

Bridging Nations: Quantifying the Role of Multilinguals in Communication on Social Media

Julia Mendelsohn¹, Sayan Ghosh², David Jurgens¹, Ceren Budak¹

¹ University of Michigan, School of Information

² University of Southern California, Department of Computer Science

juliam@umich.edu, ghoshsay@usc.edu, jurgens@umich.edu, cbudak@umich.edu

Abstract

Social media enables the rapid spread of many kinds of information, from memes to social movements. However, little is known about how information crosses linguistic boundaries. We apply causal inference techniques on the European Twitter network to quantify multilingual users' structural role and communication influence in cross-lingual information exchange. Overall, multilinguals play an essential role; posting in multiple languages increases betweenness centrality by 13%, and having a multilingual network neighbor increases monolinguals' odds of sharing domains and hashtags from another language 16-fold and 4-fold, respectively. We further show that multilinguals have a greater impact on diffusing information less accessible to their monolingual compatriots, such as information from far-away countries and content about regional politics, nascent social movements, and job opportunities. By highlighting information exchange across borders, this work sheds light on a crucial component of how information and ideas spread around the world.

Introduction

Social media facilitates worldwide diffusion of many forms of content, such as pop culture memes, protest movements, and misinformation (Bruns, Highfield, and Burgess 2013; Nissenbaum and Shifman 2018; Bridgman et al. 2021). However, most connections on social media are formed between people who share a common nationality or language (Ugander et al. 2011; Takhteyev, Gruz, and Wellman 2012). How can information spread around the world when relatively few ties cross geographic and linguistic boundaries? Multilingual users are believed to be an important piece of this puzzle, but understanding how they act as brokers in information flow across language communities remains underexplored (Hong, Convertino, and Chi 2011). We carry out a set of causal inference studies to quantify how multilingual users influence cross-lingual information exchange across Europe on Twitter and show that the role of multilingual users varies depending on the relationship between countries and the topic of content shared.

We identify multilinguals based on the languages of their authored posts and avoid making claims about users' offline




Figure 1: Consider networks of users from pairs of countries (C_x, C_y) (e.g. Germany and Turkey) with dominant languages L_x and L_y (e.g. German and Turkish). Here, multilingual user A posts in both languages. We use causal inference to quantify (i.) the structural role and (ii.) the communication influence of A in cross-lingual information exchange.

language competencies, which is an ideologically and theoretically fraught notion (Cheng et al. 2021). Using this operationalization of multilingualism, prior work has found that multilingual users tend to have ties (e.g., following links) to multiple distinct language communities, suggesting that they act as bridges in online social networks (Eleta and Golbeck 2012; Hale 2014a). We expand upon earlier population-level descriptive work by using causal inference techniques, namely propensity score stratification, to more robustly isolate the effects of individual users' multilingual behaviors on several outcomes. This approach is motivated by recent work showing propensity score stratification aligns with experimental results when measuring peer effects in link sharing behavior on social media (Eckles and Bakshy 2020). We carry out two analyses using this framework. First, we measure the **structural role** of multilinguals. Here, we quantify the extent to which these users act as bridges using betweenness centrality (Freeman 1978). We then measure multilingual users' **communication influence**. To do so, we ask how having a multilingual contact impacts the content, specifically URL domains and hashtags, that monolingual users share from other languages.

By analyzing information spread across country pairs

with different dominant languages, we show how multilinguals' influence varies based on relationships between countries. Takhteyev, Gruz, and Wellman (2012) suggest that offline country relationships (e.g., economic agreements or migration patterns) impact transnational tie frequency online. Extending this conjecture, we hypothesize that the role of multilinguals is not uniform across all country pairs. Specifically, we expect multilinguals to have a bigger impact on country pairs that are more geographically distant or have weaker bonds, in which case these users would serve as gatekeepers of otherwise inaccessible information.

We additionally consider actual content, which affects the rate and shape of information diffusion (Romero, Meeder, and Kleinberg 2011; Tsur and Rappoport 2012). However, this cannot be captured by analyses of network structure alone. Our measures of multilinguals' communication influence overcome this limitation by enabling us to compare effects across content topics. We again hypothesize multilingual users play a larger role in spreading content otherwise inaccessible to an international audience. For example, we expect multilinguals to have a bigger influence on spreading hashtags about local or national politics compared to widespread entertainment sensations.

Our contributions are as follows: in large-scale causal inference studies of European Twitter, we show that multilingually-posting users play a vital structural role and communication influence on information diffusion across languages. We compare how multilinguals' influence varies based on relationships between countries and find they have a greater effect among country pairs that are more geographically distant, especially for Western Europeans who post in Eastern European languages. We further measure how the effect of multilinguals is driven by content topic and find that they have the largest influence in spreading content about politics, developing health-related social movements, and job opportunities.

Understanding how multilinguals affect information diffusion has immense consequences for online platforms. For example, platforms may focus efforts on empowering multilingual users to spread information that can support knowledge-sharing, collaboration, crisis response, social progress, or other beneficial outcomes (Eleta and Golbeck 2014). On the other hand, they may want to discourage multilinguals from sharing dangerous content such as misinformation or abuse.

Related Work

We draw upon prior work on multilinguals' online behavior and information diffusion across social networks.

Social Networks and Information Diffusion

Online social networks tend to have relatively few ties connecting distinct national and linguistic communities, leading to structural holes (Burt 2004; Ugander et al. 2011; Hale 2012). Spreading novel information across these communities thus relies on bridges that span structural holes (Easley and Kleinberg 2010). Information spreads the most quickly across long-range bridges, where the intermediary

node greatly reduces the shortest path between two other nodes (Kossinets, Kleinberg, and Watts 2008). Bridges in a social network play a brokering role that adds to their social capital (Burt 2004). We posit that multilinguals play an important role in cross-lingual information exchange because they serve as bridges between language communities.

Prior work measures how people influence others to propagate information (Guille et al. 2013). For example, people are more likely to share information their friends share, and overall, weak ties (e.g., acquaintances) are more responsible for spreading novel information than strong ties (e.g., close friends) (Bakshy et al. 2012). Although multilinguals are not necessarily weak ties, they similarly can act as bridges between communities. Methodologically, we draw from Eckles and Bakshy (2020), who use causal inference techniques, namely propensity score stratification, to measure peer influence on link-sharing behavior and show that their observational estimates are consistent with experimental results.

Understanding information diffusion and influential brokers impacts research across disciplines, including the development of activism and protest movements (González-Bailón et al. 2011; Park, Lim, and Park 2015; Lee and Murdie 2020), spread of misinformation (Bridgman et al. 2021), product adoption (Talukdar, Sudhir, and Ainslie 2002), and online abuse (Spörlein and Schlueter 2021).

The Bridging Role of Multilinguals

Prior work has analyzed connections within and across language communities online (Hale 2012, 2014b; Samoilenko et al. 2016). Hale (2014b) argues that multilingual editors on Wikipedia are important for sharing knowledge across language communities and facilitate access to more diverse knowledge; they are more active than their monolingual counterparts overall and often write the same article across multiple language versions. However, Kim et al. (2016) suggests that language remains a barrier because multilinguals are less engaged and less likely to edit complex content in a non-primary language. While online bloggers primarily link to content within the same language, cross-lingual links do exist, and some bloggers explicitly seek to connect distinct language communities (Zuckerman 2008; Hale 2012). Authors of blogs that bridge language communities often tend to be multilingual migrants or language learners (Herring et al. 2007). On Twitter, multilinguals disseminate information across different public spheres during events such as the Arab Spring (Bruns, Highfield, and Burgess 2013).

