on the polarity of the main predicate (Karttunen [46])	<i>avoid, hesitate, refrain, attempt</i>
Report Verbs	Used to indicate that discourse is being quoted or paraphrased (Recasens et al. [67])	<i>praise, claim, dispute, feel</i>
Strong Subjectives	Add strong subjective force to the meaning of a phrase (Riloff and Wiebe [69])	<i>celebrate, dishonor, overkill, worsen</i>
Weak Subjectives	Add weak subjective force to the meaning of a phrase (Riloff and Wiebe [69])	<i>widely, unstable, although, innocently</i>

In its basic form, an ITS is modeled using a regression model (e.g., linear, logistic or Poisson) that includes only three time-based covariates as shown in Equation 1. The regression coefficients of these covariates estimate the pre-treatment trend, the level change at the treatment point, and the trend change from pre-treatment to post-treatment.

$$Y_t = \beta_0 + \beta_1 T + \beta_2 X_t + \beta_3 TX_t \quad (1)$$

Here T is the time elapsed since the start of the study; X_t is a dummy variable indicating the pre-treatment period (coded 0) or the post-treatment period (coded 1); Y_t is the outcome at time t . The pre-treatment slope, β_1 , quantifies the trend for the outcome before the treatment; the level change, β_2 , estimates the change in level that can be attributed to the treatment; the change in slope, β_3 , quantifies the difference between the pre-treatment and post-treatment slopes. This model of ITS assumes that without the treatment, the pre-treatment trend would continue unchanged into the post-treatment period, and there are no external factors systematically affecting the trends. One important step before modeling an ITS analysis is to hypothesize how the treatment would impact the outcome if it were effective, particularly whether the change will be in the slope of the trend, a change in the level, or both.

While we assume the events of NPOV tagging and subsequent NPOV corrections to be independent, editors could tag an article for NPOV as a preparation for their subsequent revisions. In such cases, NPOV tagging could be a precursor to future revisions by the same editor. Therefore, in

our analysis for RQ1, we remove all the articles where the editor who tagged an article for NPOV contributed a revision in the subsequent 40 revisions after the NPOV tagging revision.¹⁷

5 RQ1: ARTICLE-LEVEL EFFECTS OF NPOV TAGGING

5.1 Article Writing Style: RQ1

5.1.1 Quantifying Article Writing Style. We use the dataset we created for the article level analysis (§ 3.4) to study RQ1. Before measuring the writing style of each article revision, we performed a set of preprocessing steps to extract textual contribution of the article revision. We removed Wikipedia markup templates, URLs, and other non-textual content. We then tokenized the text and extracted tokens with two or more characters. To quantify article writing style, we extracted the coverage of tokens from each of the ten bias-related lexicons (§ 4.1) and computed lexicon coverage rate as the number of lexicon tokens per total tokens in the text of each article revision. We considered the latest 40 pre-treatment revisions and the earliest 40 post-treatment revisions of the NPOV articles for the ITS analysis.

5.1.2 ITS Modeling for Article Writing Style. When an article has an active NPOV tag, it stays on top of the article for all revisions until its removal (Figure 1). As this can be considered a continuous treatment, we would expect the editors to be cautious for biased language and the articles to have a continuous drop in biased language coverage. Therefore, to measure the effect of an article having an active NPOV tag on the biased language coverage, we hypothesize a slope change in the article-level biased language usage after treatment. To model this, we use the ITS regression model in Equation 1 and use a linear model to fit in the R software.

5.2 Results: RQ1

Results of the ITS analysis for RQ1 is shown in Table 2 and the trends are shown in Figure 3. Through the ITS analysis, we observe that when an active NPOV tag is on the article, there is a statistically significant change in the slope of the trend of nine out of ten lexicons used to characterize biased language.

Robustness Check: The dataset used in this analysis was collected from the snapshot of articles in Wikipedia’s *NPOV dispute* category in early 2013. Some of the articles could contain NPOV issues that are inherently difficult to resolve and could be tagged for a long period of time. As this could be a potential confound in the editing behavior of such articles, as a robustness check, we re-ran the analysis on a subset of the data after removing articles that are NPOV tagged for more than three years. The distribution of time-intervals between the addition of the NPOV tag and the data collection time, and the result of the regression analysis for the subset of articles after removing potential outliers are shown in Appendix A. The treatment effect remains significant after excluding these long-running disputes.

6 RQ2: EDITOR-LEVEL EFFECTS OF NPOV CORRECTION

Next, we explore the editor-level effects of NPOV correction through the following research questions:

RQ2: What are the editor-level effects of NPOV correction?

RQ2a: What effect does NPOV correction have on the writing style of the treatment editors?

RQ2b: What effect does NPOV correction have on the engagement of treatment editors?

¹⁷We removed 47 articles in this way, which is nearly 2.7% of the total number of articles considered for RQ1.

Table 2. Results of RQ1 Analysis. Article lexicon coverage is computed for the textual content of the article results from each revision. Coefficient β_3 indicates the change in slope after treatment. Statistical significance after correcting for multiple comparisons using Benjamini and Hochberg [5] adjustment for false discovery rate are shown; *: $p < 0.05$, **: $p < 0.01$, ***: $p < 0.001$. Percentage change is computed as the change in post-treatment lexicon coverage after 40 revisions.

Lexicon	β_3	p -value		% change
Words to Watch	-1.51e-05	1.76e-13	***	-5.44
Hedges	-5.75e-06	2.76e-08	***	-3.92
Assertives	-3.57e-06	1.89e-10	***	-6.10
Positive Words	-4.48e-05	< 2e-16	***	-7.72
Negative Words	-2.31e-05	< 2e-16	***	-4.62
Factives	1.35e-06	8.76e-01		0.90
Implicatives	-2.63e-06	1.50e-05	***	-3.79
Report Verbs	-3.18e-06	1.27e-03	**	-1.90
Strong Subjectives	-5.78e-05	< 2e-16	***	-8.74
Weak Subjectives	-4.34e-05	< 2e-16	***	-2.95

Fig. 3. Article level lexicon coverage with respect to the treatment of NPOV tagging. We computed the average lexicon coverage rate of treatment articles' 40 revisions before and after the treatment revision. Lexicons for which we observe a statistically significant drop in slope ($p < 0.05$) are marked with an asterisk (*).

