

# Differences in Social Media Usage Exist Between Western and Middle-East Countries

Jens Helge Reelfs<sup>1</sup>, Oliver Hohlfeld<sup>1</sup>, and Niklas Henckell<sup>2</sup>

<sup>1</sup> Brandenburg University of Technology, Cottbus, DE  
{lastname}@b-tu.de

<sup>2</sup> The Jodel Venture GmbH, Berlin, DE  
niklas@jodel.com

**Abstract.** In this paper, we empirically analyze two examples of a Western (DE) versus Middle-East (SA) Online Social Messaging App. By focusing on the system interactions over time in comparison, we identify inherent differences in user engagement. We take a deep dive and shed light onto differences in user attention shifts and showcase their structural implications to the user experience. Our main findings show that in comparison to the German counterparts, the Saudi communities prefer creating content in longer conversations, while voting more conservative.

## 1 Introduction

Every social networking platform depends on an active user-base. The recent pandemic disruption has even accelerated the shift of our everyday’s life into digital spaces, may it be at work or socially. What constitutes an active user-base and if cultural differences exist in this usage behavior, however, is unclear.

Social network analysis is an active field of research for more than a decade. Research provided a general understanding through the empirical and qualitative analyses of a number of different networks. Examples include structural measurements of classic online social networks [23,24,29,11] as well as more specialized variants such as microblogging [4], picture sharing [33,6], or knowledge sharing [35]. Works in this field analyzed the networks’ *structure*, mostly using graph-theory approaches. This way, they showed that social networks usually creates small-world networks [20,8]. The influence of cultural or geographic backgrounds on usage largely remain unknown.

Most studies either focus on analyzing social-media usage *worldwide* or by focusing on specific parts of the world, mostly English speaking. These works have enriched our understanding of social media. Yet it is unclear if or to what extent the obtained measures differ between different geographic regions. Cultural differences are known to exist that drive human behavior in social networks, e.g., the degree of connectivity [15] or how marketers use social media to impact purchase decisions [9]. Yet, little is known on how geographic or cultural backgrounds may impact the way users interact with a social media platform in terms of the generated traffic; that is, content creation and content voting.

In this work, we take the rare chance to analyze ground truth information provided by a social network operator to compare interactions with a social media

platform of a Western (Germany) and a Middle-Eastern (Saudi Arabia) country. We selected both countries since they represent the largest user-bases of the Jodel platform, while simultaneously representing a largely different (cultural) background. Our data sets capture the entirety of all social media interactions in both countries since the very first post. This way, we can, for the first time, shed light on whether geographic or culture specific differences exist between both countries w.r.t. how the user-bases generate and vote content.

The studied social media platform Jodel is location-based and anonymous. Most importantly, the feature of Jodel to form independent local communities enables us to compare in-country and between country effects and thereby to clearly identify country specific usage differences. Further, it does not display any form of user profiles or other user-related information that would introduce visible social credit; users solely interact framed into their physical proximity and based on their topic preferences. This results in a *pure* form of communication that is reduced to content, since any form of influence by user profiles such as social status is absent. This makes Jodel an ideal platform to study differences in content creation and voting, i.e., the entirety of active interactions with this social network. We shed light on fundamentally different user behavior and engagement patterns within such anonymous spaces having received less attention as of today, across the Kingdom of Saudi Arabia and Germany; and showcase structural implications. Our contributions are as follows.

- While not a primary focus of our work, we empirically show the very different adoption processes of a new social media platform in both countries.
- We show that, invariant to time and community size, users in Saudi-Arabia (Middle-East) behave fundamentally different to the German counterparts (Western country). They prefer creating content, but vote slightly less than the German users. This highlights, for the first time, that country-level differences in the usage culture in social media exist that create drastic differences in user behavior of the very same social media platform.
- We exemplify the implications of shown user engagement to the social media platform. E.g., we show that the availability of more content in Saudi Arabia naturally decreases the available votes per post, which can have serious impact on, e.g., distributed community moderation techniques.