Earlier work has also examined the structural role of multilinguals. Kim et al. (2014) count edges between monolingual and multilingual "lingua" groups within multilingual regions on Twitter. They find that monolingual groups tend to be connected via multilingual groups, suggesting that multilingual users are bridges between different language communities. Hale (2014a) similarly argues that multilinguals collectively play an important bridging role in the Twitter network. When removing all multilingual nodes, the largest connected component becomes smaller and the number of small components increases, and these changes are significantly larger than if the same number of randomly-selected monolinguals were removed (Hale 2014a). Eleta

and Golbeck (2012, 2014) characterize multilingual users’ ego-networks as gatekeepers (language communities connected by only a few users) and language bridges (tightly-connected language groups). The authors suggest that gatekeepers are essential individuals for spreading information across linguistic, national, and cultural boundaries.

Other work has measured multilinguals’ participation in information cascades (resharing chains) on Twitter. Jin (2017) find that multilingual behaviors of the original poster and their followers are predictive of information cascades crossing languages. Chen et al. (2021) compare monolingual and multilingual users in cascade trees of COVID-19 information, develop measures that capture users’ bridging roles based on how much they propagate information, and find that multilinguals have a bigger bridging role than monolinguals in nearly two-thirds of information cascades.

We build upon prior work in several important ways. While structural analyses suggest that multilinguals are important for cross-lingual information exchange, our causal inference studies quantify the effect of multilingual behaviors on both structural and content-sharing outcomes while accounting for possible confounds such as posting frequency and overall popularity. We further extend the evidence that online multilingual behaviors and connections vary across countries by identifying systematic variation in the effect of multilinguality across country pairs (Kim et al. 2014).

Data

We construct an undirected network with Decahose data from 2012-2020, where edges are mutual “mentions” between users. We consider all pairs of European countries (C_x, C_y) with dominant languages L_x and L_y , and select countries based on geography and Eurovision participation, a marker of cultural affiliation. Using location inference (Compton, Jurgens, and Allen 2014), we extract network subsets for all (C_x, C_y) pairs where nodes are users from C_x or C_y . We use tweets from 2018-2020 for multilingualism and content measures due to limited text availability.

Definitions Each pair of countries (C_x, C_y) is a *multilingual country pair (MCP)* with dominant languages L_x and L_y . Users who post only in L_x are L_x *monolinguals* and users who post in both L_x and L_y are (L_x, L_y) *multilinguals*. Within each (C_x, C_y) network we separately measure the role of (L_x, L_y) multilinguals located in C_x and C_y , which we refer to as *loci*. In Figure 1, *(Germany, Turkey)* is an MCP, and we measure the role of *(German, Turkish)* multilinguals within each locus *Germany* and *Turkey*.

Language identification Twitter posts present a challenge to automated language identification (LangID) due to their short length, informal style, and lack of ground truth labels (Graham, Hale, and Gaffney 2014; Williams and Dagli 2017). We use Twitter’s built-in language detector because it is computationally efficient for a massive dataset, requires few additional resources, and is trained on in-domain data.¹

¹https://blog.twitter.com/engineering/en_us/a/2015/evaluating-language-identification-performance

We validate our decision with a comparison to 5 popular LangID packages: `langdetect`², `langid.py` (Lui and Baldwin 2012), and `CLD2`³ use probabilistic models, while `fastText` (Joulin et al. 2016) and `CLD3`⁴ use neural networks. As in prior LangID evaluations, we randomly sample 1K tweets from 32 countries written in that country’s dominant language, as labeled by Twitter (Graham, Hale, and Gaffney 2014; Lamprinidis et al. 2021). Like Graham, Hale, and Gaffney (2014), we calculate intercoder agreement between all pairs of models. Table S3 (Supplemental Material) shows that Twitter’s LangID has a high agreement with other models, even at higher rates than they agree with each other.

Like most LangID models, Twitter’s LangID has relatively low coverage of the world’s languages and may struggle to distinguish similar languages (Lui and Baldwin 2014; Williams and Dagli 2017). We mitigate these limitations by selecting multilinguals and MCPs so that our analyses are not overly sensitive to individual prediction errors.

Identifying multilingual users Following Eleta and Golbeck (2014), an individual uses language L if at least 10% of their posts containing original content are written in L (i.e., excluding retweets but including quote tweets and replies). A user is *monolingual* if one language passes the 10% threshold, and *multilingual* in the (C_x, C_y) network if both L_x and L_y pass this threshold. Following Kim et al. (2014), we collect language information for all users with at least five posts in the Decahose. We additionally exclude users if over 20% of their tweets are in unidentified languages because our calculations for language use may be less accurate. We set these thresholds so that estimates of users’ multilingualism are more reliable and robust to language prediction errors of individual tweets. We emphasize that our operationalization of multilingualism is based solely on users’ expression on Twitter; we do not make claims about language knowledge or offline behavior.

Selecting network subsets for analysis From an initial set of 50 European countries, we exclude 18 after three filtering steps. First, we remove six micro-states with area under 500km² because Compton, Jurgens, and Allen (2014)’s reported errors suggest that location inference is less reliable within such small areas. Second, we restrict our analysis to countries with a single official or dominant language (i.e., used by at least 70% of the population), removing eight more countries.⁵ This step is necessary for our problem formulation, which focuses on information spread across borders. The study of users within highly-multilingual countries is beyond the scope of this paper. Third, we calculate the distribution of tweets authored by users from each country and exclude four countries where the majority of tweets’ languages are unidentified or no tweets are identified as written in the country’s dominant language. See Supplemental Material Table S1 for the list of included and filtered countries.

²<https://github.com/shuyo/language-detection>

³<https://github.com/CLD2Owners/cld2>

⁴<https://github.com/google/cld3>

⁵Data for the geographic area and population-level language usage is from the CIA World Factbook.




Figure 2: General design for causal inference studies (Ho et al. 2007; Stuart 2010). We first estimate the propensity of each unit to receive treatment and use these scores to divide the sample into 25 strata. We then compare treated and untreated units within each stratum through a weighted regression to estimate the causal effect of the treatment.

We include an MCP (C_x, C_y) if: (1) C_x and C_y have different dominant languages, and (2) the (C_x, C_y) network is sufficiently large to rigorously estimate causal effects. Specifically, an MCP network is included if it has at least 100 L_x and L_y monolinguals, 20 (L_x, L_y) multilinguals, and 100 users with at least 1 (L_x, L_y) multilingual neighbor.⁶ See Supplemental Material Table S2 for included MCPs.

Descriptive statistics Our dataset contains information for 12.6M users from 32 countries. 8.0M users pass the thresholds for posting activity and language identifiability and are considered for analysis. Of these users, 1.1 million (13.6%) tweet in both their country’s dominant language and another European language. Norway has the largest percentage of bilingual users (46.9%), while the United Kingdom has the least (3.3%). 232 MCP network subsets were selected based on our inclusion criteria, covering 30 European countries and 7.7 million unique users (Georgia and Moldova were considered, but no MCP network containing these countries was included). MCP networks range in size from 4.7K nodes (*Latvia, Lithuania*) to 3.6M nodes (*United Kingdom, Turkey*), with a median size of 319K nodes.