6.1 Editor Writing Style: RQ2a

6.1.1 Quantifying Editor Writing Style. We use the dataset we created for the editor level analysis (§ 3.5) to study RQ2a. Before measuring editor writing style, we performed a set of preprocessing steps to extract textual contribution of the treatment editors. We removed Wikipedia markup templates, URLs, and other non-textual content. We then tokenized the text and extracted tokens with two or more characters. To quantify editor writing style, we extracted the coverage of tokens from each of the ten bias-related lexicons (§ 4.1) and computed lexicon coverage rate as the number of lexicon tokens per total tokens in the text contributed by an editor in a revision. We considered the latest 40 pre-treatment edits and the earliest 40 post-treatment edits of the treatment editors for the ITS analysis.

6.1.2 Controlling for Editor Experience and Talk Page Discussion. Other factors may confound measurement of the effect of NPOV correction on editor-level writing. One such factor is the

experience of the editor: previous work shows that inexperienced editors respond differently to norm-enforcement activities such as edit reverts (e.g., [38]). To control for editor experience, we include it as a binary predictor in the regression model. Based on the criteria used in prior work [61], we consider editors who have contributed at least 250 edits prior to the treatment as *experienced editors* and others as *inexperienced editors*. In our dataset, 2,847 out of 4,378 editors had at least 250 edits prior to the treatment.

In addition to correcting content for NPOV, editors sometimes also post to the talkpage of the articles or the talkpage of the editor whose contribution was revised to discuss about the correction. These posts trigger an email prompt to the original editor. Thus, the talk page discussion acts as an additional “correction” method that could potentially amplify the treatment of correcting content with NPOV tags in the comment metadata. The discussion between the correcting and corrected editors during the treatment period could influence the future behavior of the corrected editor in terms of language usage and engagement. Therefore, we also include the presence of talk page discussion as an additional binary predictor, which indicates whether any of the following criteria is satisfied during a week prior or after the treatment: (1) correcting editor posts on corrected editor’s talk page; or (2) corrected editor posts on correcting editor’s talk page; or (3) both corrected editor and correcting editor post on the article talk page. In our dataset, 1,216 of the 4,378 NPOV corrections were accompanied by a talk page discussion.

6.1.3 ITS Modeling for Editor Writing Style. Unlike an active NPOV tag on the article, which stays as a continuous treatment for all later revisions until its removal, the editor who gets an NPOV correction only receives treatment once.¹⁸ Therefore, we hypothesize a change in the *level* of the trend of average editor writing style after treatment. We use the ITS regression model in Equation 1, excluding the term $\beta_3 TX_t$ (since we are only hypothesizing a level change) and using a linear model fit in the R software (*Model-I*). In addition to the terms in the ITS equation, as discussed in the previous subsection, we also used a second model (*Model-II*) including two binary control predictors for editor experience (*experienced*) and talk page discussion during treatment (*discussion*). We also add an interaction term between each of these additional predictors and the post-treatment indicator (i.e., interaction with dummy term X_t).

6.2 Results: RQ2a

To understand the trend of the change due to treatment, we first visualize the level change in average editor writing style from *Model-I* in Figure 4. NPOV correction is associated with a small level decrease in the usage of four out of ten lexicons used to characterize biased language: positive words, negative words, strong subjectives, and weak subjective. No significant treatment effect is observed for the remaining six lexicons. *Model-II* includes the control predictors for editor experience and talk page discussion. Relevant output from *Model-II* is shown in Table 3. The trend of level change (β_2) remains the same after including the control predictors. The coefficients for the *experienced* predictor are negative and significant for eight of the ten lexicons, which suggests that experienced editors use lower amount of non-neutral language.¹⁹ The coefficients for the effect of discussion after treatment ($D : X_t$) are not significant except for the negative words lexicon. None of the coefficients for the effect of editor experience after treatment ($E : X_t$) are significant, suggesting that editor experience does not influence the effect of NPOV correction on writing style.

¹⁸While 32.4% of the treatment editors are corrected more than once, repeated corrections are rare in the short window of 40 edits adjacent to the treatment.

¹⁹This finding also validates our choice of lexicons to characterize non-neutral language as we would expect experienced editors to better adhere to community writing style norms compared to newbies.

Fig. 4. Editor level lexicon coverage with respect to the treatment of NPOV correction. We computed the average lexicon coverage rate of treatment users' textual contributions 40 revisions before and after the treatment revision. Lexicons for which we observe a statistically significant level drop ($p < 0.05$) are marked with an asterisk (*).

Table 3. Results of RQ2 Analysis after controlling for editor experience and talk page discussion during treatment. Editor lexicon coverage is computed for their textual contribution in each revision. Coefficient β_2 indicates the change in level after treatment; D indicates whether there was any talk page discussion during treatment (reference: *discussed* = 0); E indicates whether the editor is considered experienced or not (reference: *experienced* = 0); X_t is the indicator variable for post-treatment; $D : X_t$ and $E : X_t$ are the interaction terms. Statistical significance after correcting for multiple comparisons using Benjamini and Hochberg [5] adjustment for false discovery rate are shown; *: $p < 0.05$, **: $p < 0.01$, ***: $p < 0.001$. Percentage change is computed as the level change following the treatment. Note that none of the coefficients for the $E : X_t$ interaction term are significant at $p < 0.05$.

Lexicon	β_2	% change	$D : X_t$	E	$E : X_t$
Words to Watch	-4.98e-04	-1.92	-5.90e-04	-2.98e-03	6.72e-04
Hedges	-8.76e-04	-3.85	-7.78e-05	-1.38e-03	6.59e-04
Assertives	1.58e-04	-3.32	-3.05e-04	-5.63e-04	-2.46e-04
Positive Words	-2.27e-03	**	-1.70e-04	-3.47e-03	1.48e-03
Negative Words	-2.55e-03	**	-1.92e-03	-3.66e-03	5.41e-05
Factives	3.51e-04	-0.10	-5.44e-04	1.48e-03	-3.78e-04
Implicatives	-2.78e-06	0.16	-1.72e-04	-6.44e-04	2.22e-04
Report Verbs	3.27e-04	-1.86	-7.11e-04	6.02e-04	-6.03e-04
Strong Subjectives	-2.48e-03	**	1.59e-04	-5.99e-03	-7.43e-06
Weak Subjectives	-2.70e-03	*	-2.18e-04	-7.20e-03	-1.88e-04

6.3 Editor Engagement: RQ2b

6.3.1 Quantifying Editor Engagement. To study RQ2b, we define engagement of a Wikipedia editor as any non-minor revisions they made to the article pages. We use the editor-level engagement dataset we created (§ 3.5) for this purpose and considered all treatment editors who have contributed at least ten revisions during the two months prior to the treatment.