## 2 The Jodel App

Jodel is a mobile-only messaging application which we show in Fig. 1. It is location-based and establishes local communities relative to the users' location ❶. Within these communities, users can *anonymously* post both images and textual content of up to 250 characters length ❷ (i.e., microblogging) and reply to posts forming discussion threads ❸. Posted content is referred to as “Jodels”, colored randomly ❹. Posts are only displayed to other users within close (up to ≈20km) geographic proximity ❺. Further, all communication is *anonymous* to other users since no user handles or other user-related information are displayed. Only *within* a single discussion thread, users are enumerated to enable referencing

to other users. Threads are displayed to the users in three different feeds **1** :  
i) *recent* showing the most recent threads, ii) *most discussed* showing the most discussed threads, and iii) *loudest* showing threads with the highest voting score.

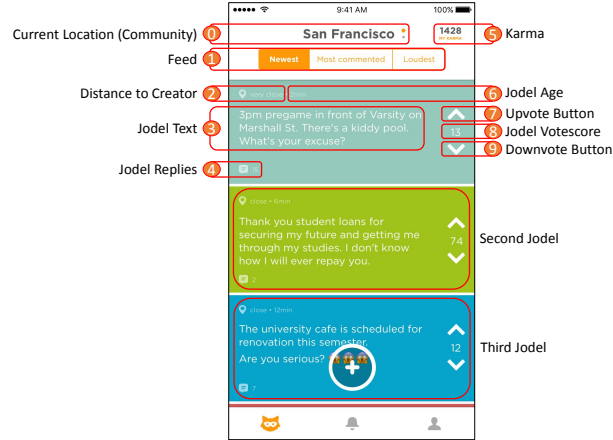


Fig. 1: **Jodel iOS mobile application.**

Jodel employs a community-driven filtering and moderation scheme to avoid adverse content. For an anonymous messaging app, community moderation is a key success parameter to prevent harmful or abusive content. In Jodel, content filtering relies on a distributed voting scheme in which every user can increase or decrease a post’s vote score by up- (+1) **7** or downvoting (-1) **9** (similar to StackOverflow). Posts reaching a cumulative vote score **8** below a negative threshold (e.g., -5) are no longer displayed. Depending on the number of vote-contributions, this scheme filters out adverse content while also potentially preferring mainstream content. To increase user engagement w.r.t. posting and voting, Jodel uses lightweight gamification by awarding *Karma* points **5**.

### 3 Dataset Description and Ethics

The Jodel network operators provided us with *anonymized* ground truth data of their network, cf. Table 1. The obtained data contains post and interaction *metadata* created within Germany and the KSA only, and spans over multiple years from the beginning of the network in 2014 up to August 2017. It is limited to metadata only without textual content and user records stripped and anonymized. The data enables us to cluster users by their anonymous ID. Further, it contains no personal information and cannot be used to personally identify users. While users consent to research via the Jodel ToS, we also inform and synchronize with the Jodel operator about evaluations we perform on their data. The structure of our available dataset includes 3 categories: interactions, content, and users.

Type	#SA	#DE	Description
User	1.2M	3.6M	User metadata
Content	469M	285M	Content (posts, replies)
Interaction	961M	3.0G	Interactions incl. user, community and type (post, reply, up-/downvoted)

Table 1: **Dataset statistics.** The data ranges from the application start in late 2014 up to the beginning of August 2017. We find the first observation in the KSA in December 2014.

**Dataset limitations.** Our dataset only includes the users’ *active* interactions with the system, where they contribute like registering, creating posts, replying, or voting. Thus, we cannot infer when or how much a user only *passively* participates—lurkers—who only consume content. Further, the vote interactions are always mapped to the date and community of the respective content creation. This prevents us from making detailed analyses depending on the voting time or place. However, due to the vivid usage of the application (multiple posts/replies per minute), we generally consider votes to be executed on the same day as their respective content. Especially since posts are only accessible via the three different feeds, where they will only stay for a very limited time, casting votes long after the content creation is usually not possible.

