

On the Joint Effect of Culture and Discussion Topics on X (Twitter) Signed Ego Networks

JACK TACCHI, Institute of Informatics and Telematics, Consiglio Nazionale delle Ricerche, Italy and Scuola Normale Superiore di Pisa, Italy

CHIARA BOLDRINI, ANDREA PASSARELLA, and MARCO CONTI, Institute of Informatics and Telematics, Consiglio Nazionale delle Ricerche, Italy

Humans are known to structure social relationships according to certain patterns, such as the Ego Network Model (ENM). These patterns result from our innate cognitive limits and can therefore be observed in the vast majority of large human social groups. Until recently, the main focus of research was the structural characteristics of this model. The main aim of this paper is to complement previous findings with systematic and data-driven analyses on the positive and negative sentiments of social relationships, across different cultures, communities and topics of discussion. A total of 26 datasets were collected for this work. It was found that contrary to previous findings, the influence of culture is not easily “overwhelmed” by that of the topic of discussion. However, more specific and polarising topics do lead to noticeable increases in negativity across all cultures. These negativities also appear to be stable across the different levels of the ENM, which contradicts previous hypotheses. Finally, the number of generic topics being discussed between users seems to be a good predictor of the overall positivity of their relationships.

CCS Concepts: • **Networks** → **Topology analysis and generation**; • **Social and professional topics** → **Cultural characteristics**; • **Information systems** → *Sentiment analysis*.

Additional Key Words and Phrases: online social networks, ego network model, signed networks, signed ego networks, sentiment analysis, cultural analysis, topic analysis

ACM Reference Format:

Jack Tacchi, Chiara Boldrini, Andrea Passarella, and Marco Conti. 2024. On the Joint Effect of Culture and Discussion Topics on X (Twitter) Signed Ego Networks. 1, 1 (February 2024), 21 pages. <https://doi.org/10.1145/nnnnnnn.nnnnnnn>

1 INTRODUCTION

Throughout human history, our ability to communicate has been a defining trait of our species. Communication has an immeasurable impact on our behaviour and its importance is evident at all levels of society. It affects everyday interactions between individuals as well as how our societies are organised. Very importantly, it is primarily through communication and interaction that we build our social networks. Given the omnipresence of communications, one might expect a great deal of variety in the resulting patterns of social relationships. However, some intriguingly common patterns can be observed, resulting from our innate cognitive limits.

Authors' addresses: Jack Tacchi, jack.tacchi@sns.it, Institute of Informatics and Telematics, Consiglio Nazionale delle Ricerche, 1 Via Giuseppe Moruzzi, Pisa, Italy, 56124 and Scuola Normale Superiore di Pisa, 7 Piazza dei Cavalieri, Pisa, Italy, 56126; Chiara Boldrini, chiara.boldrini@iit.cnr.it; Andrea Passarella, andrea.passarella@iit.cnr.it; Marco Conti, marco.conti@iit.cnr.it, Institute of Informatics and Telematics, Consiglio Nazionale delle Ricerche, 1 Via Giuseppe Moruzzi, Pisa, Italy, 56124.

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

© 2024 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM XXXX-XXXX/2024/2-ART

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

One such pattern can be observed between the proportional size of a species' neocortex and the size of a social group that that species can maintain [9]. The link between these two variables is known as the Social Brain Hypothesis and it further posits that if a social group grows beyond its maintainable size, it will inevitably start to break down into smaller, less cognitively demanding collectives. This phenomenon is so ingrained in our evolutionary neurology that it has been observed not just for humans but also for many other types of primates, and even some species of birds [10]. By extrapolating the observed group sizes of animals up to the size expected for an animal with the neocortex size of a human, one would expect our own social group size limit to be around 150 (known as Dunbar's number). Indeed, 150 is a common unit size across human social structures and has been recorded in contexts as diverse as modern-day militaries' company sizes and traditional hunter-gatherer communities from 5 different continents [11].

Additionally, when a human social network is viewed from the point of view of a single individual, Dunbar's number once again emerges. Indeed, the number of annually active relationships being maintained is seldom far from 150. If these relationships are then organised by the strength of their connection to the initial subject, a series of concentric circles of increasing size but decreasing connection strength will almost inevitably be observed [13]. This individual-focused representation of a social network is called the Ego Network Model (ENM), with the individual in question being referred to as the Ego (from which the model takes its name) and their connections being called Alters. What's more, because these circles are another emergent pattern resulting from neocortical limits, their sizes are incredibly regular. For humans, these sizes are 5 (support clique), 15 (sympathy group), 45-50 (affinity group) and 150 (active network), with each subsequent circle size increasing by a ratio of around 3. Ego networks have also been studied for Online Social Networks. Quite remarkably, despite the immediacy of establishing social relationships, data-driven studies have shown that the ENM can be found also in Online Social Networks [12]. Further, Ego Networks feature the same layered structure, and the sizes of the layers align remarkably well with what has been found in "offline" social networks. The only notable difference is the emergence of an additional inner-most layer, of average size 1.5 [12]. As previously observed [22], this confirms that Ego Network structures are determined by human cognitive limits no matter how advanced the "tools" used to facilitate social interactions. Figure 1 shows a standard representation of an Ego Network from an online context.




Fig. 1. The Ego Network Model, taken from [27], the displaying expected numbers of Alters in each circle.

Of course, determining how to measure the strength of the connection between an Ego and their Alters is of key importance to the ENM. What's more, it is not something that can be objectively

defined. Fortunately, there is a renowned definition that has become the standard for research on Ego Networks: the tie strength between two individuals is the equally weighted combination of the time spent maintaining the relationship, its emotional intensity, its level of intimacy and the reciprocal services it generates [17]. Given the obvious difficulty of objectively measuring the tie strength between two individuals along the dimensions defined in [17], traditionally, the Ego-Alter contact frequency has often been used. While it has been found that this is an excellent proxy metric [15] for relationship strength, it may miss a lot of potentially important information. One could reasonably expect contact frequency to be strongly correlated with the time spent on the relationship (perhaps explaining why it is such a good proxy metric) but the qualitative aspects of relationships might be completely ignored.

To address this problem, some recent works [26, 27] have taken a bottom-up approach by applying sentiment analysis models to generate a label for each interaction within a relationship and then, using a psychology-based threshold [16], inferring a *sign* for the relationship as a whole. This threshold, referred to here as the “Golden Interaction Ratio”, comes from observations that relationships with more than 1 negative interaction for every 5 positive interactions (roughly 17%) tend to cause myriad problems for those involved. For marriages, this means a significant increase in the probability of divorce [16], and for parental relationships, can lead to the child developing behavioural problems and/or struggling in school [19].

Using the aforementioned inferred signs, some surprising novel insights were able to be obtained, providing solid grounds for the definition of a Signed Ego Network Model (hereafter referred to as SENM). For instance, despite the remarkably predictable structure of the ENM, its relationship signs can vary dramatically. Indeed, one past paper made a concentrated effort to compare levels of negativity across multiple different cultures and communities [27]. Networks were observed to vary between 49.90% and 69.47%. What’s more, it was found that, although there were some differences due to culture, these appeared to be overwhelmed by the impact of the topic around which a community was formed (e.g. journalism or reality TV). Furthermore, the influence of the topic was stronger the more negative it was, with the differences between cultures decreasing dramatically for the most negative topics. However, in the absence of a specific topic-based community, cultural differences were still very much observable. Understanding these cultural differences becomes even more important the more different cultures interact and, in the globally connected modern world, intercultural communications occur more frequently than ever.

In addition to cultural differences, and in contradiction with previous expectations, the most negative circles of the SENM tend to be the innermost ones [26]. In the conventional ENM, the inner circles are where our most important and trusted connections lie [25], so this was a particularly surprising finding. As with the negativity of each network as a whole, a better understanding of how the negativities of the different circles within the SENM change between cultures may reveal non-trivial and impactful insights.