6.3.2 ITS Modeling for Editor Engagement. Similar to the hypothesis for editor level lexicon coverage, we hypothesize a change in the level of editor engagement after treatment. We model this using

Table 4. Results of RQ2b Analysis: predictor coefficients of the count model of the zero-inflated negative binomial regression. The dependent variable, editor engagement is computed as number of edits during each of the three-day period. We considered 60 days prior to the treatment and 60 days after the treatment. *Post-treatment* is the indicator for whether the observation is post (=1) or pre (=0) treatment. *Timebin* is the three-day window; we consider 20 timebin windows pre- and post-treatment, which includes 120 days in total. *Discussion* is the indicator whether there was talk page discussion between correcting and corrected editors during treatment (*Discussion* = 1) or not. *Experienced* is the indicator for whether the editor is experienced (>250 edits, *Experienced* = 1) or not. Interaction terms include: Post-treatment : *Discussion*(=1) and Post-treatment : *Experienced*(=1). Statistical significance after correcting for multiple comparisons using Benjamini and Hochberg [5] adjustment for false discovery rate are shown;^f: $p < 0.05$, **: $p < 0.01$, ***: $p < 0.001$.

Predictor	Coefficient	p-value	
Post-treatment	0.4675	<2e-16	***
Timebin	0.0080	<2e-16	***
Discussion(=1)	0.2018	<2e-16	***
Experienced(=1)	1.8937	<2e-16	***
Post-treatment : Discussion(=1)	0.0491	0.0404	*
Post-treatment : Experienced(=1)	-0.6618	<2e-16	***

the ITS regression model in Equation 1 excluding the slope term $\beta_3 TX_t$, but with editor fixed effects. Because the dependent variable in this model is a count variable (i.e., number of revisions per day), we would use a Poisson regression model. However, summary statistics of the dependent variable show an excessive amounts of zeros (> 40%) and over-dispersion (variance $\gg$ mean). Therefore we used zero-inflated negative binomial regression model, as implemented in the *pscl* package²⁰ in R software. Similar to RQ2a, we include binary controls for *experienced* and *discussion*.

6.4 Results: RQ2b

Results of the ITS analysis for RQ2b are shown in Table 4. The interaction between the level change and control predictors *experienced* and *discussion* are shown in Figure 5. These results show a significant increase in engagement for inexperienced editors (editors with less than 250 edits at treatment time). The decrease in engagement for experienced editors is small after including the overall temporal trend. Talk page discussion during treatment is associated with a small increase in engagement for both experienced and inexperienced editors.

7 DISCUSSION

Next, we reflect on the article-level and editor-level effects of NPOV tagging, discuss the differences in success for articles and editors, and conclude by considering implications of our findings to Wikipedia and other online communities with writing style norms.

7.1 RQ1: Article-level Effects of NPOV Tagging

When a Wikipedia article is tagged for not adhering to the community norm of neutral point of view, a banner template is displayed on the article stating that “This article has multiple issues. Please help improve it or discuss these issues on the talk page.” Articles marked with any of the NPOV tags are listed under Wikipedia’s *NPOV disputes* category and editors are encouraged to make edits to resolve the NPOV issues. When an article has an active NPOV tag, we observe

²⁰<https://cran.r-project.org/web/packages/pscl/pscl.pdf>

Fig. 5. Visualization of RQ2b regression coefficients showing the pre- and post-treatment engagement levels for the control factors *experienced* and *discussion*.

a statistically significant, gradual reduction in biased language coverage for all lexicons except factives.²¹ The reduction is larger for Wikipedia words to watch (5.4%), strong subjectives (8.7%), positive words (7.7%), and assertives (5.4%), compared to other lexicons. Using these lexicons as a proxy for biased or non-NPOV language, it is evident that non-objective language used in Wikipedia articles decreases as editors make more revisions. This demonstrates the effectiveness of marking Wikipedia articles with NPOV tag in reducing the prevalence of non-neutral language.

The dataset we used for this analysis is collected from a single snapshot of the articles in the *NPOV dispute* category. Some of the articles which are controversial and hard to resolve could be NPOV tagged for a long period of time could affect the editing behavior. As a robustness check we repeated the analysis on the subset of the data after removing articles which has an active NPOV tag for more than three years. As we report in [Appendix A](#), the main results hold after removing these potential outliers.

7.2 RQ2: Editor-level Effects of NPOV Correction

Wikipedia’s policies around NPOV language encourage editors to revise article content to correct for biased language. When editors correct non-neutral language, those revisions are logged with comments such as “correcting for NPOV language”, “removing POV”, and “edits for NPOV” [67]. Using these comments, we extract revisions which are corrected for NPOV and using the revision history of the articles, we trace back to the editor who contributed the portion of text that was corrected for NPOV.

When an editor is corrected for NPOV, there is a significant reduction in their negative words usage (13.5%); however, we do not observe any statistically significant change for six of the ten lexicons relating to biased language, including Wikipedia’s own list of “words to watch” (RQ2a). These trends remain even after controlling for editor experience and talk page discussion during treatment. For RQ2b, we observe an increase in engagement for inexperienced editors after NPOV correction. One possible explanation is that NPOV correction helps inexperienced users become aware that their contributions are monitored by others, and this awareness could motivate them

²¹Note that compared to other lexicons, the set of factives is small in size—containing only 27 terms—and hence the coverage of this lexicon is small.

to contribute more. These results for RQ2 suggest that when corrected for NPOV, the quality of the writing style of editors does not improve significantly (at least, as measured by all but one of our lexicons), while it leads to a increase in editing activities for inexperienced editors. These findings partially conflict with the observations of Halfaker et al. [38], who found that revert actions demotivate new editors and reduce the quantity of work, even as they increase the overall quality of contributions. Halfaker et al. [38] also found that reverts affect editors differently based on the experience of the editor who makes the revert. Future work could perform similar analysis to further understand the effects of NPOV correction on engagement based on the experience of correcting editors.