## 4 The Birth of the Jodel Networks in DE and the KSA

The growth patterns of social networks are less understood given that data captures from the very beginning of a social media platform are typically unavailable. We take the rare chance to begin showcasing the rate in which the Jodel platform established itself in both countries. Our first peek is relevant to better understand network activity and to define a meaningful aggregate layer for comparison, e.g., time slices, for studying cultural shifts in social media usage in the next section. We thus remark that a detailed study of the shown adoption pattern is beyond the scope of this paper.

### 4.1 Different Adoption Pattern in Germany and the KSA

We show the adoption of the Jodel network in both networks by the number of interactions over time in Figure 2. The figure shows the number of weekly interactions for Germany (solid line) and the KSA (dotted) since the very first interaction till the end of our data set in August 2017. With interaction, we refer to any interaction with the Jodel system, i.e., either posting, replying, or voting.

The adoption of Jodel in Germany is characterized by a slow but rather steady growth of network activity over time, peaking in 2016/2017. This captures the birth of the Jodel network that originated in Germany and then constantly increased in popularity. In contrast to the steady increasing activity in Germany, the adoption of Jodel in the KSA is characterized by a substantial influx of users and an increase in activity at a short time in March 2017. To our best knowledge,




Fig. 2: Jodel activity in Germany (DE) and the Kingdom of Saudi Arabia (SA) over our observation time. While the DE network enjoys a steady growth over time, the usage in the KSA has suddenly surged and remains stable.

the reason for this behavior is that Jodel went viral via social media in Saudi Arabia—in absence of any marketing campaigns of the operator itself.

While studying the reasons that were driving these adoption processes is beyond the scope of this paper, their adoption processes differ substantially. That is, referring to Table 1, in only 4 months, Jodel KSA has roughly gathered 1.2 million users, while Jodel Germany over six times longer time period accumulates to 3.6 million. Likewise, the amount of interactions equally scales between the KSA with 1 billion and Germany with 3 billion interactions. This observation and differences in adoption allow for putting aggregates, e.g., comparable time slices, for our study into perspective.

**Findings.** *Adoption pattern and thus associated traffic can differ substantially.*

#### 4.2 Different Adoption Pattern Require Comparable Time Slices

To compare the social media usage in both countries, we need to derive comparable datasets. That is, we need to enable comparisons between two different populations of posts, interactions, and users across a multitude of communities. We thus identify two main dynamics within the lifecycle of communities defining the aggregation dimension: *i)* time and *ii)* per community interaction volume.

**Slicing by i) Time.** As we are limited in the length of observations, especially for Saudi Arabia, we selected half-year long timeframes backwards from the end of observation. We show these timeframes with vertical bars in Figure 2 named by country  $DE_0, DE_1, \dots, DE_5, SA_0, SA_1$ , where the index indicates the partition’s age. This simple batching approach creates half-year slices that represent various stages within the community development. In our study, we compare these slices independently in each evaluation to account for the different adoption processes. We have experimented with higher resolutions to enrich our results with more data points (not shown), but our conclusions remain the same for our presented period length. We handle these partitions independently of each other, i.e., early-day users from  $DE_5$  may drop out of the statistics in subsequent partitions due to a lack of interactions.

**Slicing by ii) Community Interaction Volume.** Defining “a” community is not possible on Jodel given that content is displayed relative to the users location and thus differs from user to user. That is, every user might experience

a slightly different community to interact with, which cannot be reconstructed from the data. To solve this, we assign each interaction to a nearby major city or district, which generates clusters of interactions that we refer to as communities. This discretization generates an approximation of the individually experienced communities. The resulting approximation is of sufficient accuracy to study and compare the Jodel usage in different parts of the respective countries. Further, the discretization does not normalize for covered area, nor covered population.

We mitigate these inherently hard problems in normalization by simplifying our partitioning approach. By slicing all interactions into quantiles ordered by their corresponding community size, we enable a relative comparison; named *is\_q0\_25*, ... representing the corresponding quantile of all interactions, discretized into communities (leading to an approximation). We provide details of this partitioning in Table 2 describing the amount of discretized communities per country. That is, e.g., the single largest SA community is the capital Riyadh at about 30% total interaction volume—hence, it is the only community within the set of *q75\_100*. Due to dividing the interaction volume into equal parts, we encounter a heavy-tailed interaction distribution across communities resulting in only few entries within the upper quantiles; the German community size distribution qualitatively matches the SA counterpart, while the latter is largely shifted in volume within magnitudes of fewer communities (not shown).