This paper is an extension of the aforementioned cultural analysis of the SENM [27], which aims to enhance many of the previous findings with regard to differences arising from culture and topic. Specifically, by providing a rigorous analysis of numerous combinations of topic and culture. This has resulted in 4 main novel discoveries. The first of these is that, although there are some differences in negativity caused by both culture and topic, cultural differences do not seem to get “overwhelmed” by topics as easily as previously thought. Instead, the effects of culture on the negativity of a SENM appear noticeable the majority of the time and are only overpowered by *extremely controversial* topics. Next, when a network is centred around a more specific or controversial topic, it invariably displays an increase in negativity at all levels of the SENM; this effect is observable for all cultures. Thirdly, the level of negativity across the circles of the SENM do not appear to be less stable for more negative or specific topics, as was found previously. Finally, the

number of controversial topics within a dataset's top 20 most popular topics can be used to predict how negative the relationships in it are overall, regardless of how negative the corresponding individual tweets are.

2 METHODOLOGY

2.1 Signing Relationships

To determine the sign of each relationship, a bottom-up approach, as already been employed for SENM research [26, 27], was employed. First, this method groups all interactions between each Ego and Alter. Then, it leverages the extremely well-established field of sentiment analysis to provide a sentiment ("positive", "neutral" or "negative") based on the text of each interaction individually.

While the exact choice of model used for this step has varied in previous work, it has been shown that the relationship signs are not overly affected by this choice [28]. Given the use of multilingual data in this paper, a polyglot model was used: XLM-T [4]. This model is based on another multilingual model, XLM-R [6], which was trained on Wikipedia texts from 100 different languages. XLM-T was then hyper-tuned for dealing with Tweets by further training it on 198 million Tweets from over 60 languages. Using the XLM-T, a list of sentiments for each interaction in every Ego-Alter pair was computed.

The final step is to apply the Golden Interaction Ratio of 17% (mentioned in Section 1). This results in a positive sign for relationships with percentages of negative interactions below or equal to this threshold, and a negative sign for those above it.

2.2 Computing Signed Ego Networks

To generate a Signed Ego Network, the unsigned version of it first needs to be obtained. As described in Section 1, this is done by organising Alters around their corresponding Ego, based on their tie strength. In purely practical terms, this means clustering the Alters based on the Ego's contact frequency with them. Multiple algorithms have been used for clustering the Alters in Ego Network research. The one used in the current work, the MeanShift algorithm [14], is one of the most commonly used [5] and also has the added benefit of automatically determining the optimum number of circles (clusters) for an Ego.

After computing both the unsigned Ego Networks and the relationship signs (obtained via the steps outlined in Subsection 2.1), it is a simple matter of matching the relationship signs with the corresponding Ego-Alter pairs in each Ego Network in order to obtain the Signed Ego Networks.

2.3 Topic Analysis

Previous work has found that the percentage of negative relationships in a dataset can be dependent on the type of community that the dataset contains. For instance, communities based around more specific topics, such as journalism or reality TV, will likely be more negative than those of a more arbitrary or generic nature, such as geographical region [27]. Thus, an in-depth analysis of the topics within some of the larger generic datasets used in this paper was conducted (see Section 3), to see whether this effect can be replicated by focusing on the subcommunities within a generic dataset.

In order to accurately detect the main topics of discussions within multilingual data in a standardised way, the BERTopic model [18] was employed. BERTopic is a topic modelling tool that uses a mixture of transformers and TF-IDF to identify and group important topics within a collection of natural-language documents. BERTopic was compared against 5 other state-of-the-art topic models in terms of two staple metrics: Topic Coherence [21] and Topic Diversity [8]. These measure how well a model's grouped terms fit with one another and how much variety there is among grouped

words. BERTopic was found to consistently outperform the other models for Topic Coherence while also remaining very competitive for Topic Diversity [18]. To optimise BERTopic for multilingual data, the default transformer model it uses can be replaced with paraphrase-multilingual-MiniLM-L12-v2 [23]. This is a sentence-transformer model that is able to accurately process data in over 50 languages [24].

Unfortunately, paraphrase-multilingual-MiniLM-L12-v2 can only take in the first 384 tokens from each document, making it impossible to pass in an entire User Tweet Timeline (which often contains a few thousand Tweets, each of up to 280 characters) as a single document. Tweets were therefore parsed individually, meaning that each Tweet was treated as being entirely distinct from all the others, even if they were created by the same user. An unfortunate consequence of this is that some of the topics may appear to be undeservedly important if, for example, they are being spam-tweeted by an individual user. In an effort to mitigate this, and in order to get a perspective of the topics being mentioned by all users in each dataset, the top 200 topics provided by BERTopic were collected and the number of distinct users involved with them was computed. The 20 topics with the largest numbers of unique users were then chosen as the focus of analysis. Effectively, this meant that each topic was counted a maximum of once per user, rather than once per Tweet related to the topic. These topics are listed and discussed in Subsection 4.4.

3 DATASETS

3.1 Data Source

Every dataset included in this paper was collected from the X social media platform (formerly known as Twitter). Historically, this platform has been a consistent source of Ego Network data for over a decade [2, 5] and used to allow quick and easy access to a huge amount of public social communications data from users all across the globe via its API. In addition to this, the site allows its users to interact with one another in a variety of ways that make the data especially suitable for Ego Network research. Namely, these are Replies, Mentions and Retweet, which allow specific users to be tagged whenever a communication is made, in turn allowing Egos and Alters to be easily mapped and tracked to individual interactions.

All of the data used in this paper were collected using the, now defunct, free academic version of the Twitter API that, among others, provided two important endpoints. These were Twitter Search, which took in a search query and provided a stream of Tweets relating to it (in reverse chronological order), and User Tweet Timelines, which took in the ID or name of a specific user and returned the entirety of their publicly created Tweets. Of course, this latter also includes Tweets that contain no communication data between users, which is useless for the current research. Fortunately, X makes it possible to know when users are specifically contacting another user or users, via the use of “Mentions”, “Replies” and “Retweets”. Although the last of these three options is only considered a communication for this work if the authoring user also adds some text (known as a “Quote Retweet”). This is to ensure there is at least some cognitive involvement from the Ego and is a standard approach in the related literature [26].

The datasets have been split between those that existed prior to this paper and those that were collected as part of it. They are described in further detail, in Subsections 3.2 and 3.3 respectively, and some descriptive statistics are provided in Table 1 (after all preprocessing steps mentioned in Subsection 3.4).

3.2 Pre-existing Datasets

As a starting point for the current work, 11 datasets were collected from previous Ego Network research. These represent data from a mixture of different countries. Additionally, they can be

Table 1. Number of Egos, Relationships and Interactions in the active Ego Networks, after all preprocessing steps

Dataset Type	Region	Egos	Relationships	Interactions
Baseline	–	4,049	574,585	8,593,290
Geographical	Mediterranean	878	120,068	2,191,666
	South America	217	25,205	441,158
	Northern Europe	552	82,237	1,273,881
	West Africa	396	55,884	884,321
Reality TV	Italy	160	18,884	291,213
	Brazil	154	15,685	234,734
	Netherlands	230	24,082	441,694
Journalists	Italy	203	30,409	489,008
	Brazil	154	20,348	278,631
	Netherlands	1,316	179,668	2,702,275
Generic Users	Italy	2,740	266,701	2,133,608
	Brazil	8,223	820,165	6,561,320
	Netherlands	9,278	863,187	6,905,496
Weather	Italy	518	42,168	337,344
	Brazil	598	42,553	340,424
	Netherlands	255	14,427	115,416
	Nigeria	363	22,628	181,024
Football	Italy	1,320	141,565	1,132,520
	Brazil	1,024	117,199	937,592
	Netherlands	1,910	156,102	1,248,816
	Nigeria	159	14,206	113,648
Politics	Italy	2,004	218,005	1,744,040
	Brazil	482	44,333	354,664
	Netherlands	1,256	151,922	1,215,376
	Nigeria	866	65,846	526,768

easily sorted between “generic” and “specialised” users. Generic users are those who use the social media platform (in this case X) for predominantly social reasons, whereas specialised users use it for professional reasons. These two types of users have been shown to exhibit differing behaviours in online contexts [29] and it is therefore important to bear this in mind during the collection and analysis of the data.