7.3 Possible Reasons for the Differences in the Effectiveness of Norm Enforcement

While NPOV tagging helps articles to converge to the desired writing style, we do not find statistically significant evidence that these corrections encourage editors to adopt the desired writing style. One possible reason for the significant improvement in writing style of the NPOV tagged articles is the active contribution of a group of editors who are dedicated to improve the overall quality of Wikipedia articles. According to prior work on the classification of different types of work (refer § 2.5), editing tasks, and roles in Wikipedia, editors perform distinct roles, corresponding to distinct types of contributions. For example, Bryant et al. [11] found that while novice editors contribute only to articles related to their expertise, expert editors contribute to improve the overall quality of Wikipedia articles. Arazy et al. [4] identified several roles of Wikipedia editors, including quality assurance technicians, who are editors contribute towards patrolling Wikipedia and ensuring content quality. Further, Yang et al. [83] found that articles in different quality stages require different types of editors. Tools such as SuggestBot [20] recommends specific editing tasks (e.g., cleanup and rewrite) to editors based on their previous editing patterns and interests. These prior findings support our hypothesis that the NPOV tagged articles are of interest to specific editor roles, who are dedicated to make revisions to improve NPOV tagged articles.

NPOV tags are highly visible to Wikipedia editors via the banner template displayed in the article (Figure 1), often on the top, and the NPOV tagged articles are listed in a specific Wikipedia category (i.e., *NPOV disputes* category). These visible actions could attract more editors to contribute towards improving the tagged articles. However, the treatment of NPOV correction of editor contribution is not visible to the same extent because editors are not directly notified about their contribution being corrected for NPOV. Editors keep track of the revisions to their previously contributed articles via the “watchlist” tool [11], but this requires active actions from the editors.

7.4 Implications for Wikipedia and Online Language Moderation

Our findings in RQ2a show that the current regime of tag-and-correct does not help editors to significantly improve their adherence to the NPOV norm. This suggests the need for additional interventions targeted at editors. Possible strategies include explicit notification of their textual contribution being corrected for NPOV language, reminders about writing style norms, and incentives when they progress. Findings from RQ2b shows an increase in engagement after treatment for inexperienced editors; talk page discussion during treatment is also associated with a small increase in engagement for all editors. The increase in engagement for inexperienced editors could be due to the signal that the community is aware of their contributions. Interventions such as explicit notification of NPOV correction could include personalized messages [33] that could further increase the engagement of inexperienced editors. Other online collaborative content creation communities who wish to enforce linguistic style norms may consider implementing interventions similar to Wikipedia’s treatment of NPOV. Further community-level and individual-level interventions could help to converge to expected writing style norms.

In this work, we focused on NPOV, which is one of the three core content policies of Wikipedia. The other policies are “Verifiability” and “No original research”. These three policies jointly determine the quality of content acceptable in Wikipedia articles. Similar to NPOV tagging, these policies are enforced using tags such as {{citation needed}}, {{verification needed}}, {{original research}}, and {{synthesis}}. Our findings suggest that the NPOV tagging helps articles to converge to neutral language, but we did not find significant changes at individual editor language. As we discuss in § 7.3, various editor roles and the types of work performed by different editor roles could be one of the mechanisms by which NPOV tagging helps to improve article quality. A similar study could be performed to test the effectiveness of tagging for the other quality control policies of Wikipedia.

8 CHALLENGES, LIMITATIONS, AND FUTURE WORK

- **Limitations with characterizing biased language.** We use a set of style lexicons including Wikipedia’s “words to watch” to characterize non-neutral linguistic style. While similar approaches are used in prior work, these lexicons may not accurately and entirely capture non-neutral language. Machine learning could be applied to this problem by training on the edits that led to the application of an NPOV-related tag and correction [67]. We limit our analysis of language style to lexical methods; however, other aspects of language such as syntax, semantics, or pragmatics could also be studied (e.g., active vs. passive form).
- **Potential issues in accurately identifying the treatment groups.** The presence of reverts and vandalism imposed challenges in correctly identifying treatment articles and editors. We used a set of heuristics to reduce these issues (§ 3.2). However, our approach might not be perfect and could attribute an NPOV correction for an editor who is not the original contributor of the corrected text. This will result in an underestimation of treatment effects. Sophisticated algorithms to detect vandalism and reverts [34, 63], and algorithms to attribute authorship of revisioned content [30] could be used in future work to improve accurate identification of treatment editors.

Furthermore, the non-standard nature of NPOV tags and variations in the tags limit the recall of detecting tag addition and tag removal revisions. To detect NPOV corrections, following the approach of Recasens et al. [67], we searched for occurrences of “NPOV”, “POV”, or any case variations (§ 3.3) in the revision comments. However, other idiomatic terms such as “point of view”, “bias”, and “weasel words” could also be used to indicate NPOV corrections. Failing to detect these variants could result in an underestimation of the treatment effects.

- **Controlling for platform-level changes over time.** While we limit the time range of our analyses to control for significant platform-level changes, there is a possibility for internal and external factors, such as influx of new editors, site upgrades, policy changes, to influence the findings. The article dataset used in this work was collected in 2013, and the policies and normative behavior in Wikipedia may have changed since then. Future work could consider replicating this work using datasets collected in different time periods to analyze any longitudinal changes in normative behavior in the platform.
- **Assuming writing style changes in the short-term.** A potential limitation with using ITS is that, for ITS, the outcomes are expected to change either relatively quickly after an intervention is implemented or after a clearly defined lag [8]. In our work, we assume that NPOV enforcement makes the articles and editors to improve in short-term. This could be a potential limitation, because in some cases, edits may occur long after the initial tagging.
- **Restricting article level analysis to unresolved articles.** Our article dataset was collected from the snapshot of articles in the *NPOV dispute* in 2013. We use this single snapshot data and the articles in this dataset have an active NPOV tag (i.e., an unresolved NPOV issue). As a robustness check, we re-ran the regression for RQ1 using a subset of data after

removing articles with long standing disputes and the effect of NPOV tagging treatment remains significant. Using additional article data collection, the article level analysis could be expanded to cases where an NPOV tag was added and removed or where articles went through multiple NPOV tag additions/removals. This could provide insights about the effects after tag removal and repeated offenders, and future work could focus on this direction. From RQ1 we find that when an article has an active NPOV tag, subsequent revisions result in a decrease in bias language. Note that this finding does not imply that a tag is *required* to reduce biased language in articles.