We will compare the social media usage based on the resulting data sets.


Interactions quantile	#communities SA	#communities DE
$\approx 75..100$	1	14
$\approx 50..75$	4	29
$\approx 25..50$	12	114
$\approx 0..25$	78	6,678

Table 2: Interaction volume aggregation layer & amount of corresponding communities. Due to the heavy tailed distribution across the community discretization, the upper quantiles consist of fewer communities.

## 5 Geographic Differences in Jodel Usage: DE vs. KSA

Is there a systematic difference between Jodel users in Germany and in the KSA in the way they use the social media platform? In other words, do culture or country specific usage behaviors exist that uniquely define traffic profiles of the very same platform in each country? While social media usage has widely been studied, the question of in-platform variation and behavioral differences is still open. In this section, we set out shedding light on this aspect by comparing the Jodel usage in two countries with a different social and cultural background.

We study the question of (cultural) differences w.r.t. user behavior in Jodel usage by investigating differences in active user *interactions* with Jodel, i.e., posting and voting, accounting for all possible active system interactions. We base this evaluation on two factors introduced in Section 4.2: partitioning by *i*) time and *ii*) interaction type, which we discuss next.



(a) Total per user interactions per time period. Per user interaction volume is very similar over time.

(b) Per user interactions by time period (DE0, ...) and interaction type. Over time, we observe increasing engagement. There is a substantial difference in content creation and upvote activities between SA and DE.


Fig. 3: Differences in per-user interactions by time and country. *a)* While the individual users behave similar across the board w.r.t. interaction volume, *b)* attention shifts towards content creation for the SA user groups, whereas downvoting is less popular in comparison to DE.

**Overall per-user activity is country-independent.** First, we subdivide all interactions into independent half-year periods as described in Section 4.2 (DE0, ..., DE5, SA0, SA1). These independent periods enable to compare the behavior of the networks at different times to account for differences in the adoption of Jodel in DE and SA. We then aggregate the interactions to each user and present the resulting interactions per user CDF in Figure 3a. We show each half-year period as separate CDF (series) for each of the two countries. Invariant to time, we observe quite similar heavy tailed distributions; that is, most users are not very active, such that e.g., 60% of all users each have up to only 100 interactions. The distribution for SA1 deviates from the pattern since it captures a timeframe before being popular (considerably fewer data points). In general, we observe that irrespective of time and country, users follow a similar usage behavior—also in absolute terms (not shown).

**Difference: posting vs. voting.** Next, we further partition the data by the type of interaction in addition to the time slices used before. That is, we show distributions of interactions per user subdivided into the voting interactions (upvoted, downvoted) and content creation interactions (postcreated, replycreated) as a box plot in Figure 3b. Note the logarithmic y-axis. Further, the whiskers denote the 5%/95% percentile.

German users tend to increase their engagement over time at increasing platform activity regardless of the interaction type. While upvoting is the most prominent type of interaction for the German users, voting content down and replying to content are roughly equally less prominent. The SA users prefer content creation, especially replying, whereas upvoting happens less frequently.

Remarkably, all time periods within a country are determined by similar behavior. In other words, posting content is the dominant type of interaction in the KSA, while it is voting in Germany. This represents a clear difference in



(a) Accumulate vote scores CDF over time. (b) Vote interaction distributions by time period (DE0, ...) split by community size (q0\_25,...) and being a thread start (thread) or is a reply buried within a thread.

Fig. 4: Platform accumulate vote scores and vote interactions per post. *a)* Larger communities experience a stronger heavy tail in vote scores. *b)* Due to their exposure, threads usually collect more votes than replies buried within them.

platform usage that can be observed between these two countries (also regardless of community size, not shown).

The ratio of up to overall votes remains positive at a happyratio (upvotes to total votes) of 83% for DE and 71% for SA. The figures for the SA1 partition need to be taken with a grain of salt due to only few users; however, the engagement spread is higher compared to the latest timeframe SA0.