Problem Formulation

Adopting the Neyman-Rubin framework of potential outcomes (Holland 1986), we isolate the effect of multilingualism by addressing the counterfactual: how would a multilingual user u ’s influence be different if u were monolingual? Each study’s details vary (Table 1), but all have the same idea: we define a set of *units* (users), some of whom receive a *treatment* indicating multilingual behavior, and we estimate an *outcome* variable related to influence on information exchange across languages.

Measures of Influence

Are multilinguals well-positioned in networks to spread information? In Study 1, we quantify their structural role by estimating the effect of multilingual posting on betweenness centrality. A measure of the proportion of shortest paths that must go through an intermediary node, betweenness centrality quantifies the extent to which the intermediary is an in-

⁶While we deem such thresholds necessary for precise causal effect estimation, we acknowledge the arbitrariness of these numbers as a limitation of our study.

	Structural Role	Communication Influence	
	Study 1	Study 2	Study 3
Units	C_x users posting in L_x	Monolinguals in L_x from C_x	Monolinguals in L_x from C_x
Treatment	Posting multilingually in (L_x, L_y)	Having ≥ 1 (L_x, L_y) multilingual neighbor	Having ≥ 1 (L_x, L_y) multilingual neighbor
Outcome	Betweenness centrality	Sharing domain from L_y	Sharing hashtag from L_y

Table 1: Overview of problem formulation for causal inference studies. Each study focuses on the effect of multilingualism on cross-lingual information exchange, but defines different samples, treatments, and outcomes.

formation broker (Freeman 1978). Studies 2 and 3 focus on multilinguals’ influence on content flows across languages. If an L_x monolingual has a multilingual neighbor in the (C_x, C_y) network, how much more likely are they to share L_y content? We consider two forms of content: URL domains and hashtags (Hong, Convertino, and Chi 2011).

Betweenness centrality (Study 1) For each (C_x, C_y) MCP, we consider a sample of users from locus C_x who post in L_x . The treatment is if a user is multilingual in (L_x, L_y) , and the outcome is the user’s betweenness centrality in the (C_x, C_y) network, scaled by 10^6 and log-transformed.

Domain sharing (Study 2) We first identify URL domains associated with each language. We filter out the 200 overall most-frequent domains. Like stop word removal (Jurafsky and Martin 2000), excluded domains are largely uninformative for measuring content sharing across languages (e.g. *instagram.com*, *unfollowspy.com*). To account for sampling variability, we exclude domains that appear fewer than ten times (Monroe, Colaresi, and Quinn 2008). We then consider “domains associated with L_x ” to be the 100 domains most overrepresented in L_x tweets (excluding retweets) relative to all European tweets based on the weighted log odds ratio with an informative Dirichlet prior (Monroe, Colaresi, and Quinn 2008). For each MCP (C_x, C_y) and locus C_x ,

the sample is the set of L_x monolinguals from C_x . A user u is treated if they have at least one (L_x, L_y) multilingual neighbor, and the outcome is whether any of u 's Decahose tweets contain at least one domain associated with L_y , the language u does *not* use. Note that we exclude retweets for language-based measures, such as identifying multilinguals and language-specific domains/hashtags, but include retweets for information-sharing measures because they are an important component of information diffusion on Twitter.

Hashtag sharing (Study 3) We similarly identify hashtags associated with each language. Unlike domains, language-associated hashtags can change rapidly because they may refer to short-term events, such as elections, protests, or upcoming TV shows. We thus separate the Decahose data into 60 14-day intervals, and use the log-odds ratio with informative Dirichlet priors to identify H_x^i , the set of 100 hashtags most associated with L_x in interval i , after again filtering out 200 most-frequent hashtags (e.g. *#blog*, *#radio*), those occurring fewer than ten times in i , and excluding retweets. A user u "shares" a hashtag from L_x if any of u 's Decahose tweets contain at least one hashtag $h \in H_x^i$ during i or the subsequent period $i + 1$ (including retweets). Resulting hashtags reflect entertainment, sports, politics, and everyday life; see Supplemental Material Table S4 for examples of language-specific domains and hashtags.

Causal Inference Setup

We use propensity score stratification to estimate the causal effects of multilingual treatment variables on information diffusion outcomes (Rosenbaum and Rubin 1983). Our procedure, shown in Figure 2, closely follows the guidelines provided by Stuart (2010) and MatchIt (Ho et al. 2007).

We first fit a logistic regression model to calculate propensity scores, which represent the probability of receiving treatment as a function of the specified covariates (Rosenbaum and Rubin 1983). Covariates are shared for all studies and capture aspects of users' Twitter behavior that may affect our outcomes (Stuart 2010). These include verified status, network degree, follower and following counts, account age, number of tweets in the Decahose sample, "favorites" count, and post rate, all log-scaled. Covariates for each user are based on their most recent tweet in our sample.

We then conduct propensity score stratification where units are separated into 25 strata based on propensity scores and verify sufficient balance for all covariates with absolute standardized mean difference less than 0.1 (Ho et al. 2007). Although five strata have commonly been used in practice, more strata can yield less biased estimates for larger samples sizes (Eckles and Bakshy 2020). Propensity score estimation, stratification, and balance checks are performed using the MatchIt R package (Ho et al. 2007).

We compare treated and untreated units within each stratum by fitting a regression of each outcome on treatment status, weighted by the matching weights. In particular, we estimate the average treatment effect on the treated (ATT); this estimates how the outcome among treated units is different than in the counterfactual scenario where they are not treated (Stuart 2010). For all studies, we include covariate

	Outcome		
	Betweenness Centrality (Study 1)	Domain Sharing (Study 2)	Hashtag Sharing (Study 3)
# Eligible MCPs	214	158	199
# Eligible Loci	317	205	284
% Loci w/ sig. pos ATT	46.37%	56.10%	50.00%
% Loci w/ no sig. ATT	51.42%	40.49%	46.48%
% Loci w/ sig. neg ATT	2.21%	3.41%	3.52%

Table 2: Summary of causal effect of posting multilingually (Study 1) or having a multilingual friend (Studies 2 and 3) for users in each country (locus) across all multilingual country pairs (MCPs). Despite considerable variation across loci, each treatment has positive effects on information diffusion-related outcomes far more than negative effects. We estimate are average treated effect on the treated (ATT), and significance is determined at $p < 0.05$ with robust standard error estimation.

adjustment to control for direct effects that pre-treatment covariates may have on the outcome.

Because each study defines units, treatment, and outcomes differently, the specific details of ATT estimation vary. In Study 1, we fit a linear regression after propensity score stratification to estimate the difference in (scaled and log-transformed) betweenness centrality. Here, multilingualism increases betweenness centrality if $ATT > 0$. In Studies 2 and 3, the outcomes are binary so we use logistic regression to estimate ATT as an odds ratio, a measure used in prior work to compare domain and hashtag sharing behaviors across languages (Hong, Convertino, and Chi 2011). For some treated user u who is monolingual in L_x , the ATT estimates the odds that u shares a domain (hashtag) associated with L_y divided by the odds that u would share a domain (hashtag) associated with L_y in the counterfactual scenario where u is not treated. $ATT > 1$ indicates that having a multilingual (L_x, L_y) contact increases an L_x monolingual's likelihood of sharing domains (hashtags) from L_y .