3.2.1 Baseline. The first of the datasets used in this paper is referred to in this paper as Baseline because its initial purpose was to obtain a baseline measurement of an X user’s percentage of negative relationships [26]. It was collected using a snowball sampling methodology with an initial set of 31 seed users. The User Tweet Timelines of these users were then gathered, followed by those of their Alters, then of their Alters’ Alters, and so on, until the collection period ended. The seed users were randomly selected from another large Ego Network dataset, which was itself a snowball sampling that used Barack Obama’s official X account as its seed user [2]. The Baseline dataset is

composed of the User Tweet Timelines of the collected users¹ and was collected between the 27th April and the 25th May 2022.

3.2.2 Geographical. The next 4 datasets were collected using the same snowball collection methodology as the Baseline dataset [27]. However, rather than obtaining a baseline of the entirety of X, the purpose of these datasets was to gauge the negative relationships of specific regions. This was achieved by using seed users from countries in 4 geographically and culturally distinct regions: Mediterranean (Spain, France, Italy, Greece), South America (Brazil, Colombia, Venezuela), Northern Europe (Germany, Netherlands, Sweden) and West Africa (Nigeria, Senegal, Ghana)². These datasets were collected between the 16th June 2022 and the 26th July 2022.

3.2.3 Reality TV. The next set of datasets is the first that targets a specific type of user, in this case, those who follow reality TV shows. This was done by using Twitter Search to query for hashtags related to shows that were popular in 3 target countries: Italy (#XF2022, #GFVIP), Brazil (#XFactorBR, #BBB22) and the Netherlands (#HollandsGotTalent, #IkVertrek)³ [27]. The previous work did not include a dataset corresponding to the West Africa region. Seed users were manually identified for each dataset using the results of the aforementioned search and User Tweet Timelines were then collected using the same snowball sampling methodology as the Baseline. The Reality TV datasets were collected between the 21st and the 29th January 2023.

3.2.4 Journalists. The final set of pre-existing datasets used in this paper is taken from a paper that investigated the differences between generic users and journalists on X [30]. Unlike the snowball sampling method of the previous datasets, the users in this set were obtained via lists of verified journalist X accounts. The Journalist datasets contain the User Tweet Timelines of journalists from Italy, Brazil and the Netherlands respectively. Again, there was no corresponding dataset available for West Africa. These datasets were collected between the 14th and the 17th January 2018

3.3 Novel Datasets

To complement the pre-existing datasets, an additional 15 datasets were collected⁴. These were chosen to better understand phenomena observed in previous work [27]. Specifically, they were selected to compare the differences in negativity arising between topics of differing levels of controversy, and their interplay with cultural factors.

All of these datasets used the same snowball sampling collection method as the Baseline dataset. All randomly selected seed users were manually checked for suitability and replaced if they were thought to be spammers, bots, businesses or individuals from outside the desired country.

3.3.1 Generic Users. The first set of novel datasets was created to fulfil a similar rationale to that of the pre-existing Geographical datasets but, instead of collecting users from multiple countries within a large region, these datasets took a more precise focus and only collected users from specific countries. Specifically, we picked one country for each of the macro-areas in the Geographical dataset, focusing specifically on Italy, The Netherlands, Brazil and Nigeria. These countries were chosen by counting the number of active X users in the Baseline dataset for each country and selecting the highest for each region. These datasets were collected between the 3rd May and the 4th June 2023. Unfortunately, the academic version of the Twitter API was disabled around the end of this period, which prevented the collection of a Nigerian Generic Users dataset.

¹All Tweet IDs from this dataset are provided at <https://zenodo.org/records/7717006>.

²All Tweet IDs from these datasets are provided at <https://zenodo.org/records/7717047>.

³All Tweet IDs from these datasets are provided at <https://zenodo.org/records/7716860>.

⁴All Tweet IDs from these datasets are provided at <https://zenodo.org/records/10605838>.

Because it is not possible to manually check the country of every user in the datasets and, as the snowball sampling method allows for the possibility of users to be collected from outside the target country, some tests were performed after the collection of the Generic User datasets. These involved checking the self-declared location and X-detected main language of each user. For the Brazilian and Dutch datasets, the vast majority of users' locations and languages were as expected. However, the Italian dataset showed a significant proportion (around a third) of users from the UK and USA. To combat this, users were removed from this dataset if their location contained "UK", "London", "USA", "DC" or "CA" or if their main language was English. Unfortunately, this resulted in a much smaller dataset than the other two Generic Users, however, this is still one of the larger datasets used in this paper. Indeed, these datasets specifically included much larger amounts of data so that the subcommunities of these datasets could be examined (see Subsection 4.4).

3.3.2 Weather. The next set contains slightly less generic datasets focused around the weather. The initial seed users were randomly selected from users who commented on the posts of local weather forecasting X accounts. These accounts were Meteo Italia (@meteo_italia7) and meteo.it (@wwwmeteoit) for Italy, MetSul Meteorologia (@metsul) for Brazil and Weer & Radar Nederland (@weerenradar_nl) for the Netherlands. For Nigeria, it was not possible to find a weather-related account that generated more than a few comments from other users, so the seed users were collected using a Twitter Search for "weather nigeria" instead. These datasets were collected between the 21st and 25th April 2023.

3.3.3 Football. The first of the two specialised sets of novel datasets is themed on Football. For these datasets, seed users were collected from users commenting on posts made by popular local football teams: Juventus (@juventusfc) for Italy, Regatas do Flamengo (@Flamengo) for Brazil, Ajax (@AFCAjax) for the Netherlands and Enyimba (@EnyimbaFC) and Plateau United (@plateau_united) for Nigeria. Nigeria was the only country whose most popular football teams were from foreign countries, such as Manchester United, Chelsea and Barcelona. This was also why a second team was included, as it was not possible to get enough seed users using a single Nigerian football team. These datasets were collected between the 20th April and the 22nd May 2023.

3.3.4 Politics. The final set of datasets used in this paper focuses on Politics. Seed users for these were taken from users commenting on posts made by political parties in each of the target countries. In order to get a broader image of the general political discussions of each country, rather than that of any single political party, a list of seed users was generated for multiple parties for each country, and the final seeds used were selected randomly from these lists, with a minimum of 5 users from each party's list. For Italy, the chosen parties were Fratelli d'Italia (@FratellidItalia), Lega Salvini (@LegaSalvini) and Partito Democratico (@pdnetwork), for Brazil they were Partido Liberal (@PartidoLiberal), Movimento Democrático Brasileiro (@MDB_Nacional) and Partido dos Trabalhadores (@ptbrasil), for the Netherlands Volkspartij voor Vrijheid en Democratie (VVD), Democraten 66 (@D66), Christen-Democratisch Appèl (@cdavandaag) and for Nigeria they were All Progressives Congress (@OfficialAPCNig), Peoples Democratic Party (@OfficialPDPNig), Labour Party (@OfficialPDPNig) and New Nigeria Peoples Party (@OfficialNNPPNig). These datasets were collected between the 27th April and the 20th May 2023.

3.4 Preprocessing

3.4.1 Non-Human Users. Unfortunately, not all accounts on X are controlled by human individuals. When sampling such large quantities of accounts, many spammers, bots and groups (such as businesses) will inevitably end up among them. Obviously, such undesired users do not have the same cognitive constraints as a single human, meaning that they will not display the ENM structure

within their communications. Therefore, it is an important preprocessing step to identify and remove any non-human accounts. Given the common usage of X for the collection of Ego Network data, this is a common problem. The usual method of filtering out non-humans is to use a Support Vector Machine (SVM) [7], trained on a set of 500 users, to label each account as “people” or “other”. Both this filtration method and the training set have been established in previous ENM research [2]. Using k-fold cross-validation (with a k value of 5), this model originally achieved an accuracy of 81.3%. Once the SVM has been trained, it is a trivial matter to run the model over the collected data and remove the accounts predicted to be “other”.