**Findings.** *Invariant to time and community size, the SA users (Middle-East) behave fundamentally different to the DE counterparts (Western country). They heavily prefer creating content, but vote slightly less than the German users. This highlights, for the first time, that cultural patterns in social media user behavior exist that create drastic shifts in how a very same social media platform is used in each country. This finding may be considered even more interesting, given Jodel being an anonymous platform that enables a very pure form of communication; it entirely focuses on posted content in absence of any user profile.*


## 6 Structural Implications

With the identified fundamental attention shift in user behavior between SA and DE users in content creation and voting, we now aim at studying resulting implications on the platform. According to the operator, the communities in both countries are considered to be well-functioning. That is, regardless of implications arising from different usage profiles, participants are enjoying spending time on the platform to the most part; e.g., by creating content, voting, or just lurking.

### 6.1 Content Voting

**Accumulated votes overview.** Votes on posted content have two roles in Jodel: *i)* they show content appreciation to others (e.g., enable users to sort by popularity) and *ii)* enable distributed user based content moderation that





(a) Thread w/o replies over time and response ratio. Most threads can attract replies.

(b) Thread length over time by community size. With increasing engagement (*left*, Fig. 5a), communities experience longer threads on average invariant to community size.

Fig. 5: Thread engagement with responses - how many & how much. *a)* At an increasing ratio over time, within recent recent time periods, most threads attract replies. *b)* With growing amounts of interactions over time, the average thread length increases to a similar level in DE and SA.

removes content with negative vote scores (see § 2). Factors influencing vote distributions can thus have structural implications on the platform. To study differences in voting behavior, we first take a look into accumulate vote score ( $\#upvotes - \#downvotes$ ) distributions. There are two given bounds for posts gathering votes: 1) Posts beyond a negative threshold are no more displayed on the platform. 2) There is no conscious upper bound given by the system, yet posts are only temporarily displayed within the various app feeds (see § 2) and therefore, the time for interacting with them is inherently limited. Given these constraints, all communities naturally enjoy a rather positive mood. To put an emphasis on the temporal dimension, we show the CDF of vote scores to posts over time (DE0, ...) in Figure 4a. In the earlier times of DE with less activity, more posts were able to gather more votes as illustrated by the DE3.5 series resembling a broader distribution. There is a slight decrease of accumulated vote scores throughout time, hence interactions per post, for the DE communities—also observed within the KSA. The distributions become more and more long-tailed over time and appear scale-shifted. Noteworthy, a split into community sizes confirms this finding: larger communities may reach far higher absolute scores, but the distributions become more skewed correlated to the observed interaction volume (not shown).

Due to SA users producing much more content, the feeds displayed in the app also get renewed completely at a very high pace. Thus, SA posts compete harder for time to collect possible votes in comparison to their DE cousins; the feeds also promote observed long-tails. What implications does this shift have on experienced vote distributions?

**Votes per Post.** To better understand the voting interactions and the observed skew in accumulated vote scores, we next normalize observed figures to a per-post basis. The box plots in Figure 4b show various per post vote interaction

distributions across time (DE0, ...), and community size (q0\_25, ...), while further distinguishing between threads (thread) and replies.

We find the long tail of high vote scores in the rather long 95% percentile whiskers on the log scaled y-axis. Invariant to time and community size, the median German user enjoys voting on threads with median levels around 10 to 30 votes gathered by each post throughout time. Naturally given by the communication structure and app design, content buried within threads is much less appreciated; they accumulate only two to three votes in DE. As discussed before, the SA posts stand in stark contrast at three to four votes within the main Q0 timeframe. I.e., opposed to German users, the average participant within the Kingdom of Saudi Arabia cannot expect to receive any vote on her content—especially and naturally not on replies.

**Findings.** *We show that the availability of more posted content in Saudi Arabia decreases the available votes per post, which can influence community moderation techniques that depend on voting.*

## 6.2 Spinning Faster: Response Time and Volume

While voting or liking is a vital part of a social network, it can only exist because of posted content and replies. We thus next study geographic properties that influence the response time and volume.