In each study, we estimate the ATT on aggregate data from both loci in all MCP networks in order to get a single causal estimate of the role of multilinguals in cross-lingual information exchange. We additionally estimate separate ATT values for units from each locus in each MCP network, and analyze the variation across locus-specific causal effects later. To ensure sufficient data across strata, we only estimate locus-specific causal effects for locus C_x in MCP (C_x, C_y) if there are at least 100 treated units. Furthermore, we only include locus-specific causal effects for Studies 2 and 3 if at least 100 units in the combined treated and untreated sample (all L_x monolinguals) have an outcome of 1 (share at least one domain or hashtag from L_y).

Overall Effects of Multilingual Behavior

All 3 studies support our hypothesis that multilinguals play an important role in cross-lingual information exchange.

Betweenness centrality (Study 1) Multilingual (L_x, L_y) users have higher betweenness centrality in the (C_x, C_y) network than their monolingual peers, suggesting that these users serve as local bridges and thus are well-positioned to spread novel information across the network (Freeman 1978; Granovetter 1973). The overall ATT is 0.034 ($p < 0.0001$, with robust standard error estimation). In other words, multilingual posting increases the outcome of log-transformed betweenness by 0.034, which corresponds to a 13.5% increase in betweenness centrality.

URL domain sharing (Study 2) Having a multilingual (L_x, L_y) friend increases the odds of a monolingual L_x user sharing a domain from L_y by a factor of 15.57 ($p < 0.0001$). For interpretability, we also estimate the average marginal effect with a marginal effects model on the logistic regression used for estimating the ATT as an odds ratio. We find that having a (L_x, L_y) friend increases an L_x monolingual's probability of sharing an L_y domain by 20.0%.

Hashtag sharing (Study 3) Having a multilingual (L_x, L_y) friend significantly increases the odds of an L_x monolingual sharing a hashtag from L_y by a factor of 3.98 ($p < 0.0001$). Through estimating the average marginal effect, we find that this corresponds to an increase in the probability of sharing an L_y hashtag by 32.7%. While the odds ratio for hashtag sharing is lower than for domains, the probability increase is greater because extremely low domain-sharing probabilities result in inflated odds ratios.

Table 2 summarizes locus-specific causal effect estimates. Due to minimum inclusion criteria, we only estimate ATTs for a subset of all 464 loci corresponding to 232 MCPs. All three treatments increase information-sharing outcomes in about half of the loci, and decreases outcomes in under 4% percent of loci. Locus-specific estimates reinforce that multilinguals facilitate information exchange across language boundaries. However, the distribution of positive, negative, and insignificant effects across loci as shown in Table 2 suggests wide variation across MCPs and loci.

Variation across Country Pairs

All three studies show that multilingual behaviors increase cross-lingual information exchange. However, locus-specific causal estimates indicate substantial heterogeneity in the magnitude of their effects across MCPs and loci. We conduct regression analyses to characterize how geographic, demographic, economic, political, and linguistic relationships between country pairs correlate with the effects of multilingual behaviors on European Twitter.

Regression Setup

We fit linear regression models to characterize how the causal effects of multilinguals vary across countries. For a given MCP (C_x, C_y) and locus C_x , we define 3 dependent variables: the estimated causal effects (ATT) of multilingual behavior in each of the 3 described studies. Independent variables capture relationships between C_x and C_y .

Demographic variables include the population ratio of C_x to C_y from 2019 based on CEPII's Gravity Dataset

(Head and Mayer 2014; Conte, Cotterlaz, and Mayer 2021). Using 2017 World Bank migration data, we consider the fraction of C_x 's population born in C_y , C_y 's population born in C_x , and each country's population who are foreign-born.⁷

Geographic variables are the distance (in km) between C_x and C_y 's population centers and time difference: the number of hours C_y is ahead (further east) of C_x .

Economic variables include i.) the ratio of C_x to C_y 's GDP per capita, ii.) if C_x and C_y are in an RTA (Regional Trade Agreement, which includes the EU), and iii.) trade flow between C_x and C_y averaged over both country's reports and both directions, normalized by the total population of both countries. Geographic and economic variables use 2019 data from CEPII's Gravity Dataset.

Political variables, specifically material conflict between C_x and C_y , are determined by querying the GDELT event database (Leetaru and Schrodt 2013). These include the percent of C_x 's external conflict actions inflicted on C_y , and C_y 's external conflict actions inflicted on C_x .

The last fixed effect is **linguistic distance** between L_x and L_y using Glottolog (Hammarström et al. 2021). Inspired by Samoilenko et al. (2016)'s measurement of shared language families, we define a 4-level measurement of linguistic distance between L_x and L_y : i.) no relationship (e.g. Spanish and Hungarian), ii.) in the same primary family (e.g. German and Polish are Indo-European), iii.) in the same branch (e.g. English and Swedish are Germanic), and iv.) in the same sub-branch (e.g. Spanish and Italian are Romance).⁸

We avoid multicollinearity issues by ensuring that all variables' variance inflation factor is under 4. We thus exclude highly-correlated variables, such as EU membership and each country's population. We weight each regression model by the number of treated units from each locus. Finally, we scale all variables by z-score to facilitate direct comparisons.

Regression Results

Geography Multilinguals have a larger effect on cross-lingual information exchange in MCPs where C_x and C_y are further away from each other (Table 3). We visualize this pattern in Figure 3, which shows the relationship between geography and effects of multilinguals on betweenness centrality and domain sharing (see Figure S4 in Supplemental Material for the hashtag sharing map).

In Figure 3, the effect of multilinguals in MCP (C_x, C_y) and locus C_x is shown as a directed edge from C_y to C_x . Only edges corresponding to significant estimates are drawn. Negative effects are red, positive effects are blue, and greater magnitude is represented with darker and thicker edges. For example, the dark blue edge from Russia to Spain in both maps indicates that Russian-Spanish multilinguals are especially important for bringing Russian information to Spain. In contrast, the faint edge from Portugal to Spain

⁷ worldbank.org/en/topic/migrationremittancesdiasporaissues/brief/migration-remittances-data

⁸ We do not use graph-based measurements of distance because there is wide variation across language branches' structures due to an uneven interest by linguists across languages.




Figure 3: Maps of causal effects (ATTs) of multilingual treatments on information diffusion outcomes. The left map shows the ATTs of multilingual posting on betweenness centrality (Study 1), where each edge indicates ATT among users from the *destination* node who post in both the source and destination nodes' dominant languages. The right map shows the ATTs of having a multilingual friend on cross-lingual domain sharing (Study 2), where each edge indicates ATT among monolingual users from the *destination* node on sharing URL domains associated with the *source* node's dominant language. Negative effects are shown with red arrows and positive are shown with blue arrows. The magnitude of the causal effect is indicated by arrow shading and width. Only statistically significant estimates ($p < 0.05$) with robust standard error estimation are shown.

	Dependent variable:		
	Betweenness Centrality	Domain Sharing	Hashtag Sharing
Geographic distance	0.020***	0.434	0.412***
Time difference	0.015**	0.815***	0.001
Pop. C_x / C_y	-0.030***	-0.031	0.017
% C_y foreign-born	0.021**	-1.100**	-0.058
% C_y pop. born in C_x	0.017**	-0.044	-0.043
% C_x foreign-born	-0.002	-0.031	0.064
% C_x pop. born in C_y	0.007	-0.197	-0.040
GDP per capita C_x / C_y	0.038***	0.318	1.171***
RTA	0.010	-0.411	0.216*
Trade flow per capita	-0.013*	0.425	0.109
% C_x 's conflicts vs. C_y	0.019**	-0.159	-0.032
% C_y 's conflicts vs. C_x	0.002	-0.257	0.088
Linguistic distance	-0.027***	0.449	0.146*
Observations	317	205	284
R ²	0.266	0.193	0.448
Note:	*p<0.1; **p<0.05; ***p<0.01		

Table 3: Geographic, demographic, economic, political, and linguistic aspects of relations between countries C_x and C_y and their associations with the role of multilinguals in spreading information from C_y 's dominant language to C_x . Coefficients are from weighted linear regression models where the dependent variables are the locus-specific causal effects (ATTs) of multilingual behaviors in Studies 1-3.

means that Portuguese-Spanish multilinguals have a smaller role in importing Portuguese information to Spain.