This preprocessing step was performed for all datasets except the Journalists. This is because those users were collected from a verified list of known accounts of journalists and therefore did not contain any undesired types of users.

3.4.2 Irregular Egos. The second preprocessing step was to remove irregular users. Egos who spend little time on X or who engage with it infrequently will not have a fully formed Ego Network on the platform, which is problematic for the analyses of this paper. Therefore, Egos were removed if the total length of their User Tweet Timeline was less than 2,000 Tweets, if their Timeline spanned fewer than 6 months, or if they tweeted less than once every 3 days for more than half of the months they were active. These parameters have previously been shown to be appropriate for preparing data for Ego Network analysis [3].

Unlike the non-human filter, inactive users were removed from all the datasets including the journalists. This is because, although the list of journalists was verified, it is still possible that some of those users were not fully engaged with the platform when their data was collected.

3.4.3 Inactive Relationships. As stated in Section 1, an individual’s Ego Network is expected to contain up to around 150 Alters. However, this is not the number of users that an individual will interact with throughout their lifetime. Humans will interact with a significantly larger number of Alters but many of these will be one-off instances that will have little to no effect on cognitive load. As this threshold has traditionally been defined as once per year [20], Alters were removed if their Ego interacted with them, on average, less than once annually. Again, this preprocessing step was carried out for all of the datasets used in this paper.

4 RESULTS

4.1 Unsigned Ego Network Analysis

First of all, the unsigned structures of the Ego Networks in each dataset were analysed, to ensure that there are no anomalies in the datasets. The first step in this process is to calculate the mean active network sizes and mean numbers of circles. As can be seen in Table 2, the mean active network sizes of the datasets mostly fall around the 90 to 120 range. Although this is lower than the 150 that would be expected according to the standard ENM, it is common to observe slightly smaller active networks when using social media data [1]. This is because a user’s true active network will inevitably include real-world relationships that are not present online. Similarly, the mean number of circles is close to 5 for all the datasets, which is the expected number for online data [12]. In fact, 5 is the closest integer for all of the datasets except the Italian Journalist and the Dutch and Nigerian Weather datasets, for which it is 6, 4 and 4 respectively. Slight variations around 5 are commonly found in the literature [2, 12] and can be attributed to psychological and behavioural differences between individuals (for example, how social they are and, for social media datasets, how engaged they are with the given platform).

As the exact size and shape of an Ego Network can vary somewhat between individuals, it can sometimes make it difficult to compare them with one another or across a social network as a

Table 2. Mean active network sizes, number of optimum circles [95% confidence intervals] and number of Egos with 5 circles

Dataset Type	Region	Mean Network Size	Mean # Circles	# 5-Circle Egos
Baseline	–	99.05 [96.49, 101.60]	4.81 [4.78, 4.84]	1,160
Geographical	Mediterranean	109.68 [103.78, 115.58]	5.11 [5.03, 5.18]	374
	Northern Europe	119.60 [112.68, 126.51]	5.10 [5.01, 5.19]	275
	West Africa	102.32 [94.94, 109.69]	5.00 [4.90, 5.10]	206
	South America	101.94 [92.89, 110.99]	4.85 [4.73, 4.96]	130
Reality TV	Italy	103.65 [88.05, 119.25]	5.48 [5.25, 5.71]	40
	Brazil	96.63 [85.11, 108.14]	5.42 [5.22, 5.62]	48
	Netherlands	98.63 [86.41, 110.85]	5.20 [5.01, 5.39]	62
Journalist	Italy	120.10 [110.49, 129.70]	5.72 [5.20, 5.93]	51
	Brazil	116.48 [103.95, 129.01]	5.46 [5.27, 5.65]	50
	Netherlands	122.69 [118.42, 126.96]	5.45 [5.42, 5.51]	440
Generic Users	Italy	101.48 [97.36, 105.60]	4.88 [4.83, 4.94]	699
	Brazil	101.86 [99.60, 104.11]	4.96 [4.93, 4.99]	2,208
	Netherlands	102.64 [100.12, 105.16]	4.79 [4.76, 4.82]	2,274
Weather	Italy	102.86 [91.11, 114.62]	4.67 [4.55, 4.78]	106
	Brazil	90.17 [81.54, 98.80]	4.54 [4.45, 4.78]	142
	Netherlands	88.06 [73.50, 102.63]	4.28 [4.13, 4.43]	56
	Nigeria	81.60 [70.89, 92.31]	4.36 [4.24, 4.49]	68
Football	Italy	107.44 [101.74, 113.13]	5.08 [5.01, 5.15]	344
	Brazil	107.01 [101.74, 113.13]	5.27 [5.20, 5.35]	281
	Netherlands	88.65 [83.82, 93.48]	4.66 [4.60, 4.72]	457
	Nigeria	110.76 [96.58, 124.93]	4.80 [4.59, 5.01]	41
Politics	Italy	104.00 [99.49, 108.51]	5.12 [5.06, 5.18]	510
	Brazil	100.88 [92.76, 109.01]	4.91 [4.80, 5.03]	134
	Netherlands	111.31 [104.63, 118.00]	5.24 [5.16, 5.32]	332
	Nigeria	87.55 [82.27, 92.83]	4.67 [4.59, 4.75]	243

whole. To address this, it is common practice to standardise ENM analyses by focusing on Egos who have a number of circles exactly equal to 5 in their Ego Networks [2, 12]. The rightmost column of Table 2 displays the number of such Egos for each dataset. As the focus of this paper is most often on the active networks of all users, it shall be specifically stated whenever an analysis includes only Egos with 5 circles, as is the case or the circle-by-circle analysis that follows.

Next, the mean sizes of each individual circle, displayed in Table 3, are observed. Recalling from Section 1 that the outer circles of ENMs tend to be below standard in online contexts, one can see that the observed values are close to the expectations: i.e. 1-2, 5, 15, 45-50, 150. What's more, the scaling factor of 3 between each circle is very evident, despite the smaller outer circles. As with the overall network sizes, various external factors cause the outer circles to be marginally underpopulated, which is expected [2]. Therefore, these values are very in line with those of previous works on the ENM; the datasets appear to be appropriate for the computation of Ego Networks. What's more, because the sizes of the circles and the active networks, as well as the

Table 3. Mean circle sizes of Egos with exactly 5 circles

Dataset Type	Region	C ₁	C ₂	C ₃	C ₄	C ₅
Baseline	–	1.78	6.16	16.86	44.19	125.91
Geographical	Mediterranean	1.70	5.60	14.67	38.83	120.41
	South America	1.80	5.76	15.71	39.92	118.29
	Northern Europe	1.80	5.88	17.12	45.34	131.12
	West Africa	1.65	5.60	15.64	39.71	118.81
Reality TV	Italy	1.63	5.15	13.85	35.60	103.65
	Brazil	1.58	4.65	12.08	31.29	96.63
	Netherlands	1.61	5.29	14.29	37.16	98.63
Journalists	Italy	1.12	3.57	10.59	33.14	120.10
	Brazil	1.78	5.90	15.66	41.62	116.48
	Netherlands	1.66	5.51	15.54	43.08	122.69
Generic Users	Italy	1.55	4.91	12.77	32.43	96.29
	Brazil	1.63	5.00	12.98	33.31	98.29
	Netherlands	1.62	5.09	13.25	33.76	97.11
Weather	Italy	1.63	4.97	12.81	32.28	94.38
	Brazil	1.46	4.49	11.20	27.73	80.27
	Netherlands	1.34	4.29	10.41	25.80	75.96
	Nigeria	1.35	4.18	10.38	24.00	71.25
Football	Italy	1.72	5.52	14.12	35.32	104.91
	Brazil	1.63	5.24	13.43	34.67	106.67
	Netherlands	1.56	4.84	12.27	29.51	86.15
	Nigeria	1.54	5.22	12.49	32.07	101.24
Politics	Italy	1.70	5.38	13.40	33.38	101.22
	Brazil	1.70	5.07	13.31	33.11	97.55
	Netherlands	1.81	5.77	14.72	36.85	108.59
	Nigeria	1.52	4.56	11.47	27.70	83.22

mean number of circles, of all the datasets closely match those of the standard ENM, there is no discernible impact of culture or topic on these values.