In Figure 5a, we show the amounts of posts with and without replies (bars) and the ratio (lines) across time for both countries. The German communities increase their response cover over time, while it instantly is equally high for the SA communities at about 90%. I.e., 9 out of 10 users in both countries can expect getting at least a single reply on a thread.

**Response volume.** As most users receive a reply, does the total achieved thread length correlate with community size, and how does this interplay with the distribution shift in content creation? We answer this question with distributions given in Fig. 5b as a box plot, which depicts the thread length gathered per post across time (DE0, ...) and community sizes (is\_q0\_25, ...).

First, the 95% percentile whiskers indicate a long-tailed distribution in the length of threads, which we confirm (now shown). Second, the amount of replies is invariant to community size as the distributions are very similar; however, there still exists a huge spread from the 75% to the 95% percentile (whiskers) due to the long-tailed distribution. Second, we observe an increasing trend over time. This increasing engagement is also apparent when looking deeper into the interactions (cf. Section 4.2 and alike split by community size—not shown).

**Findings.** *Most posts in both countries get a reply; even at larger volumes for SA, the thread lengths are similar to DE.*

**Conversations.** Having established an understanding of the amount of replies most users experience, we get into more structural detail. We define *conversation-ness* as the ratio between replies per replier as a proxy for conversations—where lower ratios naturally depict a heterogeneous set of repliers, while higher ratios indicate fewer participants forming a back and forth conversation.



(a) Conversationness - Replies per thread participant are long-tailed and substantially larger within SA. (b) Thread engagement speed. The median timeframe between consecutive answers within a thread decreases over time in the meanwhile growing DE environment. Likewise, larger communities experience faster responses—in SA widely immediately.

Fig. 6: Thread engagement with responses - who & when. *a)* The observed overall activity increase in DE results in longer conversations over time, the SA users talk substantially longer. *b)* The average time until receiving a response reduces to only few minutes in both countries, SA still takes the lead.

We present the distributions of this ratio over time for both countries in Figure 6a as box plots; the 95% whisker indicates long-tailed distributions, which we can confirm (not shown). Over time and with increasing network activity, all German communities increase up to about 1.6 replies per thread participant within the 75% quantile in DE0. This indicates a shift from rather random single comments becoming less popular in favor of interacting with each other. Given the high preference on creating content and vividly replying, this trend is particularly apparent and reinforced in the SA communities at rates of up to two replies for the 75% quantile of SA0.

**Findings.** *Employing a conversationness metric, we identify SA users to be more conversational compared to DE.*

**Response time.** We have seen that most threads receive at least some replies. While the counts in responses may matter quantitatively, we also want to shed light on the time-dependent dynamics of the reply interactions. Fig. 6b shows a box plot of the distributions of the time between consecutive responses within a thread split by time (DE0, ...) and community size (q0.25, ...); note the log y-scale. Unfortunately, our dataset does not allow for this evaluation on vote interactions (see § 3).

From this evaluation, we gain two major insights: 1) throughout time with increasing activity and engagement, the German communities establish shorter response times down to only minutes. Having reached a sustainable community size, the response times no longer drop. 2) The SA communities instantly drop response times substantially below the German counterpart to only a single minute within most threads. Note: High response times within small communities in SA1 are due to small amounts of data; Missing series indicate no present data.

**Findings.** *In comparison to DE, the SA communities are more vividly responsive as the average response times are considerably lower.*

## 7 Related Work

The research community established a rich field of understanding human interaction within social media, yet not studying geographic differences in social media usage. Empirical studies on social media focused on the birth and growth [29,22,27], social media usage in specific regions such as the Arab Gulf states [27] specifically focused to global usage [17], information propagation [6], specific platforms such as Facebook [18,24], YouTube [5], SnapChat [33], or Twitter [13]. Such research tries to understand and identify social structures and influence [11,31]. Mathematical modelling [32] and graph methods are common techniques to analyze social ties [14,19]. Platforms may also have rather adverse effects like cyberbullying [36,12,10], or may raise privacy concerns [30].