We believe stronger treatment effects across longer distances are due to information accessibility. Information from faraway places is not readily accessible for monolinguals, so they may need to rely more on their multilingual friends to serve as information brokers. On the other hand, information between nearby countries such as Norway and Sweden may be more easily accessible with more channels for diffusion, possibly via more multilinguals, so users rely less on individual multilinguals to spread information.

In addition to distance, the time difference between C_x and C_y significantly predicts multilinguals' effects (Table 3). Multilinguals have the largest impact when C_y is 2-3 hours ahead (i.e., further east) of C_x (see Figure S1 in Supplemental Material). In other words, multilinguals have a large impact on spreading information from Eastern European languages to Western Europe. This asymmetric East-West pattern is visible in both maps of Figure 3 (e.g. the many dark blue outgoing links from Russia indicate a large influence of multilingual users from other countries who post in Russian).

Why are Western European users of Eastern European languages so influential in cross-lingual information exchange? Fully answering this question is left for future work, but we speculate that it is a consequence of historical inaccessibility to information across strict Cold War-era borders. Additionally, offline connections may explain these users' online role as bridges; Eastern migrants in Western Europe have been characterized by more transnational and circular offline networks since the early 2000s (Favell 2008).

Demographics Multilingualism increases betweenness centrality more for users in smaller countries who post in more populous countries' languages, but this relationship is

not significant for content sharing outcomes. While multilinguals in smaller countries are better positioned within MCP networks to spread novel information, they do not necessarily “import” information from the larger to smaller country.

When C_y has a higher proportion of foreign-born residents, (L_x, L_y) multilingualism has a greater impact on users’ structural role in C_x , but not communication influence. In fact, having a multilingual friend has a negative impact on sharing a domain from L_y when C_y has more foreign-born residents; perhaps this is because shared links from L_y feature more content with multicultural or international appeal, so people rely less on multilinguals as information brokers. We observe no significant effects of the foreign-born population of C_x on either the structural role or communication influence of multilinguals.

Economics and politics Multilinguals’ effects are associated with GDP per capita: multilinguals in C_x have more influence when C_x is wealthier than C_y (although this is not significant for domain sharing). However, government-level relationships including national trade agreements, trade flows, and political conflict are not significant predictors.

Linguistic similarity We observe mixed results for linguistic distance. Our information accessibility hypothesis suggests (L_x, L_y) multilinguals would play a larger role when L_x and L_y are very different because their monolingual peers would rely more heavily on them. This pattern appears to hold for communication influence outcomes but is not significant. Contrary to our hypothesis, multilinguals’ structural role is amplified when L_x and L_y are closely related. Though we do not have a clear explanation for this pattern, it may be driven by deeper differences in the structures of MCP networks connecting countries with similar dominant languages (Herring et al. 2007).

Variation across Topics

Different aspects of content, such as topic, framing, sentiment, and subjectivity, could amplify or hinder multilinguals’ influence. Therefore, here we extend Study 3 to investigate how the role of multilinguals depends on the topic of shared hashtags using unsupervised topic modeling (Kim et al. 2014; Jin 2017). Specifically, for each MCP (C_x, C_y) we measure how having a multilingual (L_x, L_y) friend affects the odds of an L_x monolingual from C_x sharing a *topic-specific hashtag* from C_y .

Measuring Multilinguals’ Role across Topics

Identifying topic-specific hashtags We train a multilingual contextualized topic model (CTM) to identify topics (Bianchi et al. 2021). This approach uses multilingual sentence-BERT (Reimers and Gurevych 2019) as input to a variational autoencoder topic model to support zero-shot topic prediction for texts in unseen languages during training (Srivastava and Sutton 2017; Bianchi et al. 2021). Crucially, the CTM uses the same topics in all languages, making direct comparisons across languages possible. The CTM is trained for 20 epochs on a random sample of 1M English-language tweets from the European decahose data. We set

Description	Example Hashtags
TV Shows	<i>skamitalia, masterchefgr, thearchers</i>
Fandoms	<i>jungkook, saveshadowhunters</i>
Art	<i>painting, vintage, architecture, arte</i>
Romance TV	<i>loveisland, poweroflovegr, liebesgeschichten</i>
TV Promo	<i>comingsoon, luciferonnetflix, skyupnext</i>
Job Promo	<i>career, hiring, jobs, startup, jobsearch</i>
Giveaways	<i>giveaway, freebiefriday, sorteo, winwin</i>
Music Promo	<i>radio, youtube, hits, newmusic, live</i>
Government, News	<i>bbcnews, afd, noafd, labour, orban</i>
Covid, Crisis, Tech	<i>covid19france, koronawirus, tech, gdpr</i>
Int’l Politics	<i>syria, venezuela, eeuu, isis, migranti</i>
Health	<i>mentalhealth, autism, clapforhns</i>
Requests	<i>stop, shoplocal, stopgiletsjaunes, helpme</i>
Equality	<i>metoo, weltfrauentag, racisme, lgbt</i>
Sports	<i>arsenal, halamadrid, futbol, fcporto</i>

Table 4: Fifteen topics of interest with example hashtags

the total number of topics to 50 and retain default values for all other hyperparameters. We use the trained CTM to predict topic distributions for all tweets containing any hashtag highly associated with any language in any time interval. Tweets are assigned to the single topic with the highest probability, and hashtags are then assigned to single topics based on the most frequent topic of tweets in which it appears.

We then manually inspect the ten most-frequent hashtags per topic to identify 15 topics of interest, where hashtags are coherent, meaningful, and reflect different types of content. These 15 topics account for 52.1% of European tweets from our dataset, and 66.8% of unique hashtags associated with any language at any time, which rises to 82.5% when accounting for hashtag frequency (see topic distributions in Figure S3 in Supplemental Material). Brief topic descriptions are in Table 4. To understand how content shapes multilinguals’ influence, we separate the 15 topics into four macro-categories: entertainment, politics, sports, and promotion (e.g., giveaways). The distribution of hashtag macro-categories across languages are shown in Figure S2 in the Supplemental Material.

Evaluation We evaluated our multilingual hashtag topic assignment method with intrusion tests, following common practice (Hoyle et al. 2021). Four hashtags were sampled from each topic t with frequency weighting, and one hashtag was sampled from one of the 49 other topics based on frequency. Annotators tried to identify which hashtag does *not* belong to t . We evaluated the 15 topics of interest in English, German, Spanish, Italian, and Turkish, selected based on annotators’ language proficiencies. We conducted ten intruder tests for each topic in English, and 5 for each topic in the other languages. Three annotators completed the intruder test for English hashtags with an average accuracy of 0.73 and interannotator agreement of 0.67 (Krippendorff α). Among tasks where at least two annotators selected the same intruder, accuracy rose to 0.78, and further rose to 0.89 for tasks where all three annotators agreed.