Therefore, all of the collected networks display Egos with network structures that are well aligned with the general findings of Ego Network research (see Section 1).

4.2 Network Negativities by Topic and Region

Next, signs were computed for every relationship of every Ego in each dataset following the methodology detailed in Subsection 2.1. This allows the mean negativity for each dataset to be easily calculated and compared. To aid in this comparison, the datasets' overall mean negativities (including all Egos regardless of their number of circles) have been organised into two tables: Table 4 for the pre-existing datasets and Table 5 for the novel ones. (Also, recall from Section 3.2, that the values of the datasets in the Geographical column of Table 4 are not just the countries listed in the Region column but also those of culturally-similar, neighbouring countries). The datasets in both tables are arranged into rows by region and into columns by type of user (Table 4) or topic (Table 5). Both rows and columns are then sorted by negativity. When organised in this way, each negativity is larger than the one below it and to its left, with the only exceptions being Brazilian Journalists,

Table 4. Mean user negativities for datasets taken from previous papers, arranged by region and user type and ordered by negativity. Ranges between the Brazilian, Italian and Dutch datasets, ranges between all the datasets and ranges between all the topics are also displayed.

Region	Geographical	Journalists	Reality TV	Range
Brazil	65.67 ¹	64.93	69.47	4.89
Italy	60.08 ¹	63.87	64.97	4.54
Netherlands	54.66 ¹	57.65	68.36	13.71
Nigeria	50.29 ¹	–	–	–
Baseline	49.90	–	–	–
Range _{BIN}	11.01	7.27	4.50	–
Range _{all}	15.78	7.27	4.50	–

¹Note: The Geographical datasets are not exclusively the counties listed in the Region column, but also include a few neighbouring countries (see Subsection 3.3)

Table 5. Mean user negativities for datasets collected for this paper, arranged by region and topic and ordered by negativity. Ranges between the Italian, Brazilian and Dutch datasets, ranges between all the datasets and ranges between all the topics are also displayed.

Region	Generic	Weather	Football	Politics	Range
Italy	66.75	74.24	77.96	84.34	10.10
Brazil	64.58	69.88	71.41	81.24	11.36
Netherlands	54.55	56.67	64.09	78.25	21.57
Nigeria	–	60.73	58.45	68.30	9.86
Range _{IBN}	12.19	17.57	13.87	6.09	–
Range _{all}	12.19	17.57	19.51	16.03	–

Dutch Reality TV and Nigerian Weather. This strongly illustrates that both the geographical culture and the communities/topics that individuals are engaged with have a pronounced impact on the percentage of negative relationships that they maintain⁵. What’s more, a clear pattern naturally emerges in the order of the topics in Table 5: from least controversial (Generic) to most controversial (politics). This strongly illustrates that the more a topic is controversial, the higher its negativity, and this holds across all cultures. Additionally, the negativities of the country-specific Generic Users are very close to those of the Geographical datasets, which also include users from neighbouring countries. This would suggest users from the countries that have been grouped together in the latter are fairly similar in terms of negative relationship percentage.

While the results of the previous paragraph are unsurprising given the results of previous works (namely [27]), what is surprising are the ranges of the novel datasets. Specifically, the range column in both tables shows the range of negativity across the corresponding rows (i.e. for a given country). Likewise, range rows in the tables show the ranges across columns (i.e., across topics/communities). Note that, to compare values across the two tables, we compute also the range rows only across countries that have complete datasets in the new datasets, i.e. all but Nigeria. As mentioned in Section 1, it has been observed that cultural differences in negativity tend to be overpowered by

⁵It is worth noting that Italy and Brazil swap positions between the two tables. Suggesting that the order of geographical cultures based on negativity can vary slightly between user types and/or topics.

the influence of the communities with which we are involved and that this effect is stronger the more negative or controversial a social group or topic is. Indeed, this effect appears visible in the range rows of Table 4: the more negative types of users display lower ranges (i.e. less difference between the different cultures). However, this effect is not visible for the novel datasets, where the ranges seem to oscillate without any correlation with this expectation. Instead, the ranges appear to not be strictly determined by the topics. However, focusing on the ranges for the 3 countries that have available datasets in each novel category (i.e. excluding Nigeria), there is a topic with a lower range than the Generic dataset: Politics. This, of course, is the most controversial of all the datasets collected, which may indicate that cultural effects may only be suppressed by topics that are exceedingly polarising. In the datasets of Table 4, the same holds for Reality TV, which can also be considered as a topic eliciting quite controversial discussions. Thus, it appears that the impact of culture is not *always* being “overwhelmed” by that of topics or subcommunities with which an individual is engaged as previously thought [27]. Rather the strength of the influence a given topic or community has on an individual depends somewhat on their culture and that, although topics can overwhelm cultural differences in extreme cases (e.g. Politics), this is not as common as previously thought. Intuitively, this conclusion seems logical as values and priorities can change dramatically between cultures, whereas politics can lead to strong and contradicting opinions in almost any culture. Indeed, the ranges could be taken as a measurement of how much the value of a given topic varies from culture to culture, rather than an inverse measure of how controversial it is. Therefore, future research may want to take steps to consider which topics could be considered controversial for each culture separately.

In contrast to those of the topics, the ranges of the countries (range columns) are more consistent: with all countries except the Netherlands displaying ranges that are relatively close to one another in both tables. This suggests that, while the impact a topic has on negativity across a set of cultures can vary drastically, the cultural influence on negativity across a set of topics may be, at least somewhat, predictable.

4.3 Signed Ego Network Analysis

Next, the distribution of the relationship signs across each circle is observed. As with the unsigned analysis of the circles, here the focus is on Egos with exactly 5 circles. The mean numbers and percentages of negative relationships for each of the 5 circles are displayed in Table 6. The most negative circle of each dataset is emphasised in bold. Thus, one can see that Circle 1 is the most negative circle for 7 of the datasets, Circle 2 for 13 of them and Circle 3 for the remaining 6. This is in line with previous findings on SENMs, as the inner circles are often found to contain the highest density of negative relationships [26].

Similarly in line with previous findings is that negativity percentages tend to be higher across all circles for datasets which are centred around a more specific topic. For example, the circle negativities of the Italian Generic Users are 67.78, 70.40, 69.50, 66.58 and 62.18 and each of these negativities is less than that of the corresponding circle for the Italian Weather dataset, 76.99, 76.66, 79.68, 78.87 and 74.87, which in turn are lower than the corresponding Italian Football negativities, 84.80, 87.84, 87.75, 86.40, 79.50, and so on. Therefore, there is a clearly visible increase in negativity for more specific topics and this increase is observable for all cultures and at all levels of the SENM.

Observing the variation in negativity across the first 4 circles⁶ and the corresponding range reveals further insights. As mentioned in Section 1, it has been observed that the percentages of

⁶Circle 5 is not included because it has previously been found that Alters in this circle usually have an average number of interactions of between 3 and 4, which is not considered reliable given the 1:5 ratio of the Golden Interaction Ratio (described in Section 1) used to sign the relationships [28].

negative relationships change less from circle to circle for users related to more generic sets of topics, as opposed to users related to more controversial topics [27]. This can be seen for the previous datasets, where the Baseline and Geographical datasets have, with one exception (Northern Europe), ranges between 3 and 5. Conversely, more specific datasets tended to have larger ranges, all of which are above 5 and half of which are over 10. It has been hypothesised that the increased ranges may be due to increased levels of user engagement and/or more polarising topics, resulting in a greater diversity of relationship negativities [27]. For the novel datasets, the previous findings are replicated for the generic dataset, Generic Users and Weather, which all have ranges between 2 and 6. However, what is more surprising is that the ranges of the specific datasets, Football and Politics, also fall within this interval (the only exception to this is Nigerian Football). Furthermore, Politics, which would be expected to have the most variation between circles, instead, as was the case for the ranges between cultures in Subsection 4.2, show quite small ranges. This further supports the suggestion that extremely controversial topics can overpower differences in negativities, causing them to converge, between cultures as previously seen, and also between different individuals within the same culture, as seen here.