A recent body of research aims at understanding anonymous social networks. The desire for anonymity can result in throwaway accounts [16] and can also manifest in anonymous self disclosures [2]. Anonymous content platforms have been detailed w.r.t. content [25] and user behavior [1,26,7]. Other empirical work focuses specifically on location based anonymous platforms, e.g., Whisper [34] or Yik Yak [21,28], or analyzes its local content [3].

We complement these works providing a new unique view on the lifecycle of various Jodel communities in Germany and Saudi Arabia with a special focus on happening interactions, differences and resulting platform implications.

## 8 Conclusions

In this paper, we show that the usage behavior of users in Germany (DE) fundamentally differs from users in Saudi Arabia (KSA) in the anonymous and location based Jodel network. This study is enabled by the feature of Jodel to form independent local communities enables us to compare in-country and between country effects and thereby to clearly identify country specific usage differences. We empirically characterizes usage behavior based on ground truth user interaction data provided by the operator. While we can rule out marketing effects by the operator, our findings motivate future work that study root causes.

We find that, independent of time and community size, KSA users prefer content creation (posting & responding), while German users tend to interact slightly more passively (voting). Other than this shift towards content, due to the users in both regions else behaving identical on a per-user measure; we find rather identical community engagement. However, due to much more content being available within SA per user, posts compete harder in gathering votes than the German counterparts, which can have implications for vote-based content moderation schemes. Further, the average number of replies also does not increase in comparison; Still, reply times are much smaller due to higher activity. The average Saudi user tends towards having longer conversations. Overall, we identify time- and geographic-invariant differences between DE and SA user engagement as the latter substantially focus on creating content, giving a slight lead in voting to the Germans. This provides a new interaction-based perspective on geographic difference of social media usage that have not yet been studied.

## References

1. Bernstein, M.S., Monroy-Hernández, A., Harry, D., André, P., Panovich, K., Vargas, G.G.: 4chan and/b: An analysis of anonymity and ephemerality in a large online community. In: ICWSM (2011)
2. Birnholtz, J., Merola, N.A.R., Paul, A.: Is it weird to still be a virgin: Anonymous, locally targeted questions on facebook confession boards. In: SIGCHI (2015)
3. Black, E.W., Mezzina, K., Thompson, L.A.: Anonymous social media—understanding the content and context of yik yak. SIGCHI (2016)
4. Bollen, J., Mao, H., Pepe, A.: Modeling public mood and emotion: Twitter sentiment and socio-economic phenomena. ICWSM (2011)
5. Brodersen, A., Scellato, S., Wattenhofer, M.: Youtube around the world: geographic popularity of videos. In: WWW. pp. 241–250 (2012)
6. Cha, M., Mislove, A., Gummadi, K.P.: A measurement-driven analysis of information propagation in the flickr social network. In: WWW (2009)
7. Correa, D., Silva, L.A., Mondal, M., Benevenuto, F., Gummadi, K.P.: The many shades of anonymity: Characterizing anonymous social media content. In: ICWSM (2015)
8. Freeman, L.: The development of social network analysis. *A Study in the Sociology of Science* (2004)
9. Goodrich, K., de Mooij, M.: How social are social media? a cross-cultural comparison of online and offline purchase decision influences. *Journal of Marketing Communications* **20**(1-2), 103–116 (2014). <https://doi.org/10.1080/13527266.2013.797773>
10. Hosseinmardi, H., Ghasemianlangroodi, A., Han, R., Lv, Q., Mishra, S.: Towards understanding cyberbullying behavior in a semi-anonymous social network. In: ASONAM (2014)
11. Kairam, S., Brzozowski, M., Huffaker, D., Chi, E.: Talking in circles: selective sharing in google+. In: SIGCHI (2012)
12. Kayes, I., Kourtellis, N., Quercia, D., Iamnitchi, A., Bonchi, F.: The social world of content abusers in community question answering. In: WWW (2015)
13. Kouloumpis, E., Wilson, T., Moore, J.D.: Twitter sentiment analysis: The good the bad and the omg! In: ICWSM (2011)
14. Kumar, R., Novak, J., Tomkins, A.: Structure and evolution of online social networks. In: KDD. ACM (2010)
15. LaRose, R., Connolly, R., Lee, H., Li, K., Hales, K.D.: Connection overload? a cross cultural study of the consequences of social media connection. *Information Systems Management* **31**(1), 59–73 (2014). <https://doi.org/10.1080/10580530.2014.854097>
16. Leavitt, A.: This is a throwaway account: Temporary technical identities and perceptions of anonymity in a massive online community. In: CSCW (2015)
17. Leskovec, J., Horvitz, E.: Planetary-scale views on a large instant-messaging network. In: WWW (2008)
18. Lewis, K., Kaufman, J., Gonzalez, M., Wimmer, A., Christakis, N.: Tastes, ties, and time: A new social network dataset using facebook. com. *Social networks* **30**(4), 330–342 (2008)
19. Magno, G., Comarela, G., Saez-Trumper, D., Cha, M., Almeida, V.: New kid on the block: Exploring the google+ social graph. In: IMC (2012)
20. Manku, G.S., Naor, M., Wieder, U.: Know thy neighbor’s neighbor: the power of lookahead in randomized p2p networks. In: STOC (2004)
21. McKenzie, G., Adams, B., Janowicz, K.: Of oxen and birds: Is yik yak a useful new data source in the geosocial zoo or just another twitter? In: SIGSPATIAL (2015)

