

How Inclusive Are Wikipedia’s Hyperlinks in Articles Covering Polarizing Topics?

Cristina Menghini

Computer Science

Brown University

Providence, Rhode Island, USA

cristina_menghini@brown.edu

Aris Anagnostopoulos

Sapienza University

Rome, Italy

aris@diag.uniroma1.it

Eli Upfal

Computer Science

Brown University

Providence, Rhode Island, USA

eli_upfal@brown.edu

Abstract—Wikipedia relies on an extensive review process to verify that the content of each individual page is unbiased and presents a “neutral point of view.” Less attention has been paid to possible biases in the hyperlink structure of Wikipedia, which has a significant influence on the user’s exploration process when visiting more than one page. The evaluation of hyperlink bias is challenging because it depends on the global view rather than the text of individual pages.

In this paper, we focus on the influence of the interconnect topology between articles describing complementary aspects of polarizing topics. We introduce a novel measure of *exposure to diverse information* to quantify users’ exposure to different aspects of a topic throughout an entire surfing session, rather than just one click ahead. We apply this measure to six polarizing topics (e.g., *gun control* and *gun right*), and we identify cases in which the network topology significantly limits the exposure of users to diverse information on the topic, encouraging users to remain in a *knowledge bubble*. Our findings demonstrate the importance of evaluating Wikipedia’s network structure in addition to the extensive review of individual articles.

Index Terms—wikipedia, diversity, web

I. INTRODUCTION

Knowledge on Wikipedia is distributed across articles interconnected via hyperlinks. According to Wikipedia’s Linking Manual [1], “*Internal links can add to the cohesion and utility of Wikipedia, allowing readers to deepen their understanding of a topic by conveniently accessing other articles.*” Consequently, users are *directly* exposed to an article’s content and *indirectly* exposed to the content of the pages it points to.

Wikipedia’s pages offer high-quality content with emphasis on an unbiased, neutral point of view (NPOV) [2]–[4], thanks to numerous policies and guidelines [5], [6]. Although it provides tools to support the community for curating pages, it lacks a systematic way to contextualize them within the more general articles’ network. Indeed, it is hard to evaluate the extent to which the current hyperlinks satisfy their purpose, especially in connecting articles related to a broad topic.

The majority of users who look for a specific information are likely to find their answers on the first Wikipedia page they are visiting [7], whereas about 20% of Wikipedia’s users follow hyperlinks within Wikipedia to develop a broad view

of a subject.¹ It is therefore important to investigate whether the link structure leads users to visit pages presenting broad and diverse aspects of their topic of interest. We initiate this important study by concentrating on polarizing topics spanning across multiple articles.

For example, consider the topic *abortion*, which is distributed across multiple articles on Wikipedia. Because of its *polarizing* nature, we recognize pages about events, people, or organizations that are associated either with *pro-choice* or *pro-life*. For instance, the page *Abortion-rights movements* describes organizations related to *pro-choice* view. In this particular page, we identify 15 links pointing to articles about *pro-choice* subjects and only 3 hyperlinks directed to *pro-life* related pages. Furthermore, if we consider articles at distance 2 from the page *Abortion-rights movements*, then there are 4 times more pages associated with *pro-choice* than articles associated with *pro-life* subjects. Similar counting, starting from the page *Anti-abortion movements*, shows 18 outgoing links to *pro-life* pages and only 1 to a *pro-choice* article. At distance two we have 15 times more pages related to the category *pro-life* than pages related to *pro-choice*.

The example above demonstrates unbalanced hyperlink structure in Wikipedia that may influence users’ exposure to diverse information on a topic. On another, albeit very different, platform a recent work [9] empirically showed that its recommender system contributes to radicalizing users’ pathways. Given the major role of Wikipedia as a popular primary source of knowledge, it is important to evaluate the effect of its hyperlink structure on user navigation, to guarantee a balanced access to well-rounded knowledge.

Evaluating the influence of the hyperlink topology is challenging because it requires a broad view of the network topology, not just the text of a single article. Such a view, and a technique to analyzing it, is not readily available to Wikipedia editors. In this work, we develop an algorithmic approach to quantify users’ exposure among a set of articles. Then, we audit the extent to which the current Wikipedia’s link structure allows users to browse different stances of polarizing topics.

¹For the English Wikipedia, the number of unique devices is around 800 million per month. If a device corresponds to a user, around 160 million of users click at least one link throughout their visit on Wikipedia [8].

Our main contributions are the following:

- We initiate the study of the hyperlink network’s role in driving users to explore articles of different categories. We investigate this on a set of polarizing topics.
- We design two metrics, the *exposure to diverse information* and the (*mutual*) *exposure to diverse information*, which quantify the likelihood of visiting pages belonging to different sets of articles click-after-click (Sect. IV).
- We apply our new measures to Wikipedia hyperlink subgraphs related to six polarizing topics. We identify cases in which the Wikipedia’s hyperlink network significantly limits the exposure of users to diverse information on the topic (Sect. V).

The code to replicate the analysis is available at <https://github.com/CriMenghini/WikiNetBias>.

II. RELATED WORKS

Improving Wikipedia. Previous works proposed semi-automated procedures to improve Wikipedia’s quality by checking the veracity of references [10], [11], suggesting articles’ structure [12], looking for hoaxes [13], or recommending links [14], [15]. Among these tools, none provides a measure to evaluate the link-based relationship across articles of diverse categories. In this work, we define such metrics (Sect. IV-B).

Wikipedia Navigation. The literature still lacks a model that generalizes Wikipedia’s users’ behavior. Previous studies [16]–[19] focused on modeling and predicting human navigation relying on traces from games [19]–[23]. Even though such games provide valuable insights on how users exploit links to move across concepts, other studies showed that users display different behavioral patterns depending on their information needs and the links’ position within pages [7], [24], [25]. We exploit such insights to define a general model mimicking localized and in-depth topic exploration (Sect. IV).