These latter observations would suggest that contrary to previous conclusions, user type does not have an observable impact on the variations in negativity across circles. However, as mentioned when discussing the ranges of the overall negativities of each dataset (Subsection 4.2), how controversial a given topic is may vary significantly between different cultures. This may be an interesting avenue for future research.

Table 6. Mean number of negative relationships for Egos with 5 circles for each circle, with percentages in parentheses and the most negative circle of each dataset in bold. The percentage range between C_1 and C_4 is also displayed.

Dataset Type	Region	C_1	C_2	C_3	C_4	C_5	Range $_{C_1, C_4}$
Baseline	–	1.00 (56.25%)	3.63 (58.84%)	9.72 (57.64%)	24.28 (54.95%)	63.71 (50.60%)	3.90
Geographical	Mediterranean	1.25 (73.58%)	4.06 (72.54%)	10.38 (70.77%)	27.07 (69.70%)	76.85 (63.82%)	3.88
	South America	1.37 (76.42%)	4.42 (76.76%)	12.00 (76.38%)	28.71 (71.93%)	75.03 (63.43%)	4.83
	Northern Europe	1.26 (69.86%)	3.94 (67.05%)	11.04 (64.48%)	27.45 (60.54%)	70.67 (53.89%)	9.32
	West Africa	0.92 (55.40%)	3.18 (56.81%)	8.75 (55.94%)	21.25 (53.51%)	60.80 (51.17%)	3.30
Reality TV	Italy	1.08 (66.15%)	3.85 (74.76%)	10.58 (76.35%)	26.38 (74.09%)	71.38 (68.86%)	10.20
	Brazil	1.31 (82.89%)	3.83 (82.51%)	9.92 (82.07%)	24.31 (77.70%)	67.73 (70.09%)	5.20
	Netherlands	1.15 (71.00%)	3.97 (75.00%)	10.97 (76.75%)	27.73 (74.61%)	67.42 (68.36%)	5.75
Journalists	Italy	1.00 (89.47%)	3.12 (87.36%)	8.67 (81.85%)	25.80 (77.87%)	84.80 (70.61%)	11.60
	Brazil	1.14 (64.04%)	4.42 (74.92%)	11.94 (76.25%)	30.20 (72.56%)	77.02 (66.12%)	12.20
	Netherlands	1.19 (71.60%)	3.90 (70.74%)	10.78 (69.34%)	28.13 (65.30%)	71.50 (58.27%)	6.31
Generic Users	Italy	1.05 (67.78%)	3.45 (70.40%)	8.87 (69.50%)	21.59 (66.58%)	59.87 (62.18%)	3.82
	Brazil	1.19 (72.65%)	3.68 (73.72%)	9.48 (73.04%)	23.43 (70.72%)	63.15 (64.33%)	3.00
	Netherlands	0.93 (56.97%)	3.01 (59.12%)	7.87 (59.39%)	19.44 (57.59%)	51.82 (53.36%)	2.42
Weather	Italy	1.25 (76.88%)	3.81 (76.66%)	10.21 (79.68%)	25.46 (78.87%)	70.66 (74.87%)	3.02
	Brazil	1.18 (81.16%)	3.65 (81.32%)	8.88 (79.26%)	21.42 (77.27%)	56.26 (70.09%)	4.05
	Netherlands	0.84 (62.67%)	2.79 (65.00%)	6.70 (64.32%)	16.20 (62.77%)	45.14 (59.43%)	2.33
	Nigeria	1.01 (75.00%)	3.10 (74.30%)	7.54 (72.66%)	16.74 (69.73%)	45.57 (63.96%)	5.27
Football	Italy	1.46 (84.80%)	4.85 (87.84%)	12.39 (87.75%)	30.51 (86.40%)	83.41 (79.50%)	3.04
	Brazil	1.34 (82.28%)	4.34 (82.81%)	10.92 (81.31%)	27.13 (78.25%)	74.55 (69.89%)	4.56
	Netherlands	1.05 (67.13%)	3.43 (70.90%)	8.69 (70.84%)	20.35 (68.97%)	54.65 (63.44%)	3.77
	Nigeria	1.20 (77.78%)	4.12 (78.97%)	8.95 (71.68%)	21.02 (65.55%)	59.51 (58.78%)	13.42
Politics	Italy	1.56 (91.82%)	4.95 (92.01%)	12.25 (91.43%)	29.94 (89.70%)	84.98 (83.95%)	2.31
	Brazil	1.54 (90.79%)	4.71 (92.79%)	12.36 (92.88%)	29.99 (90.56%)	80.57 (82.60%)	2.32
	Netherlands	1.54 (85.14%)	5.03 (87.13%)	12.78 (86.78%)	31.09 (84.38%)	84.89 (78.17%)	2.75
	Nigeria	1.16 (76.69%)	3.47 (76.01%)	8.70 (75.84%)	20.57 (74.26%)	57.11 (68.62%)	2.44

4.4 Topic Analysis

Given the clear differences in rates of negative relationships due to topics displayed over the previous subsections, a further analysis was conducted to see whether dividing a generic dataset into some of its constituent topics would yield the same findings. As mentioned in Subsection 2.3, BERTopic can be used to obtain a list of key terms that are being used for each of these topics, making it possible to check each individual tweet to see if it is related to any of the top topics. Specifically, it was used to obtain the main topics being discussed in each of the 3 novel Generic Users datasets (those for Italy, Brazil and the Netherlands). This led to a list of IDs of both the tweets and the users involved, which in turn allowed the topics to be matched to 2 negativity metrics. The first of these is the percentage of negative relationships within the Ego Network of users that tweeted in relation to each topic, taken as the mean of all the users. The resulting value provides a gauge the negative impact a topic has on the relationships of those who are engaged with it. The second metric is simply the percentage of related tweets which are negative, which reveals how negative each topic is in isolation, i.e. irrespective of the surrounding network. These metrics are displayed in Table 7 (mean percentage of negative relationships) and Table 8 (percentage of negative tweets). In addition, some specific categories of topics have been chosen to be focused on in more detail. These categories represent a mix between the specific topics of the other datasets collected for this paper (Politics, Football and Generic) as well as two additions: COVID, which provides a unique opportunity to analyse a single event which has had an impact on every single person across the globe, and Religion, which was noted as a specific topic of interest for future research in previous work [27]. Both the aforementioned tables have been colour-coded to make these categories more visible.

Comparing the two negativity metrics provides a deeper understanding of how each topic affects negative relationships. For example, the words “peggio” and “gemist” are keywords in the topics most likely to be in a negative tweet, for the Italian and Dutch Generic Users datasets respectively (Table 8). This is unsurprising given that they mean “worse” or “the worst” in Italian and “excrement” in Dutch. However, they drop down to the 4th and 6th most negative topics in terms of impact on a user’s relationships (Table 7). Revealing that, although they are frequently used in negative contexts, their influence in terms of relationships is lower than expected. By comparison, “salvini” and “biden”, both well-known politicians, have the greatest negative impact on users’ relationships out of any topics of the Italian and Dutch datasets while being at the 4th and 6th positions for tweet negativity. This shows that, although they are less frequently used in negative contexts than “peggio” and “gemist”, they are much stronger indicators of negativity in a user’s surrounding network.