22. Mislove, A., Koppula, H.S., Gummadi, K.P., Druschel, P., Bhattacharjee, B.: Growth of the flickr social network. In: WOSN (2008)
23. Mislove, A., Marcon, M., Gummadi, K.P., Druschel, P., et al.: Measurement and analysis of online social networks. In: IMC (2007)
24. Nazir, A., Raza, S., Chuah, C.N.: Unveiling facebook: a measurement study of social network based applications. In: IMC (2008)
25. Papasavva, A., Zannettou, S., De Cristofaro, E., Stringhini, G., Blackburn, J.: Raiders of the lost kek: 3.5 years of augmented 4chan posts from the politically incorrect board. In: Proceedings of the International AAAI Conference on Web and Social Media. vol. 14, pp. 885–894 (2020)
26. Reelfs, J.H., Bergmann, M., Hohlfeld, O., Henckell, N.: Understanding & predicting user lifetime with machine learning in an anonymous location-based social network. In: WWW LocWeb Workshop (2021). <https://doi.org/10.1145/3442442.3451887>
27. Reyae, S., Ahmed, A.: Growth pattern of social media usage in arab gulf states: An analytical study. Social Networking (2015)
28. Saveski, M., Chou, S., Roy, D.: Tracking the yak: An empirical study of yik yak. In: ICWSM (2016)
29. Schiöberg, D., Schneider, F., Schiöberg, H., Schmid, S., Uhlig, S., Feldmann, A.: Tracing the birth of an osn: Social graph and profile analysis in google+. In: WebSci (2012)
30. Stutzman, F., Gross, R., Acquisti, A.: Silent listeners: The evolution of privacy and disclosure on facebook. J. of privacy and confidentiality (2013)
31. Tang, J., Sun, J., Wang, C., Yang, Z.: Social influence analysis in large-scale networks. In: SIGKDD (2009)
32. Van Mieghem, P., Blenn, N., Doerr, C.: Lognormal distribution in the digg online social network. The European Physical Journal B **83**(2), 251 (2011)
33. Vaterlaus, J.M., Barnett, K., Roche, C., Young, J.A.: snapchat is more personal: An expl. study on snapchat behaviors and young adult interspers. relationships. Computers in Human Behavior (2016)
34. Wang, G., Wang, B., Wang, T., Nika, A., Zheng, H., Zhao, B.Y.: Whispers in the dark: analysis of an anonymous social network. In: IMC (2014)
35. Wang, S., Lo, D., Jiang, L.: An empirical study on developer interactions in stackoverflow. In: SAC (2013)
36. Whittaker, E., Kowalski, R.M.: Cyberbullying via social media. Journal of School Violence (2015)