9 CONCLUSION

While the Wikipedia NPOV tagging system has been in place for a long time, the impact of this tagging system has not been quantitatively measured. In this paper, we studied the effects of NPOV norm enforcement on Wikipedia using a corpus of NPOV-tagged articles and a set of style lexicons to characterize biased language. Focusing on the causal effect of NPOV norm enforcement both at the article and editor level, we found that NPOV tagging helps Wikipedia articles to converge to the accepted writing style, but no statistically significant improvements at editor level. Our findings suggest design improvements and interventions for Wikipedia and other online communities with strict language norms.

10 APPENDIX A: ROBUSTNESS CHECK FOR RQ1

The distribution of time-intervals between the addition of the NPOV tag and the data collection time is shown in [Figure 6](#).

Fig. 6. The distribution of time-intervals between the addition of the NPOV tag and the data collection time.

[Table 5](#) shows results of the ITS analysis for RQ1 on the subset of data after removing articles which are NPOV tagged for more than three years. The treatment effect remains significant after excluding these long-running disputes.

Table 5. Results of RQ1 Analysis robustness check on the subset of data after removing articles which are tagged for more than three years. Article lexicon coverage is computed for the textual content of the article results from each revision. Coefficient β_3 indicates the change in slope after treatment. Statistical significance after correcting for multiple comparisons using Benjamini and Hochberg [5] adjustment for false discovery rate are shown; *: $p < 0.05$, **: $p < 0.01$, ***: $p < 0.001$. Percentage change is computed as the change in post-treatment lexicon coverage after 40 revisions.

Lexicon	β_3	p -value		% change
Words to Watch	-1.36e-05	3.79e-15	***	-5.23
Hedges	-5.19e-06	3.66e-10	***	-3.21
Assertives	-3.15e-06	4.39e-10	***	-5.25
Positive Words	-2.62e-05	< 2e-16	***	-4.13
Negative Words	-1.25e-05	4.07e-12	***	-2.52
Factives	1.21e-06	3.16e-01		-0.06
Implicatives	-9.84e-07	3.66e-02	*	-2.08
Report Verbs	-5.92e-06	1.26e-04	**	-3.79
Strong Subjectives	-3.75e-05	< 2e-16	***	-8.74
Weak Subjectives	-2.80e-05	< 2e-16	***	-1.55

11 ACKNOWLEDGMENTS

We thank the anonymous reviewers for their helpful and constructive feedback. We thank Amy Bruckman for helpful discussions; Scott Appling, Stevie Chancellor, Sandeep Soni, and Ian Stewart for their valuable feedback on the paper. This research was supported by Air Force Office of Scientific Research award FA9550-14-1-0379, by National Institutes of Health award R01-GM112697, and by National Science Foundation awards 1452443 and 1111142.

REFERENCES

- [1] Khalid Al Khatib, Hinrich Schütze, and Cathleen Kantner. 2012. Automatic Detection of Point of View Differences in Wikipedia. In *Proceedings of the 24th International Conference on Computational Linguistics (COLING 2012)*. Association for Computational Linguistics, 33–50.
- [2] Maik Anderka and Benno Stein. 2012. A Breakdown of Quality Flaws in Wikipedia. In *Proceedings of the 2nd Joint WICOW/AIRWeb Workshop on Web Quality*. ACM, 11–18. <https://doi.org/10.1145/2184305.2184309>
- [3] Maik Anderka, Benno Stein, and Nedim Lipka. 2012. Predicting Quality Flaws in User-Generated Content: The Case of Wikipedia. In *Proceedings of the 35th International ACM SIGIR Conference on Research and Development in Information Retrieval*. ACM, 981–990. <https://doi.org/10.1145/2348283.2348413>
- [4] Ofer Arazy, Felipe Ortega, Oded Nov, Lisa Yeo, and Adam Balila. 2015. Functional Roles and Career Paths in Wikipedia. In *Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work & Social Computing*. ACM, 1092–1105. <https://doi.org/10.1145/2675133.2675257>
- [5] Yoav Benjamini and Yosef Hochberg. 1995. Controlling the False Discovery Rate: A Practical and Powerful Approach to Multiple Testing. *Journal of the Royal Statistical Society. Series B (Methodological)* (1995), 289–300.
- [6] Yochai Benkler. 2006. *The Wealth of Networks: How Social Production Transforms Markets and Freedom*. Yale University Press.
- [7] Kelly Bergstrom. 2011. “Don’t Feed the Troll”: Shutting Down Debate About Community Expectations on Reddit.com. *First Monday* 16, 8 (2011). <https://doi.org/10.5210/fm.v16i8.3498>
- [8] James Lopez Bernal, Steven Cummins, and Antonio Gasparrini. 2017. Interrupted Time Series Regression for the Evaluation of Public Health Interventions: A Tutorial. *International Journal of Epidemiology* 46, 1 (2017), 348–355.
- [9] Ivan Beschastnikh, Travis Kriplean, and David W McDonald. 2008. Wikipedian Self-Governance in Action: Motivating the Policy Lens.. In *Proceedings of the International AAAI Conference on Web and Social Media (ICWSM)*.
- [10] Pierre Bourdieu. 1991. *Language and Symbolic Power*. Harvard University Press.