Subsequently, the negativities of the top topics are viewed when grouped into the aforementioned categories (Politics, Religion, Football, COVID and Generic). As was just touched upon, Politics is by far the most negative category of topic in all three datasets. They are also the most specific topics, with all of the topics referencing specific people (Matteo Salvini, Matteo Renzi, Vladimir Putin and Joe Biden), places (The Hague) or problems (Operação Lava Jato and rising fuel prices/Petroleo Brasileiro SA). The only non-specific political topic is “democrazia” (“democracy” in Italian), which is also the least negative (in terms of placement). Thus, the higher negativity of the political category of topics may be due to their specificity, which, as previously mentioned, can lead to a greater number of conflicting opinions.

Next, the two smallest categories: Religion and Football. Topics relating to these categories only appeared in the Italian and Brazilian datasets for the former and in the Italian for the latter. The religious topics were focused on specific individuals (the pope and Jesus) whereas Football is a generic topic relating to the sport as a whole (“calcio” being the Italian for football). Despite the

Table 7. Top 20 topics for the 3 Generic Users datasets and the mean percentage of negative relationships of users who are engaged with each topic, ordered by negativity. The topics are colour-coded: red for Politics, green for COVID, yellow for Religion, purple for Football and blue for Generic.

Index	Italy (66.75%)		Brazil (64.58%)		Netherlands (54.55%)	
	Topic	Negativity	Topic	Negativity	Topic	Negativity
1	salvini	89.17	lava	81.03	biden	64.03
2	renzi	86.88	putin	80.50	haag	61.42
3	amazon	79.07	gasolina	76.93	lachen	60.05
4	peggio	77.89	menino	76.49	nope	59.96
5	odio	76.97	verdades	75.93	hond	58.77
6	democrazia	76.93	jesus	75.53	gemist	57.07
7	nomi	75.03	cabelo	75.45	trein	56.42
8	pizza	74.75	perdi	74.88	vakantie	55.90
9	virus	74.26	fã	73.44	slapen	55.74
10	papa	74.17	festa	73.00	coronavirus	54.74
11	concordo	73.71	meme	72.77	filmpje	53.18
12	calcio	72.99	máscara	72.01	gold	51.53
13	thread	69.35	netflix	71.13	aflevering	51.15
14	dibattito	68.54	barato	70.99	facebook	50.93
15	caffè	66.71	gato	68.91	seizoen	50.38
16	natale	66.47	pizza	67.54	koffie	49.78
17	sogno	66.28	facebook	67.52	verjaardag	48.69
18	coronavirus	66.25	artista	66.11	interviews	48.40
19	facebook	63.84	natal	64.99	fotos	46.71
20	serata	62.33	dm	63.96	anniversary	40.22

religious terms being just as specific as the political ones, they are visibly less negative overall, suggesting that specificity alone is not enough to explain the differing negativities. Football appears just below the corresponding religious topic in its dataset in both of the tables.

The COVID category shows the biggest difference between the two tables. In Table 7 these topics appear towards the lower half of the table and, the least negative non-Generic term in each of the three datasets belongs to the COVID category. However, in Table 8, it is almost the exact opposite, with the most negative non-Generic term belonging to this category for both the Italian and Dutch datasets. This difference between the tables may suggest that, while individual Tweets related to COVID do tend to be very negative, tweeting negatively about COVID does not have an overly strong negative effect on an individual’s surrounding relationships. This highlights the difference between a negative topic and a controversial one. Despite COVID being a very negative category, there appear to be fewer differing opinions about it (i.e. there is a fairly strong consensus that COVID overall is bad) and, therefore, fewer disagreements compared to the other selected categories.

Counting the total number of topics in the final category, Generic, for each dataset, Italy has the fewest (13), followed by Brazil (15) and then the Netherlands (17). These numbers are reversely correlated with the order of the overall negativities of these datasets, which are 66.75%, 64.58% and 54.55% respectively. Thus, it appears that the more specific topics a user is engaged with, the higher the percentage of negative relationships they have is expected to be. This further supports previous work, which came to the same conclusion [27].

Table 8. Top 20 topics for the 3 Generic Users datasets and the mean negativity of all corresponding tweets, ordered by negativity. The topics are colour-coded: red for Politics, green for COVID, yellow for Religion, purple for Football and blue for Generic.

Index	Italy (66.75%)		Brazil (64.58%)		Netherlands (54.55%)	
	Topic	Negativity	Topic	Negativity	Topic	Negativity
1	peggio	95.33	putin	85.59	gemist	55.59
2	odio	90.09	perdi	82.11	coronavirus	45.90
3	virus	83.03	lava	68.50	haag	45.15
4	salvini	81.18	jesus	59.87	hond	41.87
5	renzi	80.94	máscara	59.01	trein	37.23
6	democrazia	60.09	verdades	49.71	biden	36.02
7	coronavirus	59.40	gasolina	48.55	lachen	34.38
8	papa	54.53	menino	47.64	slapen	34.31
9	calcio	53.02	meme	41.45	nope	29.83
10	nomi	51.67	gato	40.71	vakantie	25.66
11	pizza	47.01	cabelo	37.08	facebook	25.46
12	amazon	43.00	festa	36.05	filmpje	23.18
13	concordo	38.53	fã	35.97	gold	19.07
14	dibattito	33.88	facebook	35.67	seizoen	16.92
15	caffè	32.23	barato	35.27	fotos	16.88
16	facebook	30.36	netflix	34.91	interviews	16.75
17	natale	28.42	pizza	32.03	koffie	16.13
18	thread	26.81	artista	30.95	aflevering	11.59
19	sogno	24.39	natal	22.58	anniversary	10.16
20	serata	11.38	dm	16.74	verjaardag	8.72

Table 9. Mean user and tweet negativities, by category, of the top 20 topics in the 3 novel Generic Users datasets, ordered by negativity. The categories are colour-coded: red for Politics, green for COVID, yellow for Religion, purple for Football and blue for Generic.

Italy		Brazil		Netherlands	
User	Tweet	User	Tweet	User	Tweet
84.32	74.07	79.48	67.55	62.73	45.90
74.17	71.22	75.53	59.87	54.74	40.59
72.99	54.53	72.01	59.01	52.64	26.10
70.84	53.02	70.87	38.24	-	-
70.25	39.91	-	-	-	-

Finally, the categories of topics were grouped together and their mean values were calculated in terms of both tweet negativity (Table 8) and user negativity (Table 7). These values are displayed in Table 9.

Similar to the previous results, Politics is the most negative overall and is the most negative for all columns except the Dutch Tweet negativity, where it is beaten by COVID. Religion is usually the second most negative for the two datasets in which it appears, Italy and Brazil, and it is followed by Football in the former. COVID is the least consistent, appearing in 4 of the 5 possible positions.

Finally, Generic is the least negative, appearing in the least negative position for all columns except the Italian user negativity, where it appears above COVID.

As these are the means of other values previously discussed in detail, it is unsurprising that they do not provide any additional results. Their main value is to confirm the overall negativity of very controversial topics (such as Politics), and the lower negativity of neutral topics (Generic) with respect to the more specific topics (Football, Religion, COVID).

5 CONCLUSION

Overall, this paper provides a thorough analysis of relationship negativities between cultures and topics in a systemised manner using 26 datasets, 15 of which were especially collected for this task. These datasets were specifically gathered to obtain a more precise understanding of previously observed phenomena. This paper also builds on previous works to provide a significantly improved understanding of how negativities can be affected at each level of the SENM as well as at the level of the overall network. The results provide further support for some previous hypotheses (e.g. more specific or polarising topics lead to greater negativities at all levels of the SENM) and redefine others (e.g. how easily cultural effects are overwhelmed by topics).

Specifically, the main take-home messages of the analysis presented in this paper are: (i) there appear to be some differences in negativity due to both culture and topic. However, cultural differences do not seem to get “overwhelmed” by topics as easily as previously thought, rather, the impact of culture on negativity is important in the majority of cases and is only overpowered by *extremely controversial* topics; (ii) networks centred around a specific topic display a dramatic increase in negativity across all cultures and this effect is stronger the more negative or controversial the topic; (iii) the stability of negativity across the different circles of the SENM does not seem to decrease for more negative or specific topics, as was found previously; (iv) the number of controversial topics that appear among those most talked about within a dataset is an excellent indicator of how negative (in terms of relationships) that dataset is, and this is true even when tweets relating to those topics aren’t overly negative themselves.