Wikipedia Categorization. When dealing with polarizing topics, one needs to distinguish between pages belonging to different side of the topic. Because of the topic granularity, it is hard to rely on automated techniques to categorize articles. Thus, we refrain from using supervised tools as ORES² or topic modeling [26], in favor of a mining procedure employed in [27], exploiting actual Wikipedia’s categories (Sect. III).

Polarization on Social Media. Many works aim to quantify polarization on social media [28]–[34]. Random walk controversy [33] quantifies to what extent opinionated users are exposed to their own opinion rather than the opposite. Bubble radius [34] works on bipartite information networks and estimates the expected number of clicks to navigate from a page v to any page of opposing opinion. We focus on the metrics that better relate to our metric of *exposure to diverse information* (ExDIN). Differently from these two metrics, our measure of *exposure to diverse information* works on multi-partite information networks and quantifies users’ exposure to diverse information click-after-click.

²<https://www.mediawiki.org/wiki/ORES>




Fig. 1: On the left, the original Wikipedia graph. On the right, the final topic-induced network. The dashed circles in W are the set of nodes used to build the topic-induced network G . The colors red and blue refer to the sets P and \bar{P} , respectively. Green and yellow are N and s respectively. We keep the image tidy and do not specify the edges’ direction.

Cultural bias on Wikipedia. Recent works found the presence of cultural bias in the same articles of different languages [35] and gender biases [36], [37]. These content-based analyses prove that Wikipedia can be subjected to bias. We decided to investigate bias on a novel topological perspective.

III. PRELIMINARIES

We encode a topic into a topic-induced network, a subgraph of the entire English Wikipedia’s graph $W = (A, L)$ (Fig. 1). The nodes of the graph are *articles* [38], and edges are links connecting pages, *wikilinks*.³ Among all articles, we identify a set of pages $\mathcal{T} \subset A$, about the topic. We partition these pages into two sets P and \bar{P} (i.e., $P \cap \bar{P} = \emptyset$ and $P \cup \bar{P} = \mathcal{T}$), each gathering articles about the same side of the topic. In addition to the articles in \mathcal{T} , we collect in \mathcal{N} the pages at one-hop distance from them. In this way, we can consider the chances of moving across partitions via articles not necessarily related to the topic. To reduce the complexity of our analysis, we cluster all the pages in $A \setminus (\mathcal{T} \cup \mathcal{N})$, into one super node s . Note that s is only connected to vertices in \mathcal{N} . For each node $v \in \mathcal{N}$, we can have multiple edges going to s , which we compress into one. Respectively, s can have multiple links to node $v \in \mathcal{N}$, compressed into one as well.

A topic-induced network is the directed weighted graph $G = (V, E)$, whose set of articles V is $\mathcal{T} \cup \mathcal{N} \cup \{s\}$ of size $n + 1$, and edges E are the links connecting them. The edge weights are *transition probabilities* stored in a row-stochastic matrix $M \in [0, 1]^{|V| \times |V|}$, whose entry $m_{i,j}$ is the probability that from page i a reader moves to j , and it is set to 0 if $(i, j) \notin E$.

In practice, to build a topic-induced network, we first extract the entire Wikipedia’s network from a complete English Wikipedia dump.⁴ Then, we only retain the graph induced by \mathcal{T} , whose articles are selected and partitioned according to the strategy adopted in [27]. For instance, the topic *abortion* polarizes into *pro-life* (P) and *pro-choice* (\bar{P}) articles. The pro-life subcorpus consists of all articles categorized under

³We exclude links within the same page and resolve all the redirects [39]. We do consider links in the infoboxes, which are summary standardized tables at the top-right corner of articles.

⁴We refer to the dump of September 2020.

Topic	P	\bar{P}	Seed P	Seed \bar{P}
Abortion	Pro-life	Pro-choice	Anti-abortion movement	Abortion-rights movement
Cannabis	Prohibition	Activism	Cannabis prohibition	Cannabis activism
Guns	Control	Rights	Gun control advocacy groups	Gun rights advocacy groups
Evolution	Creationism	Evolutionary biology	Creationism	Evolutionary biology
Racism	Racism	Anti-racism	Racism	Anti-racism
LGBT	Discrimination	Support	Discrimination against LGBT people	LGBT rights movement

TABLE I: For each topic, the table indicates the partitions P and \bar{P} to which each standing corresponds. Moreover, we report the seed category for each partition.

Topic	$ V \setminus \{s\} $	$ P $	$ \bar{P} $	$ \mathcal{N} $	$ E $	$ E_{P \rightarrow P} $	$ E_{\bar{P} \rightarrow P} $	$ E_{\mathcal{N} \rightarrow P} $	$ E_{\mathcal{N} \rightarrow \bar{P}} $
<i>Abortion</i>	56056	469	291	55296	2.1M	205	97	21396	29889
<i>Cannabis</i>	32743	45	231	32470	1.1M	8	6	656	27823
<i>Guns</i>	65743	167	187	65393	2.5M	98	115	56702	16608
<i>Evolution</i>	84788	342	1334	83113	1.99M	391	135	15601	58720
<i>Racism</i>	129963	1024	1022	127953	4.8M	746	560	74354	58195
<i>LGBT</i>	150563	459	640	149479	4.6M	195	143	92975	81706

TABLE II: Networks’ statistics.

the *seed* category “Anti-abortion movement” and its subcategories. Similarly, we obtain the pro-choice corpus starting from the category “Abortion-rights movement.” Because we want the partitions to be disjoint, articles belonging to both “Anti-abortion movement” and “Abortion-rights movement” are assigned to \mathcal{N} .