Estimating causal effects We extend Study 3’s design to compare multilinguals’ effects across topics. For a given

Language	English	Spanish	German	Turkish	Italian
Accuracy	0.731	0.747	0.813	0.587	0.627

Table 5: Annotator accuracy on topic intrusion tests, averaged over 3 annotators for English.

MCP (C_x, C_y) and locus C_x , units are L_x monolinguals from C_x and the treatment is having at least one multilingual (L_x, L_y) contact. For each topic t , the outcome is whether a user shares a hashtag that belongs to topic t and is associated with L_y . Our analysis focuses on overall causal effect estimates for all 15 topics. We also estimate locus-specific effects for each macro-category, with a summary of results and maps in the Supplemental Material (Figure S4).




Figure 4: Causal effect (ATT) of having a multilingual friend on cross-lingual hashtag sharing across topics. All estimates have heteroskedasticity-consistent standard errors below 0.03 and are significant ($p < 0.0001$). The x-axis shows log odds ratios, with 0 being no effect of having a multilingual friend, and the dashed line at 1.38 shows the overall effect on hashtag sharing. Colors represent macro-categories.

Results

Figure 4 shows the effect of multilinguals by hashtag topic and macro-category. Three key findings support our hypothesis that multilinguals play a greater role in spreading information that is otherwise less accessible to their peers.

First, multilinguals have a greater communication influence on the cross-lingual diffusion of political content than entertainment. Political discourse likely occurs more in regional or country-centric public spheres (Schünemann 2020), so there is a greater reliance on multilingual individuals to broker political information across borders. Correspondingly, of all political topics, multilinguals play a smaller role for topics with widespread transnational popularity and awareness, such as: *Equality*, which includes hashtags for global gender equality movements like #metoo and #8m, and *International Politics*, which includes country names and non-European political organizations. Contrary to most political hashtags, entertainment hashtags of-

ten reflect globally-popular phenomena (e.g., K-pop fandoms) despite being associated with specific languages, and more cross-lingual information cascades involve entertainment content (Jin 2017). We believe that individual multilingual contacts play a smaller role in entertainment because there are more ways for that content to spread.

Second, we argue that multilingual individuals are crucial for nascent social movements to gain global traction, but have a relatively small influence within well-established transnational movements. Figure 4 indicates that multilinguals have a large impact on spreading health-related information, which includes hashtags advocating for mental health awareness, COVID-related activism, and disability rights. In contrast to racial and gender equality movements reflected in the *Equality* topic, organizing for disability rights gained traction as a social movement later than for race and gender both offline and on Twitter, where the public sphere about disability is still growing (Scotch 1988; Sarkar et al. 2021). The initial adoption of burgeoning social movements across countries relies heavily on direct contacts, such as multilingual friends (McAdam and Rucht 1993). Information diffusion about more well-established social movements do not depend on multilingual bridges because it occurs via many channels of communication, including news and television media (McAdam and Rucht 1993).

Third, multilinguals play an especially important role in sharing information about job searches and career opportunities. This parallels Granovetter (1973)’s argument about the strength of weak ties for job-seeking purposes. Even though monolingual users’ ties with multilinguals are not necessarily weak, multilinguals similarly serve as bridges between different parts of a social network and thus facilitate access to novel information, such as job opportunities that users may not have otherwise been aware of.

Discussion

Gaining a complete picture of global information diffusion requires understanding how information crosses languages. We design three studies to investigate how multilingual users participate in this process, which we use to quantify their structural role and communication influence in information exchange across European languages on Twitter. For each pair of countries with different dominant languages, we construct networks where users are connected if they have mutually “mentioned” each other. We use these networks in Study 1 to quantify the extent to which multilinguals’ positions in these networks facilitate spreading novel information, as measured by betweenness centrality. In Studies 2 and 3, we quantify how having a multilingual contact influences the odds of monolingual users posting content from the other country’s dominant language.

Results from all three studies show that multilinguals play an outside role in cross-lingual information exchange compared to their monolingual peers. Effects vary widely but systematically across country pairs and topics. We conduct regression analyses to measure how the role of multilinguals is associated with demographic, geographic, economic, political, and linguistic aspects of the relationship between country pairs. To compare multilinguals’ influence across

topics, we augment our study design for hashtag sharing with multilingual contextualized topic modeling.

In general, multilingual individuals have a greater influence on the spread of information that is otherwise *less* accessible to their monolingual peers, as they play more of a gatekeeping role. Multilinguals have a greater effect on information diffusion between dominant languages of countries that are geographically far apart, with Western European multilinguals who post in Eastern European languages having an especially big influence. We identify a similar pattern for topics, where multilinguals have greater influence on cross-lingual information exchange for topics discussed in more restricted public spheres: national or regionally-oriented politics over entertainment which can have international appeal, nascent health-related social movements over established racial and gender equality movements, and job opportunities previously known only to small communities.

We acknowledge that this work has important limitations. First, our studies do not account for multilinguals who use minority languages (e.g., Basque) or reside in highly multilingual countries (e.g., Switzerland). Imperfect performance of location inference and language detection also limited the set of countries and languages studied. Furthermore, we make the simplifying assumption that tweets are written in one language, which does not adequately account for code-mixing within posts and the users who engage in such practices (Androutsopoulos 2015). While code-mixed tweets are a relatively small percent of Twitter communication (Rijhwani et al. 2017), accurately recognizing these tweets at scale has proven challenging due to the absence of labeled data for training models (Jurgens, Tsvetkov, and Jurafsky 2017). As the performance and efficiency of language detection of code-mixed tweets improve, we anticipate that incorporating such information would be fruitful, and analyzing the relationship between code-mixing strategies and information diffusion could yield interesting theoretical insights.

To avoid making assumptions about people's offline language usage or competence, we intentionally define multilinguals based on their performance on Twitter. However, this presents a significant limitation: users who only tweet in one language but understand multiple may also play an important, and perhaps different, role in information diffusion. Future research could employ different methodologies to highlight these users, such as linking social media activity with survey data about users' language backgrounds.

Further research can also improve upon our study designs. We adopt a traditional causal inference setup, which considers treatment status binary to emulate randomized experiments. Thus, all of our studies involve collapsing underlying continuous variables into binary indicators. A possible next step would involve adapting our studies to account for continuous treatments; this would facilitate investigation of how cross-lingual information exchange is impacted by a user's degree of multilingualism (Study 1) or the number and/or strength of a user's ties to multilinguals (Studies 2 and 3).

Beyond addressing these limitations, there are numerous avenues for future work. For example, we adopt a microscopic perspective on information diffusion by examining how multilingualism impacts individual users' roles in in-

formation diffusion; we focus on local influence because it is more precise and less random than observations of information cascades (Bakshy et al. 2011). Nevertheless, an interesting extension that takes a macroscopic perspective, perhaps involving simulations of cross-lingual information cascades (Chen et al. 2021), could help contextualize how these individual-level effects, in aggregate, shape the global flow of information. Another future direction would involve considering other forms of information that may spread via different mechanisms than URLs and hashtags, such as meme templates, images, videos, and text outside of hashtags. There is likely variation in the role of multilinguals across semantic dimensions beyond topic, such as emotional valence or misinformation. Finally, future research can assess the generalizability of our findings beyond the scope of European Twitter by applying our methodology to study other regions, languages, and platforms.