- [11] Susan L Bryant, Andrea Forte, and Amy Bruckman. 2005. Becoming Wikipedian: Transformation of Participation in a Collaborative Online Encyclopedia. In *Proceedings of the 2005 ACM Conference on Supporting Group Work*. ACM, 1–10.
- [12] Gary Burnett and Laurie Bonnici. 2003. Beyond the FAQ: Explicit and Implicit Norms in Usenet Newsgroups. *Library & Information Science Research* 25, 3 (2003), 333–351. [https://doi.org/10.1016/S0740-8188\(03\)00033-1](https://doi.org/10.1016/S0740-8188(03)00033-1)
- [13] Ewa S Callahan and Susan C Herring. 2011. Cultural Bias in Wikipedia Content on Famous Persons. *Journal of the Association for Information Science and Technology* 62, 10 (2011), 1899–1915. <https://doi.org/10.1002/asi.21577>
- [14] Joshua McCann Kyle Frye Jed R. Brubaker Casey Fiesler, Jialun “Aaron” Jiang. 2018. Reddit Rules! Characterizing an Ecosystem of Governance. In *Proceedings of the International AAAI Conference on Web and Social Media (ICWSM)*.
- [15] Stevie Chancellor, Andrea Hu, and Munmun De Choudhury. 2018. Norms Matter: Contrasting Social Support Around Behavior Change in Online Weight Loss Communities. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. ACM. <https://doi.org/10.1145/3173574.3174240>
- [16] Eshwar Chandrasekharan, Umashanthi Pavalanathan, Anirudh Srinivasan, Adam Glynn, Jacob Eisenstein, and Eric Gilbert. 2017. You Can’t Stay Here: The Efficacy of Reddit’s 2015 Ban Examined Through Hate Speech. *Proceedings of the ACM on Human-Computer Interaction* 1, CSCW, Article 31 (Dec. 2017), 22 pages. <https://doi.org/10.1145/3134666>
- [17] Justin Cheng, Cristian Danescu-Niculescu-Mizil, and Jure Leskovec. 2015. Antisocial Behavior in Online Discussion Communities.. In *Proceedings of the International AAAI Conference on Web and Social Media(ICWSM)*. 61–70.
- [18] Thomas Chesney. 2006. An Empirical Examination of Wikipedia’s Credibility. *First Monday* 11, 11 (2006). <https://doi.org/10.5210/fm.v11i11.1413>
- [19] Andrea Cifforilli. 2003. Phantom Authority, Self-Selective Recruitment and Retention of Members in Virtual Communities: The Case of Wikipedia. *First Monday* 8, 12 (2003).
- [20] Dan Cosley, Dan Frankowski, Loren Terveen, and John Riedl. 2007. SuggestBot: Using Intelligent Task Routing to Help People Find Work in Wikipedia. In *Proceedings of the 12th International Conference on Intelligent User Interfaces*. ACM, 32–41. <https://doi.org/10.1145/1216295.1216309>
- [21] Cristian Danescu-Niculescu-Mizil, Robert West, Dan Jurafsky, Jure Leskovec, and Christopher Potts. 2013. No Country for Old Members: User Lifecycle and Linguistic Change in Online Communities. In *Proceedings of the 22nd International Conference on World Wide Web*. ACM, 307–318. <https://doi.org/10.1145/2488388.2488416>
- [22] Sanmay Das, Allen Lavoie, and Malik Magdon-Ismael. 2016. Manipulation Among the Arbiters of Collective Intelligence: How Wikipedia Administrators Mold Public Opinion. *ACM Transactions on the Web (TWEB)* 10, 4 (2016), 24.
- [23] Luca De Alfaro and Michael Shavlovsky. 2013. Attributing authorship of revisioned content. In *Proceedings of the 22nd international conference on World Wide Web*. ACM, 343–354.
- [24] Angelo Di Iorio, Fabio Vitali, and Stefano Zacchiroli. 2008. Wiki Content Templating. In *Proceedings of the 17th International Conference on World Wide Web*. ACM, 615–624. <https://doi.org/10.1145/1367497.1367581>
- [25] Nora A Draper. 2018. Distributed intervention: Networked Content Moderation in Anonymous Mobile Spaces. *Feminist Media Studies* (2018), 1–17.
- [26] Michelle Broder Van Dyke. 2015. Reddit Users Revolt After Site Bans “Fat People Hate” And Other Communities. <https://www.buzzfeed.com/mbvd/reddit-users-revolt-after-site-bans-fat-people-hate-and-oth>. *Buzzfeed News* (2015).
- [27] Antonella Elia. 2006. An Analysis of Wikipedia Digital Writing. In *Proceedings of the Workshop on NEW TEXT Wikis and Blogs and Other Dynamic Text Sources*.
- [28] William Emigh and Susan C Herring. 2005. Collaborative Authoring on the Web: A Genre Analysis of Online Encyclopedias. In *System Sciences, 2005. HICSS’05. Proceedings of the 38th Annual Hawaii International Conference on. IEEE*, 99a–99a.
- [29] Oliver Fersckhe, Iryna Gurevych, and Marc Rittberger. 2013. The Impact of Topic Bias on Quality Flaw Prediction in Wikipedia. In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, Vol. 1. 721–730.
- [30] Fabian Flöck and Maribel Acosta. 2014. WikiWho: Precise and Efficient Attribution of Authorship of Revisioned Content. In *Proceedings of the 23rd International Conference on World Wide Web*. ACM, 843–854. <https://doi.org/10.1145/2566486.2568026>
- [31] Andrea Forte and Amy Bruckman. 2005. Why Do People Write for Wikipedia? Incentives to Contribute to Open-Content Publishing. *Proceedings of the 2005 ACM Conference on Supporting Group Work* 5 (2005), 6–9.
- [32] Andrea Forte and Amy Bruckman. 2006. From Wikipedia to the Classroom: Exploring Online Publication and Learning. In *Proceedings of the 7th International Conference on Learning Sciences*. International Society of the Learning Sciences, 182–188.
- [33] R Stuart Geiger, Aaron Halfaker, Maryana Pinchuk, and Steven Walling. 2012. Defense Mechanism or Socialization Tactic? Improving. In *Proceedings of the Sixth International AAAI Conference on Weblogs and Social Media*.
- [34] R Stuart Geiger and David Ribes. 2010. The Work of Sustaining Order in Wikipedia: The Banning of a Vandal. In *Proceedings of the 2010 ACM Conference on Computer Supported Cooperative Work*. ACM, 117–126.