In fact, as a consequence of Wikipedia’s Neutral Point of View (NPOV) policy [4], we assume articles’ content to “fairly and proportionately represent all the significant views that have been published by reliable sources on the topic.” Moreover, as subcategories are often redundant or not entirely related to the parent category, we check them manually, discarding categories whose names do not include topic-specific keywords.

A. Topic-Induced Networks

We collect the topic-induced networks related to six different polarizing topics: *abortion*, *cannabis*, *guns*, *evolution*, *LGBT*, and *racism* (Tab. I).

1) **Partitions:** In Tab. II,⁵ we observe that the size of P and \bar{P} differs substantially, for all the topics but *racism* and *guns*. The disproportionate number of articles does not imply an imbalance in content representation, but it can affect the partition’s exposure within the entire Wikipedia network. The sizes of P and \bar{P} are not linear in the number of edges across partitions. For instance, although the nodes in *pro-life* are twice as many as those in *pro-choice*, the links pointing to *pro-choice* are 36% more than those pointing to *pro-life*. This happens, with different magnitude, also for *guns* and *LGBT*.

2) **Hyperlinks across partitions:** The direct exposure of users in P to pages in \bar{P} , depends on the number of links

⁵We add to the set \mathcal{N} the articles assigned to both partitions. The size of such intersections is: 2 (*abortion*), 3 (*cannabis*), 2 (*evolution*), 1 (*guns*), 5 (*LGBT*), 7 (*racism*). Because we do not remove these articles, they act as bridges connecting P and \bar{P} in sessions longer than one click.




Fig. 2: Percentage of edges across and within partitions in each topic-induced network (red) and a random graph with the same degree distribution (orange). Topics in order are: *abortion*, *cannabis*, *guns*, *evolution*, *racism*, and *LGBT*.

p-values	$\mathcal{N} \rightarrow P$ vs. $\mathcal{N} \rightarrow \bar{P}$	Incoming higher avg.	$P \rightarrow \mathcal{N}$ vs. $\bar{P} \rightarrow \mathcal{N}$	Outgoing higher avg.
<i>Abortion</i>	$8.4 \cdot 10^{-2}$	-	$9.6 \cdot 10^{-1}$	-
<i>Cannabis</i>	$6.5 \cdot 10^{-8}(**)$	Activism	$1.6 \cdot 10^{-13}(***)$	Activism
<i>Guns</i>	$1.3 \cdot 10^{-4}(**)$	Control	$5.1 \cdot 10^{-2}$	-
<i>Evolution</i>	$7.2 \cdot 10^{-3}(**)$	Creationism	$4.9 \cdot 10^{-5}(**)$	Creationism
<i>Racism</i>	$6.2 \cdot 10^{-6}(**)$	Anti-racism	$3.3 \cdot 10^{-7}(**)$	Anti-racism
<i>LGBT</i>	$1.4 \cdot 10^{-2}(**)$	Discrimination	$1.4 \cdot 10^{-5}(**)$	Discrimination

TABLE III: We report the p-values of t-tests ($\alpha = 0.05$) on (1) the number of links from \mathcal{N} to P and \bar{P} (first column), and (2) the number of links to \mathcal{N} from P and \bar{P} (third column). In the second and fourth columns we indicate which partition is significantly more connected to the rest of Wikipedia. Statistics are computed after bootstrapping the distributions of from \mathcal{N} to P and \bar{P} and vice versa.

connecting the two partitions.⁶ To study their connectivity, we compare the portion of links in pages of P pointing to P and \bar{P} , with the same quantities expected on a random graph with the same degree sequence. In Fig. 2, we observe that most of the hyperlinks point to pages of the same partition. On average fewer than 25% of links point toward the opposing partition, which is against the 50% expected on a random graph. The differences between the real and expected number of hyperlinks highlight that (1) links are, obviously, not randomly placed, (2) the strength of connections within and between partitions is skewed w.r.t. the distribution of edges conditioned on the number of nodes and their degree. Furthermore, we speculate that the higher number of hyperlinks directed to pages of the same partition is due to the intrinsic clustered nature of Wikipedia [40], [41].

3) **Topic connectivity to the rest of Wikipedia:** We briefly investigate the connectivity between P (resp. \bar{P}) with the rest of pages connected to it (i.e., \mathcal{N}). In Tab. III, we observe that for all the topics, but *abortion*, the average number of links coming from articles in \mathcal{N} and pointing to articles in \mathcal{N} is significantly higher for one of the two partitions.

4) **Distribution of across-partition links:** If across-partition links are uniformly placed within articles of a partition, users starting from an arbitrary node in the partition have the chance

⁶We note that the number of edges across *cannabis*’s partitions is low, nevertheless we keep the topic because on sessions longer than 1 click there are other paths connecting the partitions.




Fig. 3: Percentage of articles in P connected to \bar{P} . Topics are: *abortion, cannabis, guns, evolution, racism* and *LGBT*.

to visit pages about another branch of the topic. However, we observe that in our networks only a small subset of the pages expose their visitors to another branch of the topic: Fig. 3 shows that the average percentage of pages connecting to the other partition is 25% (average of 1.8 links per page), thus most of the nodes are not connected to the other partition.

5) **Weight distribution of across-partition links:** The likelihood of traversing a link connected to the other side is conditioned on the number of links in a page. Fig. 4 shows that, for each topic, there is one partition whose average probability of traversing an across-partition hyperlink is statistically higher than the other partition (according to t -tests with $\alpha = 0.05$). For instance, the average chance to go from *creationism* to *evolutionary biology* are significantly lower than moving in the opposite direction.