Broader Impact

Understanding the role of multilinguals in information diffusion has immense consequences. Platforms like Twitter can empower multilinguals to spread information that supports positive outcomes such as knowledge-sharing, collaboration, crisis response, or social progress, thus enabling different language communities to benefit from a truly global social network (Eleta and Golbeck 2014; Hale 2014a). Our study not only highlights this potential but also identifies how it varies across topics and with respect to the geographical, linguistic, and political relationship between the countries. For instance, our research suggests that multilinguals can be better utilized to spread political news as opposed to entertainment. We also see that their importance is more pronounced for supporting information spread across countries further away from each other. Such findings not only highlight the contexts where multilinguals already play an important role but also help us identify the barriers for cross-lingual diffusion; in such situations, platforms may benefit more from technologies such as machine translation.

Our research can also help platforms address dangerous consequences of global networks by focusing efforts on nudging multilinguals to mitigate the spread of harmful information such as misinformation, conspiracy theories, or online abuse. Past network science research shows the value of betweenness centrality in identifying nodes that can limit the spread of such information (Golovchenko, Hartmann, and Adler-Nissen 2018). Here, we show that multilingual nodes tend to have high betweenness centrality. Furthermore, our study shows that multilinguals play a particularly important role in the spread of political topics, common targets for malicious actors aiming to spread propaganda and disinformation.

Despite the potential for positive impact, we acknowledge the ethical risks of this work. Rather than stem the flow of harmful content, our work may inspire malicious agents to target and manipulate multilinguals into propagating such information. In addition, our focus on users of politically and socially-dominant language varieties and use of automated language detection excludes people whose posts contain endangered or minority languages, non-prestige dialects

of dominant languages, or code-mixing. Although our work does not present direct harm to individuals, these decisions systematically exclude marginalized groups whose online behavior deserves equal consideration.

To promote transparency and future research, we publicly share data, code, and models but take steps to preserve user privacy. The datasheets shared for causal effect estimation include only variables necessary to replicate our results. We do not share user IDs, raw text, or other personally-identifiable information. While the location inference tool used presents a privacy risk by inferring users' specific geo-coordinates, we only store information at the country level.

Conclusion

By developing a set of causal inference studies that measure users' structural role and communication influence, we show that multilingually-posting users on European Twitter are particularly important for information diffusion across languages. These users have an especially large influence in situations where they serve more as gatekeepers in information flow, particularly in spreading information from places and topics that are otherwise inaccessible to their monolingual peers. This work is crucial for understanding how information is shared around the world, and has implications for platforms to support beneficial consequences of global social networks while mitigating potential harms. Publicly-available code, models, and aggregated data can be found at: <https://github.com/juliamendelsohn/bridging-nations>.

Acknowledgments

This research was supported by the National Science Foundation (Grants IIS-1815875 and IIS-2007251) and through funding from the Volkswagen Foundation.

References

- Androutsopoulos, J. 2015. Networked multilingualism: Some language practices on Facebook and their implications. *International Journal of Bilingualism*, 19(2): 185–205.
- Bakshy, E.; Hofman, J. M.; Mason, W. A.; and Watts, D. J. 2011. Everyone's an influencer: quantifying influence on twitter. In *Proceedings of the fourth ACM international conference on Web search and data mining*, 65–74.
- Bakshy, E.; Rosenn, I.; Marlow, C.; and Adamic, L. 2012. The role of social networks in information diffusion. In *WWW*, 519–528.
- Bianchi, F.; Terragni, S.; Hovy, D.; Nozza, D.; and Fersini, E. 2021. Cross-lingual Contextualized Topic Models with Zero-shot Learning. In *EACL*, 1676–1683.
- Bridgman, A.; Merkley, E.; Zhilin, O.; Loewen, P. J.; Owen, T.; and Ruths, D. 2021. Infodemic pathways: Evaluating the role that traditional and social media play in cross-national information transfer. *Frontiers in Political Science*, 3: 20.
- Bruns, A.; Highfield, T.; and Burgess, J. 2013. The Arab Spring and social media audiences: English and Arabic Twitter users and their networks. *American behavioral scientist*, 57(7): 871–898.
- Burt, R. S. 2004. Structural holes and good ideas. *American journal of sociology*, 110(2): 349–399.
- Chen, N.; Chen, X.; Zhong, Z.; and Pang, J. 2021. The Burden of Being a Bridge: Understanding the Role of Multilingual Users during the COVID-19 Pandemic. *arXiv preprint arXiv:2104.04331*.
- Cheng, L. S.; Burgess, D.; Vernooij, N.; Solís-Barroso, C.; McDermott, A.; and Namboodiripad, S. 2021. The problematic concept of native speaker in psycholinguistics: Replacing vague and harmful terminology with inclusive and accurate measures. *Frontiers in Psychology*, 3980.
- Compton, R.; Jurgens, D.; and Allen, D. 2014. Geotagging one hundred million twitter accounts with total variation minimization. In *2014 IEEE international conference on big data (big data)*, 393–401. IEEE.
- Conte, M.; Cotterlaz, P.; and Mayer, T. 2021. The CEPII Gravity Database. *CEPII Working Paper*, 1–53.
- Easley, D.; and Kleinberg, J. 2010. *Networks, crowds, and markets*, volume 8. Cambridge university press Cambridge.
- Eckles, D.; and Bakshy, E. 2020. Bias and High-Dimensional Adjustment in Observational Studies of Peer Effects. *Journal of the American Statistical Association*.
- Eleta, I.; and Golbeck, J. 2012. Bridging languages in social networks: How multilingual users of Twitter connect language communities? *Proceedings of the American Society for Information Science and Technology*, 49(1): 1–4.
- Eleta, I.; and Golbeck, J. 2014. Multilingual use of Twitter: Social networks at the language frontier. *Computers in Human Behavior*, 41: 424–432.
- Favell, A. 2008. The new face of East–West migration in Europe. *Journal of ethnic and migration studies*, 34(5): 701–716.
- Freeman, L. 1978. Centrality in social networks conceptual clarification. *Social networks*, 1(3): 215–239.
- Golovchenko, Y.; Hartmann, M.; and Adler-Nissen, R. 2018. State, media and civil society in the information warfare over Ukraine: citizen curators of digital disinformation. *International Affairs*, 94(5): 975–994.
- González-Bailón, S.; Borge-Holthoefer, J.; Rivero, A.; and Moreno, Y. 2011. The dynamics of protest recruitment through an online network. *Scientific reports*, 1(1): 1–7.
- Graham, M.; Hale, S. A.; and Gaffney, D. 2014. Where in the world are you? Geolocation and language identification in Twitter. *The Professional Geographer*, 66(4): 568–578.
- Granovetter, M. 1973. The strength of weak ties. *American journal of sociology*, 78(6): 1360–1380.
- Guille, A.; Hacid, H.; Favre, C.; and Zighed, D. A. 2013. Information diffusion in online social networks: A survey. *ACM Sigmod Record*, 42(2): 17–28.
- Hale, S. A. 2012. Net increase? Cross-lingual linking in the blogosphere. *Journal of Computer-Mediated Communication*, 17(2): 135–151.
- Hale, S. A. 2014a. Global connectivity and multilinguals in the Twitter network. In *CHI*, 833–842.
- Hale, S. A. 2014b. Multilinguals and Wikipedia editing. In *WebSci*.
- Hammarström, H.; Forkel, R.; Haspelmath, M.; and Bank, S. 2021. Glottolog Database 4.5.
- Head, K.; and Mayer, T. 2014. Gravity equations: Workhorse, toolkit, and cookbook. In *Handbook of international economics*, volume 4, 131–195. Elsevier.
- Herring, S. C.; Paolillo, J. C.; Ramos-Vielba, I.; Kouper, I.; Wright, E.; Stoerger, S.; Scheidt, L. A.; and Clark, B. 2007. Language networks on LiveJournal. In *HICSS*, 79–79. IEEE.
- Ho, D. E.; Imai, K.; King, G.; and Stuart, E. A. 2007. Matching as nonparametric preprocessing for reducing model dependence in parametric causal inference. *Political analysis*, 15(3): 199–236.