- [35] Shane Greenstein and Feng Zhu. 2012. *Collective Intelligence and Neutral Point of View: the Case of Wikipedia*. Technical Report. National Bureau of Economic Research. <https://doi.org/10.3386/w18167>
- [36] James Grimmelman. 2015. The Virtues of Moderation. *Yale JL & Tech.* 17 (2015), 42.
- [37] Aaron Halfaker, R Stuart Geiger, Jonathan T Morgan, and John Riedl. 2013. The Rise and Decline of an Open Collaboration System: How Wikipedia's Reaction to Popularity is Causing its Decline. *American Behavioral Scientist* 57, 5 (2013), 664–688. <https://doi.org/10.1177/0002764212469365>
- [38] Aaron Halfaker, Aniket Kittur, and John Riedl. 2011. Don't Bite the Newbies: How Reverts Affect the Quantity and Quality of Wikipedia Work. In *Proceedings of the 7th International Symposium on Wikis and Open Collaboration*. ACM, 163–172. <https://doi.org/10.1145/2038558.2038585>
- [39] Manoj Harpalani, Michael Hart, Sandesh Singh, Rob Johnson, and Yejin Choi. 2011. Language of Vandalism: Improving Wikipedia Vandalism Detection Via Stylometric Analysis. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies: Short papers-Volume 2*. Association for Computational Linguistics, 83–88.
- [40] Livnat Herzig, Alex Nunes, and Batia Snir. 2011. An Annotation Scheme for Automated Bias Detection in Wikipedia. In *Proceedings of the 5th Linguistic Annotation Workshop*. Association for Computational Linguistics, 47–55.
- [41] Michael A Hogg and Scott A Reid. 2006. Social Identity, Self-Categorization, and the Communication of Group Norms. *Communication Theory* 16, 1 (2006), 7–30. <https://doi.org/10.1111/j.1468-2885.2006.00003.x>
- [42] Joan B Hooper. 1974. *On Assertive Predicates*. Indiana University Linguistics Club.
- [43] Ken Hyland. 2005. *Metadiscourse: Exploring Interaction in Writing*. A&C Black.
- [44] Dell Hymes. 1972. On Communicative Competence. *Sociolinguistics* 269293 (1972), 269–293.
- [45] Mikolaj Jan Piskorski and Andreea Gorbatai. 2017. Testing Coleman's Social-Norm Enforcement Mechanism: Evidence from Wikipedia. *Amer. J. Sociology* 122, 4 (2017), 1183–1222. <https://doi.org/10.1086/689816>
- [46] Lauri Karttunen. 1971. Implicative Verbs. *Language* (1971), 340–358. <https://doi.org/10.2307/412084>
- [47] Brian Keegan, Darren Gergle, and Noshir Contractor. 2013. Hot Off the Wiki: Structures and Dynamics of Wikipedia's Coverage of Breaking News Events. *American Behavioral Scientist* 57, 5 (2013), 595–622. <https://doi.org/10.1177/0002764212469367>
- [48] Sara Kiesler, Robert Kraut, Paul Resnick, and Aniket Kittur. 2012. *Regulating Behavior in Online Communities. Building Successful Online Communities: Evidence-Based Social Design*. MIT Press, Cambridge, MA (2012).
- [49] Amy Jo Kim. 2000. *Community Building on the Web: Secret Strategies for Successful Online Communities*. Addison-Wesley Longman Publishing Co., Inc.
- [50] Paul Kiparsky and Carol Kiparsky. 1970. Fact. ed. M. Bierwisch and K. Heidolph. *Progress in Linguistics* (1970).
- [51] Travis Kriplean, Ivan Beschastnikh, and David W McDonald. 2008. Articulations of Wikiwork: Uncovering Valued Work in wikipedia Through Barnstars. In *Proceedings of the 2008 ACM conference on Computer Supported Cooperative Work*. ACM, 47–56. <https://doi.org/10.1145/1460563.1460573>
- [52] Cliff Lampe and Paul Resnick. 2004. Slash (dot) and Burn: Distributed Moderation in a Large Online Conversation Space. In *Proceedings of the SIGCHI conference on Human Factors in Computing Systems*. ACM, 543–550.
- [53] Cliff Lampe, Paul Zube, Jusil Lee, Chul Hyun Park, and Erik Johnston. 2014. Crowdsourcing Civility: A Natural Experiment Examining the Effects of Distributed Moderation in Online Forums. *Government Information Quarterly* 31, 2 (2014), 317–326. <https://doi.org/10.1016/j.giq.2013.11.005>
- [54] Andrew Lih. 2004. Wikipedia as Participatory Journalism: Reliable Sources? Metrics for Evaluating Collaborative Media as a News Resource. *Nature* 3, 1 (2004).
- [55] Nedim Lipka and Benno Stein. 2010. Identifying Featured Articles in Wikipedia: Writing Style Matters. In *Proceedings of the 19th International Conference on World Wide Web*. ACM, 1147–1148. <https://doi.org/10.1145/1772690.1772847>
- [56] Bing Liu, Mingqing Hu, and Junsheng Cheng. 2005. Opinion Observer: Analyzing and Comparing Opinions on the Web. In *Proceedings of the 14th International Conference on World Wide Web*. ACM, 342–351. <https://doi.org/10.1145/1060745.1060797>
- [57] Tanushree Mitra Mattia Samory. 2018. Conspiracies Online: User Discussions in a Conspiracy Community Following Dramatic Events. In *Proceedings of the International AAAI Conference on Web and Social Media (ICWSM)*.
- [58] Uwe Matzat and G Rooks. 2014. Styles of Moderation in Online Health and Support Communities: An Experimental Comparison of Their Acceptance and Effectiveness. *Computers in Human Behavior* 36 (2014), 65–75. <https://doi.org/10.1016/j.chb.2014.03.043>
- [59] Hacker News. 2018. Hacker News FAQ. <https://news.ycombinator.com/newsfaq.html> [Online; accessed 22-July-2018].
- [60] Oded Nov. 2007. What Motivates Wikipedians? *Commun. ACM* 50, 11 (2007), 60–64. <https://doi.org/10.1145/1297797.1297798>
- [61] Katherine Panciera, Aaron Halfaker, and Loren Terveen. 2009. Wikipedians Are Born, Not Made: A Study of Power Editors on Wikipedia. In *Proceedings of the ACM 2009 International Conference on Supporting Group Work*. ACM, 51–60. <https://doi.org/10.1145/1531674.1531682>