IV. METRICS

In this section we introduce the metrics to quantify the exposure to diversity, accounting for (1) the across-partition edges distribution over nodes, (2) the likelihood of traversing a link toward the other partition, and (3) the average exposure to diversity of all pages in a partition, considering navigation sessions of at least one click.

A. Model of Readers' Behavior

To comprehensively measure how much the network topology exposes users to diversity, we should consider both the graph topology and how readers navigate the network. Indeed, the exposure to diverse information might vary for users who behave differently in terms of navigation session length and next-link choices. So far, there are no models that generalize the navigation behavior of Wikipedia users. Thus, on top of previous findings [24], [25], in Sect. IV-A1 and IV-A2, we define a parametric model that simulates a wide range of users' navigation sessions by embedding different behaviors. We emphasize that the scope of this model is not to perfectly replicate users' behavior on Wikipedia. Rather, we want to see how users simulated from a reasonable and general model get exposed to diverse information.

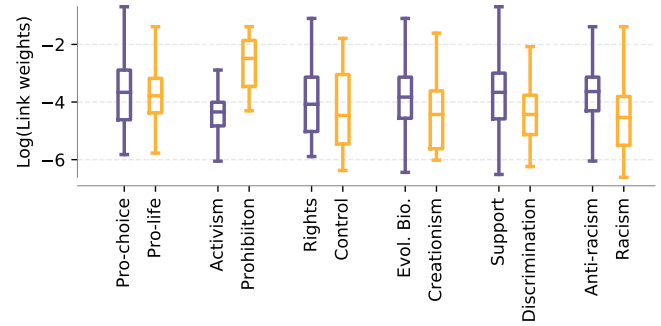


Fig. 4: Distributions of across-partition links weights. Topics are: *abortion, cannabis, guns, evolution, racism* and *LGBT*.

1) **Model Clicks Within Pages (CwP):** When readers visit a page, they have the possibility of clicking any link. However, according to the information needs they want to satisfy, each of the links may have a different click-probability [7]. We characterize the probability of “clicking a link j within an article i ” in three ways. First, let i be an article in V and $j \in N_{out}(i)$, where $N_{out}(i)$ is the set of pages to which i has a link. We define $pos(j|i)$ as the rank of j among all links in i , and $r(j|i) = |N_{out}(i)| - pos(j|i)$, such that a higher value indicates a higher ranking position. We consider links in the infoboxes as at the top of the article, according to results in [24], [25]. Moreover, we introduce $\tanh x = \frac{e^{2x}-1}{e^{2x}+1}$, which we use to transform ranking positions to values between 0 and 1, such that links at the top of the page are assigned similar scores. For instance if two links are adjacent in a line, likely their probability of being clicked is similar.

We embed *clicks within pages* models (CwP) into G by setting its transition matrix M in one of the following modes:

- 1) M^u (*Uniform*), whose entry $m(i, j) = \frac{1}{|N_{out}(i)|}$ mimics readers who click each link uniformly at random;
- 2) M^p (*Position*), whose entry $m(i, j) = \frac{\tanh r(j|i)}{\sum_{j \in N_{out}(i)} \tanh r(j|i)}$ captures readers who click with higher probability links appearing on the top of the page. This model is based on previous works showing that the links' position is a good predictor to determine its success [18], [42];
- 3) M^c (*Clicks*), whose entry $m(i, j) = \frac{c_{i,j}}{\sum_{j \in N_{out}(i)} c_{i,j}}$ is the observed probability that users in i will click the link toward j . The quantity $c_{i,j}$ counts how many times on average real users clicked the hyperlink from page i to j , from August 2019 to September 2020.⁷ For the links never clicked, we set $c_{ij} = 10$, the minimum number of times that the link must be clicked to be included in the dataset [44]. This smoothing factor allows one to assign a positive weight to links rarely clicked.

2) **Readers Navigation Model:** To characterize the users' sessions, we define a stochastic process with $|V|$ states, which,

⁷Wikipedia's clickstream data is publicly available and preserves users' privacy [15], [43]. Data description at Research:Wikipedia_clickstream.




Fig. 5: Navigation model for different α .

for each click, approximates the probability of reaching any of the articles starting at random from $p \in P$ (or from \bar{P}). We consider the process $\{X^\ell; \ell = 0, 1, \dots, L\}$, on the set of nodes V induced by the transition matrix M with starting state X^0 selected from the probability distribution $\pi_P^0 = \{(\pi_P)_i\} \in \mathbb{R}^{1 \times n}$ over V . Assuming that the user session length (the number of clicks) is finite, we evaluate the process on a finite number of steps L . Thus, $\Pr(X^\ell = j) = (\pi_P^\ell)_j$, where the (row) vector π_P^ℓ is given by the following variation of the Personalized Random Walk with Restart (RWR).

Definition 1 (Navigation Model): Let M_0 be the transition matrix embedding a click-within-pages model, π_P^0 the distribution of the starting state over P , and $\alpha \in [0, 1]$ the restart parameter. We have

$$\pi_P^1 = \pi_P^0 \cdot M_0 \quad (1)$$

and, for $\ell \geq 1$,

$$\pi_P^{\ell+1} = (1 - \alpha)\pi_P^\ell \cdot M_\ell + \alpha(\pi_P^0 \cdot M_\ell), \quad (2)$$

where $M_\ell = \text{norm}((D(M_{\ell-1})^T)^T)$ and $D = \text{diag}(\mathbf{1} + \pi_P^{\ell-1})^{-1}$. The operator $\text{norm}(M)$ transforms matrix M into a right-stochastic matrix by normalizing each row independently such that it sums to 1.