- Holland, P. W. 1986. Statistics and causal inference. *Journal of the American statistical Association*, 81(396): 945–960.
- Hong, L.; Convertino, G.; and Chi, E. H. 2011. Language matters in twitter: A large scale study. In *Fifth international AAAI conference on weblogs and social media*.
- Hoyle, A.; Goel, P.; Hian-Cheong, A.; Peskov, D.; Boyd-Graber, J.; and Resnik, P. 2021. Is Automated Topic Model Evaluation Broken? The Incoherence of Coherence. *NeurIPS*.
- Jin, H. 2017. Detection and characterization of influential cross-lingual information diffusion on social networks. In *WWW*.
- Joulin, A.; Grave, E.; Bojanowski, P.; Douze, M.; Jégou, H.; and Mikolov, T. 2016. FastText.zip: Compressing text classification models. *arXiv preprint arXiv:1612.03651*.
- Jurafsky, D.; and Martin, J. H. 2000. *Speech and language processing - an introduction to natural language processing, computational linguistics, and speech recognition*. Prentice Hall.
- Jurgens, D.; Tsvetkov, Y.; and Jurafsky, D. 2017. Incorporating Dialectal Variability for Socially Equitable Language Identification. In *Proceedings of ACL*, 51–57.
- Kim, S.; Park, S.; Hale, S. A.; Kim, S.; Byun, J.; and Oh, A. H. 2016. Understanding editing behaviors in multilingual Wikipedia. *PLoS one*, 11(5): e0155305.
- Kim, S.; Weber, I.; Wei, L.; and Oh, A. 2014. Sociolinguistic analysis of twitter in multilingual societies. In *HT*, 243–248.
- Kossinets, G.; Kleinberg, J.; and Watts, D. 2008. The structure of information pathways in a social communication network. In *KDD*.
- Lamprinidis, S.; Bianchi, F.; Hardt, D.; and Hovy, D. 2021. Universal Joy: A data set and results for classifying emotions across languages. In *The 16th Conference of the European Chapter of the Association for Computational Linguistics*. Association for Computational Linguistics.
- Lee, M.; and Murdie, A. 2020. The Global Diffusion of the #MeToo Movement. *Politics & Gender*, 1–29.
- Leetaru, K.; and Schrodt, P. A. 2013. Gdelt: Global data on events, location, and tone, 1979–2012. In *ISA annual convention*, volume 2, 1–49. Citeseer.
- Lui, M.; and Baldwin, T. 2012. langid.py: An off-the-shelf language identification tool. In *Proceedings of the ACL 2012 system demonstrations*, 25–30.
- Lui, M.; and Baldwin, T. 2014. Accurate language identification of Twitter messages. In *Proceedings of the 5th workshop on language analysis for social media (LASM)*, 17–25.
- McAdam, D.; and Rucht, D. 1993. The cross-national diffusion of movement ideas. *The Annals of the American Academy of Political and Social Science*, 528(1): 56–74.
- Monroe, B. L.; Colaresi, M. P.; and Quinn, K. M. 2008. Fightin’ words: Lexical feature selection and evaluation for identifying the content of political conflict. *Political Analysis*, 16(4).
- Nissenbaum, A.; and Shifman, L. 2018. Meme templates as expressive repertoires in a globalizing world: A cross-linguistic study. *Journal of Computer-Mediated Communication*, 23(5): 294–310.
- Park, S. J.; Lim, Y. S.; and Park, H. W. 2015. Comparing Twitter and YouTube networks in information diffusion: The case of the “Occupy Wall Street” movement. *Technological forecasting and social change*, 95: 208–217.
- Reimers, N.; and Gurevych, I. 2019. Sentence-BERT: Sentence Embeddings using Siamese BERT-Networks. In *EMNLP-IJCNLP*.
- Rijhwani, S.; Sequiera, R.; Choudhury, M.; Bali, K.; and Maddila, C. S. 2017. Estimating code-switching on twitter with a novel generalized word-level language detection technique. In *Proceedings of the 55th annual meeting of the association for computational linguistics (volume 1: long papers)*, 1971–1982.
- Romero, D. M.; Meeder, B.; and Kleinberg, J. 2011. Differences in the mechanics of information diffusion across topics: idioms, political hashtags, and complex contagion on twitter. In *WWW*, 695–704.
- Rosenbaum, P. R.; and Rubin, D. B. 1983. The central role of the propensity score in observational studies for causal effects. *Biometrika*, 70(1): 41–55.
- Samoilenko, A.; Karimi, F.; Edler, D.; Kunegis, J.; and Strohmaier, M. 2016. Linguistic neighbourhoods: explaining cultural borders on Wikipedia through multilingual co-editing activity. *EPJ data science*, 5: 1–20.
- Sarkar, T.; Forber-Pratt, A. J.; Hanebutt, R.; and Cohen, M. 2021. “Good morning, Twitter! What are you doing today to support the voice of people with #disability?”: disability and digital organizing. *Journal of Community Practice*, 29(3): 299–318.
- Schünemann, W. J. 2020. Ready for the world? Measuring the (trans-) national quality of political issue publics on Twitter. *Media and Communication*, 8(4): 40–52.
- Scotch, R. K. 1988. Disability as the basis for a social movement: Advocacy and the politics of definition. *Journal of Social Issues*, 44(1): 159–172.
- Spörlein, C.; and Schlueter, E. 2021. Ethnic Insults in YouTube Comments: Social Contagion and Selection Effects During the German “Refugee Crisis”. *European Sociological Review*, 37(3).
- Srivastava, A.; and Sutton, C. 2017. Autoencoding Variational Inference for Topic Models. In *ICLR*.
- Stuart, E. A. 2010. Matching methods for causal inference: A review and a look forward. *Statistical science: a review journal of the Institute of Mathematical Statistics*, 25(1): 1.
- Takhteyev, Y.; Gruzd, A.; and Wellman, B. 2012. Geography of Twitter networks. *Social networks*, 34(1): 73–81.
- Talukdar, D.; Sudhir, K.; and Ainslie, A. 2002. Investigating new product diffusion across products and countries. *Marketing Science*, 21(1): 97–114.
- Tsur, O.; and Rappoport, A. 2012. What’s in a hashtag? Content based prediction of the spread of ideas in microblogging communities. In *WSDM*.
- Ugander, J.; Karrer, B.; Backstrom, L.; and Marlow, C. 2011. The anatomy of the facebook social graph. *arXiv preprint arXiv:1111.4503*.
- Williams, J.; and Dagli, C. 2017. Twitter language identification of similar languages and dialects without ground truth. In *Proceedings of the fourth workshop on NLP for similar languages, varieties and dialects (VarDial)*, 73–83.
- Zuckerman, E. 2008. Meet the bridgebloggers. *Public Choice*, 134(1-2): 47–65.