As with any piece of research, there are some important limitations to bear in mind while considering the results. For instance, it would be impossible to analyse all countries, topics and types of users, and the choice of which ones were included will influence how much can be gleaned from the results. However, as specific attention was made to collect data from a range of culturally diverse regions, as well as topics at different levels of polarity, it can reasonably be expected that these results are as accurate as feasibly possible.

ACKNOWLEDGMENTS

This work is partially supported by the European Union – Horizon 2020 Program under the scheme “INFRAIA-01-2018-2019 – Integrating Activities for Advanced Communities”, Grant Agreement n.871042, “SoBigData++: European Integrated Infrastructure for Social Mining and Big Data Analytics” (<http://www.sobigdata.eu>). This work is partially supported by SoBigData.it, which receives funding from European Union – NextGenerationEU – National Recovery and Resilience Plan (Piano Nazionale di Ripresa e Resilienza, PNRR) – Project: “SoBigData.it – Strengthening the Italian RI for Social Mining and Big Data Analytics” – Prot. IR0000013 – Avviso n. 3264 del 28/12/2021. C. Boldrini was supported by PNRR - M4C2 - Investimento 1.4, Centro Nazionale CN00000013 - “ICSC -National Centre for HPC, Big Data and Quantum Computing” - Spoke 6, funded by the European Commission under the NextGeneration EU programme. A. Passarella’s and M. Conti’s work was partly funded by the PNRR - M4C2 - Investimento 1.3, Partenariato Esteso PE00000013 - “FAIR”, funded by the European Commission under the NextGeneration EU programme.

REFERENCES

- [1] Valerio Arnaboldi, Marco Conti, Andrea Passarella, and Robin I. M. Dunbar. 2017. Online social networks and information diffusion: The role of ego networks. *Online Social Networks and Media* 1 (2017), 44–55.
- [2] Valerio Arnaboldi, Marco Conti, Andrea Passarella, and Fabio Pezzoni. 2013. Ego networks in twitter: an experimental analysis. In *Proceedings IEEE INFOCOM*. 3459–3464.
- [3] Valerio Arnaboldi, Andrea Passarella, Marco Conti, and Robin IM Dunbar. 2015. *Online social networks: human cognitive constraints in Facebook and Twitter personal graphs*. Elsevier, The Netherlands.
- [4] Francesco Barbieri, Luis Espinosa Anke, and Jose Camacho-Collados. 2022. Xlm-t: Multilingual language models in twitter for sentiment analysis and beyond. In *Proceedings of the Thirteenth Language Resources and Evaluation Conference*. 258–266.
- [5] Chiara Boldrini, Mustafa Toprak, Marco Conti, and Andrea Passarella. 2018. Twitter and the press: an ego-centred analysis. In *Companion Proceedings of the The Web Conference 2018*. 1471–1478.
- [6] Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Unsupervised cross-lingual representation learning at scale. *ICLR 2019* (2019).
- [7] Corinna Cortes and Vladimir Vapnik. 1995. Support-vector networks. *Machine learning* 20 (1995), 273–297.
- [8] Adji B Dieng, Francisco JR Ruiz, and David M Blei. 2020. Topic modeling in embedding spaces. *Transactions of the Association for Computational Linguistics* 8 (2020), 439–453.
- [9] Robin IM Dunbar. 1992. Neocortex size as a constraint on group size in primates. *Journal of human evolution* 22, 6 (1992), 469–493.
- [10] Robin IM Dunbar. 1998. The social brain hypothesis. *Evolutionary Anthropology: Issues, News, and Reviews: Issues, News, and Reviews* 6, 5 (1998), 178–190.
- [11] Robin I. M. Dunbar. 1993. Coevolution of neocortical size, group size and language in humans. *Behavioral and brain sciences* 16, 4 (1993), 681–694.
- [12] Robin I. M. Dunbar, Valerio Arnaboldi, Marco Conti, and Andrea Passarella. 2015. The structure of online social networks mirrors those in the offline world. *Social networks* 43 (2015), 39–47.
- [13] Robin I. M. Dunbar and Matt Spoor. 1995. Social networks, support cliques, and kinship. *Human nature* 6 (1995), 273–290.
- [14] Keinosuke Fukunaga and Larry Hostetler. 1975. The estimation of the gradient of a density function, with applications in pattern recognition. *IEEE Transactions on information theory* 21, 1 (1975), 32–40.
- [15] Eric Gilbert and Karrie Karahalios. 2009. Predicting tie strength with social media. In *Proceedings of the SIGCHI conference on human factors in computing systems*. 211–220.
- [16] John Gottman, John Mordechai Gottman, and Nan Silver. 1995. *Why marriages succeed or fail: And how you can make yours last*. Simon and Schuster.
- [17] Mark S Granovetter. 1973. The strength of weak ties. *American journal of sociology* 78, 6 (1973), 1360–1380.
- [18] Maarten Grootendorst. 2022. BERTopic: Neural topic modeling with a class-based TF-IDF procedure. *arXiv preprint arXiv:2203.05794* (2022).
- [19] Betty Hart and Todd R Risley. 1995. *Meaningful differences in the everyday experience of young American children*. Paul H Brookes Publishing.
- [20] Russell A Hill and Robin IM Dunbar. 2003. Social network size in humans. *Human nature* 14, 1 (2003), 53–72.
- [21] Jey Han Lau, David Newman, and Timothy Baldwin. 2014. Machine reading tea leaves: Automatically evaluating topic coherence and topic model quality. In *Proceedings of the 14th Conference of the European Chapter of the Association for Computational Linguistics*. 530–539.
- [22] Giovanna Miritello, Esteban Moro, Rubén Lara, Rocío Martínez-López, John Belchamber, Sam GB Roberts, and Robin IM Dunbar. 2013. Time as a limited resource: Communication strategy in mobile phone networks. *Social networks* 35, 1 (2013), 89–95.
- [23] Nils Reimers and Iryna Gurevych. 2019. Sentence-bert: Sentence embeddings using siamese bert-networks. *arXiv preprint arXiv:1908.10084* (2019).
- [24] Nils Reimers and Iryna Gurevych. 2020. Making monolingual sentence embeddings multilingual using knowledge distillation. *arXiv preprint arXiv:2004.09813* (2020).
- [25] Alistair G Sutcliffe, Di Wang, and Robin IM Dunbar. 2015. Modelling the role of trust in social relationships. *ACM Transactions on Internet Technology (TOIT)* 15, 4 (2015), 1–24.
- [26] Jack Tacchi, Chiara Boldrini, Andrea Passarella, and Marco Conti. 2022. Signed ego network model and its application to Twitter. *IEEE BigData 2022* (2022).
- [27] Jack Tacchi, Chiara Boldrini, Andrea Passarella, and Marco Conti. 2023. Cultural Differences in Signed Ego Networks on Twitter: An Investigatory Analysis. In *Companion Proceedings of the ACM Web Conference 2023*. 1039–1049.

- [28] Jack Tacchi, Chiara Boldrini, Andrea Passarella, and Marco Conti. 2024. Keep Your Friends Close, and Your Enemies Closer: Structural Properties of Negative Relationships on Twitter. arXiv:2401.16562 [cs.SI]
- [29] Mustafa Toprak, Chiara Boldrini, Andrea Passarella, and Marco Conti. 2021. Structural Models of Human Social Interactions in Online Smart Communities: the Case of Region-based Journalists on Twitter. *Online Social Networks and Media* 30 (2021).
- [30] Mustafa Toprak, Chiara Boldrini, Andrea Passarella, and Marco Conti. 2022. Journalists' ego networks in Twitter: Invariant and distinctive structural features. *Online Social Networks and Media* 30 (2022), 100207.

Received 31 January 2024