This process is a variation of the standard random-surfer model that differs for the update of the transition matrix at each step. The vector π_P^ℓ represents the likelihood that each node is reached at step ℓ if the session starts uniformly at random from a node in P . Assuming that readers do not click multiple times the same link within a session, we desire to deflate the probability of reaching nodes that, at step $\ell + 1$, have already been visited with high probability. We achieve this by dividing the rows of M by the vector of probabilities $\pi_P^{\ell+1}$, where 1 is a smoothing factor, and then normalize the matrix to obtain the updated stochastic matrix to use in the next iteration. Looking deeper into the model:

- For $\alpha = 0$ (Fig. 5(a)), the readers' clicks depend only on the CwP model. In this case, especially if related articles are not densely connected, the exploration can quickly lead to articles less related to the starting page.
- For $\alpha = 1$ (Fig. 5(b)), readers locally explore articles likely semantically related to each other [1], and the model emulates a *star-like* behavior, which consists in sequentially opening links from the starting page.
- For $0 < \alpha < 1$ (Fig. 5(c)), the readers' choices depend on the CwP model and, occasionally, they go back to the

initial page. The more α is close to 1 the more users show a star-like behavior. The closer α is to 0 the more users navigate in a more Depth First Search-oriented fashion. The model emulates (1) readers who sequentially explore articles and then jump back to the starting page, or (2) readers keeping open multiple paths.

Wikipedia does not have a button that allows readers to go back to the previous page. Thus, to *jump back* consists of clicking the browser's back button, until the session starting page. The restart parameter indirectly embeds the back button, which for the absence of *back-links* on Wikipedia does not appear in the graph.

B. Quantification of Exposure to Diverse Information

The *exposure to diverse information* aims to quantify how much the network structure allows readers to reach one, or multiple sets of articles, depending on their behavior. It is built upon both the CwP and the *navigation* models, and its application generalizes to arbitrary sets of nodes in a graph.

Definition 2 (Exposure to diverse information (ExDIN)): Given two sets of pages P, \bar{P} in V , let π_P^ℓ be the vector indicating the probability distribution of reaching any node in V at step ℓ ($\ell \geq 1$) starting from a random page in P . We say that the exposure of P to \bar{P} is

$$e_{P \rightarrow \bar{P}}^\ell = \sum_{j \in \bar{P}} \Pr(X^\ell = j) = \sum_{j \in \bar{P}} (\pi_P^\ell)_j \quad (3)$$

and it describes the probability that a reader in P reaches an arbitrary node in \bar{P} at the ℓ th click.

Definition 2 can be extended to multiple sets. Assume that we want to measure how much the set P is exposed to three sets of nodes, Q, Z , and L . The total exposure to the three sets is the ExDIN computed setting $\bar{P} = Q \cup Z \cup L$. Otherwise, if we want to have the ExDIN w.r.t. to each set, namely, $e_{P \rightarrow Q}, e_{P \rightarrow Z}$, and $e_{P \rightarrow L}$, we take π_P^ℓ and sum up the probabilities of the nodes within each set.

To quantify the extent to which the exposure to diverse information is balanced across P and \bar{P} , we introduce the *mutual exposure to diverse information*.

Definition 3 (Mutual exposure to diverse information (M-ExDIN)): Let $e_{P \rightarrow \bar{P}}^\ell$ and $e_{\bar{P} \rightarrow P}^\ell$ be the exposure to diverse information of sets P and \bar{P} . The mutual exposure between the sets is

$$\epsilon^\ell = \frac{\min\{e_{P \rightarrow \bar{P}}^\ell, e_{\bar{P} \rightarrow P}^\ell\}}{\max\{e_{P \rightarrow \bar{P}}^\ell, e_{\bar{P} \rightarrow P}^\ell\}} \in [0, 1]. \quad (4)$$

If either $e_{P \rightarrow \bar{P}}^\ell$ or $e_{\bar{P} \rightarrow P}^\ell$ is 0, then $\epsilon = 0$.

The closer ϵ is to 1, the more balanced are the probabilities of moving from one set to the other are. Thus, the network topology does not favor connections from one set to the other. Otherwise, the network structure tends to favor the navigation from one partition toward the other. With this view, if the network structure facilitates to move from one of the sets to the other, we may say that the network topology is *biased* toward a direction. Thus, M-ExDIN is a measure of the network's bias w.r.t. two sets of nodes, at each click of a session. If




Fig. 6: *Local exposure to diversity*. Each plot shows the adjusted-ExDIN (%) across partitions. On the main diagonal we have flows within the same partition, e.g., $P \rightarrow P$. The off-diagonal reports the probability of moving across sides, e.g., $P \rightarrow \bar{P}$. The y -axis indicates the source and the x -axis is the destination. To each row corresponds the exposure to diverse information computed for different CwP. Darker colors indicate higher probability of being in the corresponding square in one click. The values in the brackets are the 90% confidence intervals. Topics are: *abortion, cannabis, guns, evolution, racism, and LGBT*.

the sizes of sets P and \bar{P} are unbalanced, we obtain higher probabilities for large partitions. To take the sizes into account, we introduce a strategy to compute the adjusted-ExDIN. Given the two partitions P and \bar{P} , we set a sample size equal to $z = \min\{|P|, |\bar{P}|\}$. From the two sets we sample with replacement P'_i and \bar{P}'_i of size z , respectively from the initial partitions. Hence, we bootstrap $e_{P' \rightarrow \bar{P}'}$ estimating the value of the adjusted-ExDIN.

V. EXPOSURE TO DIVERSE VIEWPOINTS

A. Local Exposure to Diversity

The *local* exposure to diversity is the possibility of accessing articles about another branch of the topic within one click. We measure it using exposure to diverse information, setting $\ell = 1$, which describes the static connectivity among partitions accounting for users' choices and the network topology.