- [62] Martin Potthast, Benno Stein, and Robert Gerling. 2008. Automatic Vandalism Detection in Wikipedia. In *European Conference on Information Retrieval*. Springer, 663–668. https://doi.org/10.1007/978-3-540-78646-7_75
- [63] Reid Priedhorsky, Jilin Chen, Shyong Tony K Lam, Katherine Panciera, Loren Terveen, and John Riedl. 2007. Creating, Destroying, and Restoring Value in Wikipedia. In *Proceedings of the 2007 International ACM Conference on Supporting Group Work*. ACM, 259–268. <https://doi.org/10.1145/1316624.1316663>
- [64] Sheizaf Rafaeli and Yaron Ariel. 2008. Online Motivational Factors: Incentives for Participation and Contribution in Wikipedia. *Psychological Aspects of Cyberspace: Theory, Research, Applications* (2008), 243–267.
- [65] Mahmudur Rahman, Bogdan Carbutar, Jaime Ballesteros, George Burri, and Duen Horng Chau. 2014. Turning the Tide: Curbing Deceptive Yelp Behaviors. In *Proceedings of the 2014 SIAM International Conference on Data Mining*. SIAM, 244–252. <https://doi.org/10.1137/1.9781611973440.28>
- [66] Joseph M Reagle Jr. 2010. “Be Nice”: Wikipedia Norms for Supportive Communication. *New Review of Hypermedia and Multimedia* 16, 1-2 (2010), 161–180.
- [67] Marta Recasens, Cristian Danescu-Niculescu-Mizil, and Dan Jurafsky. 2013. Linguistic Models for Analyzing and Detecting Biased Language. In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*.
- [68] Yuqing Ren, F Maxwell Harper, Sara Drenner, Loren Terveen, Sara Kiesler, John Riedl, and Robert E Kraut. 2012. Building Member Attachment in Online Communities: Applying Theories of Group Identity and Interpersonal Bonds. *Mis Quarterly* (2012), 841–864.
- [69] Ellen Riloff and Janyce Wiebe. 2003. Learning Extraction Patterns for Subjective Expressions. In *Proceedings of the 2003 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, 105–112. <https://doi.org/10.3115/1119355.1119369>
- [70] Joachim Schroer and Guido Hertel. 2009. Voluntary Engagement in an Open Web-based Encyclopedia: Wikipedians and Why They Do It. *Media Psychology* 12, 1 (2009), 96–120. <https://doi.org/10.1080/15213260802669466>
- [71] Kay Kyeongju Seo. 2007. Utilizing Peer Moderating in Online Discussions: Addressing the Controversy Between Teacher Moderation and Nonmoderation. *The American Journal of Distance Education* 21, 1 (2007), 21–36. <https://doi.org/10.1080/08923640701298688>
- [72] Pnina Shachaf and Noriko Hara. 2010. Beyond Vandalism: Wikipedia Trolls. *Journal of Information Science* 36, 3 (2010), 357–370. <https://doi.org/10.1177/0165551510365390>
- [73] Sara Sood, Judd Antin, and Elizabeth Churchill. 2012. Profanity Use in Online Communities. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, 1481–1490. <https://doi.org/10.1145/2207676.2208610>
- [74] Jakob Voss. 2005. Measuring Wikipedia. In *Proceedings of 10th International Conference of the International Society for Scientometrics and Informetrics*.
- [75] Claudia Wagner, David Garcia, Mohsen Jadidi, and Markus Strohmaier. 2015. It’s a Man’s Wikipedia? Assessing Gender Inequality in an Online Encyclopedia. In *Proceedings of the International AAAI Conference on Web and Social Media (ICWSM)*. 454–463.
- [76] Howard T Welsler, Dan Cosley, Gueorgi Kossinets, Austin Lin, Fedor Dokshin, Geri Gay, and Marc Smith. 2011. Finding Social Roles in Wikipedia. In *Proceedings of the 2011 iConference*. ACM, 122–129.
- [77] Wikipedia. 2018. Neutral Point of View. https://en.wikipedia.org/wiki/Wikipedia:Neutral_point_of_view [Online; accessed 15-April-2018].
- [78] Wikipedia. 2018. Wikipedia. <https://en.wikipedia.org/wiki/Wikipedia> [Online; accessed 15-April-2018].
- [79] Wikipedia. 2018. Wikipedia Words to Watch. https://en.wikipedia.org/wiki/Wikipedia:Manual_of_Style/Words_to_watch [Online; accessed 15-April-2018].
- [80] Wikipedia. 2018. Wikipedia:Expectations and Norms of the Wikipedia Community. https://en.wikipedia.org/wiki/Wikipedia:Expectations_and_norms_of_the_Wikipedia_community [Online; accessed 15-April-2018].
- [81] Wikipedia. 2018. Wikipedia:Manual of Style. https://en.wikipedia.org/wiki/Wikipedia:Manual_of_Style [Online; accessed 15-April-2018].
- [82] Kevin Wise, Brian Hamman, and Kjerstin Thorson. 2006. Moderation, Response Rate, and Message Interactivity: Features of Online Communities and Their Effects on Intent to Participate. *Journal of Computer-Mediated Communication* 12, 1 (2006), 24–41. <https://doi.org/10.1111/j.1083-6101.2006.00313.x>
- [83] Diyi Yang, Aaron Halfaker, Robert E Kraut, and Eduard H Hovy. 2016. Who Did What: Editor Role Identification in Wikipedia. In *Proceedings of the International AAAI Conference on Web and Social Media (ICWSM)*. 446–455.
- [84] Heng-Li Yang and Cheng-Yu Lai. 2010. Motivations of Wikipedia Content Contributors. *Computers in Human Behavior* 26, 6 (2010), 1377–1383. <https://doi.org/10.1016/j.chb.2010.04.011>

Received April 2018; revised July 2018; accepted September 2018