In Fig. 6, the local exposure to diversity on each topic-induced network shows the following:⁸ When considering only the graph topology (i.e., M^u): (1) The networks' topology facilitates users to remain in a *knowledge bubble* (i.e., same partition) hindering the exploration of the topic's diverse stances. For every topic the probability of visiting pages of the opposing partition is on average 12 times lower than staying in the initial one. (2) One of the two partitions induces higher chances of remaining within the same bubble. On average, among all topics but *cannabis*, one of the two partitions has 2 times more chances of keeping users within its articles.

After embedding users' past behavior (i.e., M^u): (3) The probability that readers keep visiting pages about the same topic click-after-click (i.e., remaining within a *knowledge bubble*) is higher than when users click links uniformly at random (i.e., M^u). Indeed, on average, users have 1.5 more

chances of moving within \mathcal{T} . The probability increases significantly, suggesting to start a discussion on the importance of exposing users to diverse information. (4) The likelihood of exploring diverse content slightly increases, showing that real users clicked pages of the opposing partitions more than what described by the uniform model. On average, users have 1.6 times more chances of moving to the opposing sides. (5) The discrepancy between the probability of moving within and outside the initial partition is on average even 2.6 times larger than when using M^u .

From past users' clicks we observe the preference of navigating across pages of the same partition. So far, users' behavior has always been justified by their information needs. From the observations (1) and (2), we know that the network's topology potentially favors the visit of pages of same standing. Is it possible that users' next-click choices are influenced by the hyperlinks network?

B. Dynamic Exposure to Diversity

We expand the analysis to sessions longer than one click to study the users' dynamic exposure to diversity. In Fig. 7, we observe that: (1) Within a few steps (4–5 clicks), users are more exposed to the pages of the same partition (i.e., the ratio of exposures is smaller than 0). (2) The current structure of the network favors users starting from any random page to reach one partition more easily than the other. Can we consider it a bias of the network? Is it fair that one side of a polarizing topic is less reachable than the opposite from any random node in the graph? (3) If users navigate according to a star-like random navigation model, the ratio between moving outside and within the partition stays steady or slightly increases. (4) All topic-induced networks are topologically biased; indeed none of them does the network provides an equal exposure across partitions (i.e., mutual exposure to diverse information is lower than 100%). (5) The local bias is smaller than the overall bias of the network. In general, the mutual exposure to diverse information decreases for longer sessions. (6) The

⁸We omit the plots showing result when using the model embedding the links' position. In general, there are not significant differences compared to the uniform model. For some topics, such as *guns*, the links' position plays a more significant role, worsening the user exposure to diverse information.

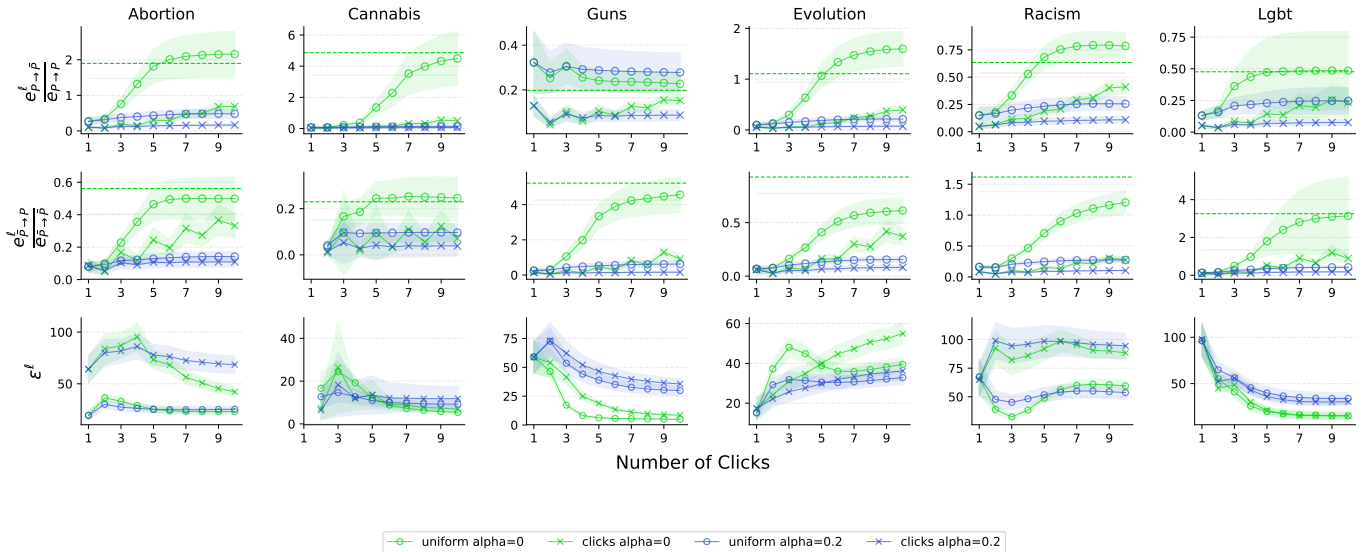


Fig. 7: (Adjusted) Dynamic Exposure to Diversity for sessions of 1 to 10 clicks. The x -axis indicates the number of clicks. In the first two rows, the y -axis is the ratio of the probability of moving from P to \bar{P} and P , and vice versa. Values lower than 1 indicate that remaining in the *knowledge bubble* is more likely than visiting pages of diverse content. Values close to zero quantify the chances of staying in a *knowledge bubble*. The y -axis of the third row indicates the mutual exposure to diversity (%). Values closer to 0 show that the current network topology might be biased toward the partition P or \bar{P} . Colors indicate the navigation model (set by α). Shapes encode the CwP model. The bands indicate the standard deviation. The dashed lines represent the convergence values for infinite sessions.

general topology of the graph makes pages related to *liberal* standings more accessible.⁹

The lack of mutual exposure might depend on many factors such as *underlinked* articles [1] or missing words to attach links. An in-depth investigation of this condition may be an interesting future work.

C. Factors Related to Exposure to Diversity

1) *ExDIN and homophily*: We measure the Pearson correlation between the homophily¹⁰ of article i and the users' exposure to diversity starting a session in i . Locally, starting from a page with high homophily decreases users' probabilities of being exposed to diverse content. Indeed, we observe a negative correlation for 1–3 clicks sessions (avg. -0.45), and an average correlation close to 0, for longer sessions.

2) *ExDIN and centrality*: We measure the Pearson correlation between the degree centrality of article i and the users' exposure to diversity starting a session in i . The number of incoming links of an initial page does not play a role in determining the exposure to diversity, especially within a few clicks (in-degree correlation 0.07). On the other hand, within 1–3 clicks, the lower out-degree mildly increases the

users exposure to diversity (average -0.25). Under the uniform model, it is because the links' transition probabilities of pages with a few links are larger than pages with a lot of links.

VI. DISCUSSIONS

In this work, we look for the first time at Wikipedia's hyperlink structure to measure its influence on users' exposure to diverse information. By employing two Wikipedia-tailored metrics, we quantify the likelihood of visiting pages representing different aspects of a topic throughout a navigation session. Our findings indicate that the current network topology often limits exposure to diverse information and incentivizes users to remain in *knowledge bubbles*.

The ultimate goal of this work is to draw attention and initiate a discussion about the importance of evaluating the hyperlink structure as part of Wikipedia's goal to provide a natural point of view presentation, even for polarizing subjects. Our observations raise a number of interesting questions for the Wikipedia community. As an example, consider a page about an *anti-abortion organization*. It seems natural that this page has more hyperlinks to pages related to anti-abortion subjects than to pages related to abortion rights. This is reasonable and aligns with the current purpose of Wikipedia's internal links, but is it still reasonable, and conforms with the goal of a natural point of view? Similarly, we observe that in the directed hyperlink graph, it is often more likely to reach an article about B starting from an article about A, than reaching an article about A starting with an article about B, when A and

⁹We omit the star-like model ($\alpha = 0$), because its value is steady around the value for $\ell = 1$, and we omit the position CwP model because it shows trends similar to uniform.

¹⁰We use the EI Homophily index [45]: $EI(v \in P) = \frac{|ext_P| - |int_P|}{|ext_P| + |int_P|}$, where ext_P is the set of edges from P to the rest of the network, and int_P is the number of edges pointing to P .

B represent two aspects of a topics. Again, some imbalance is reasonable, but it can keep users locked in an information bubble. How do we distinguish between the two cases?

We expect ours and future findings to motivate work on editors' support tools for contextualizing pages within their neighborhood in the hyperlink network, and suggest hyperlink modifications to improve access to diverse content.

ACKNOWLEDGEMENTS

The project is supported by DARPA LwLL program, by the ERC Advanced Grant 788893 AMDROMA, the EC H2020RIA project "SoBigData++" (871042), and the MIUR PRIN project ALGADIMAR. The views and conclusions contained herein are those of the authors and should not be interpreted as representing the official policies or endorsements, either expressed or implied, of DARPA or the U.S. Government.

REFERENCES

- [1] Wikipedia, "Linking," in *Wikipedia:Manual of style/linking*.
- [2] B. Keegan, D. Gergle, and N. Contractor, "Hot off the wiki: dynamics, practices, and structures in wikipedia's coverage of the tōhoku catastrophes," in *Proc. of the 7th international symposium on Wikis and open collaboration*, 2011.
- [3] A. Piscopo and E. Simperl, "What we talk about when we talk about wikidata quality: a literature survey," in *Proc. of the 15th International Symposium on Open Collaboration*, 2019.
- [4] Wikipedia, "Neutral point of view," in *Wikipedia:Neutral point of view*.
- [5] I. Beschastnikh, T. Kriplean, and D. W. McDonald, "Wikipedian self-governance in action: motivating the policy lens," in *ICWSM*, 2008.
- [6] A. Forte, V. Larco, and A. Bruckman, "Decentralization in wikipedia governance," *Journal of Management Information Systems*, 2009.
- [7] P. Singer, F. Lemmerich, R. West, L. Zia, E. Wulczyn, M. Strohmaier, and J. Leskovec, "Why we read wikipedia," in *Proc. of the 26th International Conference on World Wide Web*, 2017.
- [8] Wikipedia, "Statistics," in *Wikipedia:Statistics*.
- [9] M. H. Ribeiro, R. Ottoni, R. West, V. A. Almeida, and W. Meira Jr, "Auditing radicalization pathways on youtube," in *Proc. of FAccT 2020*.
- [10] M. Redi, B. Fetahu, J. Morgan, and D. Taraborelli, "Citation needed: A taxonomy and algorithmic assessment of wikipedia's verifiability," in *The World Wide Web Conference*, 2019.
- [11] B. Fetahu, K. Markert, W. Nejdl, and A. Anand, "Finding news citations for wikipedia," in *Proc. of the 25th ACM International on Conference on Information and Knowledge Management*, 2016.
- [12] T. Piccardi, M. Catasta, L. Zia, and R. West, "Structuring wikipedia articles with section recommendations," in *The 41st International ACM SIGIR Conference*, 2018.
- [13] S. Kumar, R. West, and J. Leskovec, "Disinformation on the web: Impact, characteristics, and detection of wikipedia hoaxes," in *Proc. of the 25th International Conference on World Wide Web*, 2016.
- [14] A. Paranjape, R. West, L. Zia, and J. Leskovec, "Improving website hyperlink structure using server logs," in *Proc. of the Ninth ACM International Conference on Web Search and Data Mining*, 2016.
- [15] E. Wulczyn, R. West, L. Zia, and J. Leskovec, "Growing wikipedia across languages via recommendation," in *Proc. of the 25th International Conference on World Wide Web*, 2016.
- [16] D. Helic, M. Strohmaier, M. Granitzer, and R. Scherer, "Models of human navigation in information networks based on decentralized search," in *Proc. of the 24th ACM conference on hypertext and social media*, 2013.
- [17] P. Gildersleve and T. Yasseri, "Inspiration, captivation, and misdirection: Emergent properties in networks of online navigation," in *International Workshop on Complex Networks*, 2018.
- [18] D. Lamprecht, K. Lerman, D. Helic, and M. Strohmaier, "How the structure of wikipedia articles influences user navigation," *New Review of Hypermedia and Multimedia*, 2017.
- [19] P. Singer, T. Niebler, M. Strohmaier, and A. Hotho, "Computing semantic relatedness from human navigational paths: A case study on wikipedia," in *International Journal on Semantic Web and Information Systems* 9, 2013.
- [20] R. West and J. Leskovec, "Human wayfinding in information networks," in *Proc. of the 21st international conference on World Wide Web*, 2012.
- [21] A. T. Scaria, R. M. Philip, R. West, and J. Leskovec, "The last click: Why users give up information network navigation," in *Proc. of the 7th ACM international conference on Web search and data mining*, 2014.
- [22] A. Dallmann, T. Niebler, F. Lemmerich, and A. Hotho, "Extracting semantics from random walks on wikipedia: Comparing learning and counting methods," in *Wiki@ICWSM*, 2016.
- [23] T. Koopmann, A. Dallmann, L. Hettinger, T. Niebler, and A. Hotho, "On the right track! analysing and predicting navigation success in wikipedia," in *Proc. of the 30th ACM Conference on Hypertext and Social Media*, 2019.
- [24] D. Dimitrov, P. Singer, F. Lemmerich, and M. Strohmaier, "Visual positions of links and clicks on wikipedia," in *Proc. of the 25th International Conference Companion on World Wide Web*, 2016.
- [25] —, "What makes a link successful on wikipedia?" in *Proc. of WWW 2017*.
- [26] D. Blei, L. Carin, and D. Dunson, "Probabilistic topic models," *IEEE signal processing magazine*, 2010.
- [27] F. Shi, M. Teplitskiy, E. Duede, and J. A. Evans, "The wisdom of polarized crowds," *Nature human behaviour*, 2019.
- [28] L. A. Adamic and N. Glance, "The political blogosphere and the 2004 us election: divided they blog," in *Proc. of the WWW-2005 Workshop on the Weblogging Ecosystem*, 2005.
- [29] A. Cossard, G. De Francisci Morales, K. Kalimeri, Y. Mejova, D. Paolotti, and M. Starnini, "Falling into the echo chamber: The Italian vaccination debate on Twitter," in *Proc. of the International AAAI Conference on Web and Social Media*, 2020.
- [30] M. D. Conover, J. Ratkiewicz, M. Francisco, B. Gonçalves, F. Menczer, and A. Flammini, "Political polarization on twitter," in *Fifth international AAAI conference on weblogs and social media*, 2011.
- [31] S. Flaxman, S. Goel, and J. M. Rao, "Filter bubbles, echo chambers, and online news consumption," *Public opinion quarterly*, 2016.
- [32] P. C. Guerra, W. Meira Jr, C. Cardie, and R. Kleinberg, "A measure of polarization on social media networks based on community boundaries," in *7th International AAAI Conference on Weblogs and Social Media*, 2013.
- [33] K. Garimella, G. D. F. Morales, A. Gionis, and M. Mathioudakis, "Quantifying controversy on social media," *ACM Transactions on Social Computing*, 2018.
- [34] S. Haddadan, C. Menghini, M. Riondato, and E. Ufpl, "Republik: Reducing polarized bubble radius with link insertions," in *Proc. of WSDM 2021*.
- [35] E. S. Callahan and S. C. Herring, "Cultural bias in wikipedia content on famous persons," *Journal of the American society for information science and technology*, 2011.
- [36] E. Graells-Garrido, M. Lalmas, and F. Menczer, "First women, second sex: Gender bias in wikipedia," in *Proc. of the 26th ACM Conference on Hypertext & Social Media*, 2015.
- [37] C. Wagner, E. Graells-Garrido, D. Garcia, and F. Menczer, "Women through the glass ceiling: gender asymmetries in wikipedia," *EPJ Data Science*, 2016.
- [38] Wikipedia, "Namespace," in *Wikipedia:Namespace*.
- [39] —, "Redirect," in *Wikipedia:Redirect*.
- [40] U. Brandes, P. Kenis, J. Lerner, and D. Van Raaij, "Network analysis of collaboration structure in wikipedia," in *Proc. of the 18th international conference on World wide web*, 2009.
- [41] D. Lizorkin, O. Medelyan, and M. Grineva, "Analysis of community structure in wikipedia," in *International conference on World wide web*, 2009.
- [42] D. Dimitrov, P. Singer, F. Lemmerich, and M. Strohmaier, "What makes a link successful on wikipedia?" in *Proc. of the 26th International Conference on World Wide Web*, 2017.
- [43] D. Dimitrov and F. Lemmerich, "Democracy and difference: Different topic, different traffic: How search and navigation interplay on wikipedia," *The Journal of Web Science* 6, 2019.
- [44] E. Wulczyn and D. Taraborelli, "Wikipedia clickstream," <https://doi.org/10.6084/m9.figshare.1305770.v22>, 2017.
- [45] D. Krackhardt and R. N. Stern, "Informal networks and organizational crises: An experimental simulation," *Social Psychology Quarterly*, 1988.