(In)visible Cities: An Exploration of Social Identity, Anonymity and Location-Based Filtering on Yik Yak

Matthew Heston[1], Jeremy Birnholtz[1]
[1]Northwestern University

Abstract
The recent popularity of mobile, anonymous social interaction applications that filter content based on location (i.e., people see what others nearby have posted) has led to concerns about anonymity enabling negative behavior such as abuse or cyber-bullying. Anonymity can also have positive effects, however, and the social identity model of deindividuation effects (SIDE) suggests that people's behavior in anonymous online environments can be affected by whether they identify more with the community or more as a disconnected individual. In location-filtered apps, however, online anonymity is combined with a physically local audience. This complicates our understanding of online anonymity because identities users feel can stem both from the online and offline environments. We present an exploratory mixed-methods study of Yik Yak, a popular location-filtered app. Results suggest that people use language to invoke a geographically local group identity via shared references. People's responses to content, moreover, seem to be related to their identity as members of the online community.
Keywords: Location; anonymity; CMC; mobile applications
Contact: heston@u.northwestern.edu, jeremyb@northwestern.edu

1   Introduction

Location-aware mobile apps, enabled by widespread GPS technology, have become common (Licoppe, 2015). They let users connect with nearby friends (Guha & Birnholtz, 2013; Lindqvist, Cranshaw, Wiese, Hong, & Zimmerman, 2011), find nearby strangers for dates or other social encounters (Blackwell, Birnholtz, & Abbott, 2015; Toch & Levi, 2013) and, more recently, interact anonymously with nearby others (e.g., apps like Whisper and Yik Yak). We refer to this last category as location-filtered social apps, because content is visible primarily to others within a certain geographic radius of the poster's location.

Location-filtered social apps have received significant recent attention in the popular press because they are often seen to facilitate negative interactions, such as cyber-bullying (Binns, 2013; E. Whittaker & Kowalski, 2015). In some ways, of course, this is not surprising in light of research suggesting that anonymity can cause disinhibition, freeing people to behave negatively (S. Whittaker, Terveen, Hill, & Cherny, 2003). The freedoms of disinhibition can also be positive, however, and allow people to feel more comfortable discussing potentially sensitive topics, such as health concerns or stigmatized identities (Barak, Boniel-Nissim, & Suler, 2008; Rains, 2014; Turkle, 1995).

This tension between the positive and negative attributes of anonymity highlights aspects of online anonymity that are different when the audience is primarily geographically local. In particular, location-filtered anonymity affects the interplay of users' social identities. The Social Identity Model of Deindividuation Effects, or SIDE (Reicher, Spears, & Postmes, 1995), suggests that immersion in groups increases the salience of group identities relative to individual identities. For example, a person may think of herself more as a "support group member" than as a disconnected individual and thus behave in more pro-social ways (Dabbish, Farzan, Kraut, & Postmes, 2012; Postmes, Spears, Sakhel, & Groot, 2001).

We argue that location-filtered apps add an additional social identity, that of the local community, to this mix. In this light a recent rally at the University of Michigan to promote positive behavior on the app Yik Yak can be seen as a way to increase the salience of the "Michigan student" identity (Allen, 2015). This addition of a local social identity makes it difficult to directly apply existing models of online interaction to location-filtered social apps. Examining behavior on these apps in detail will allow us to both extend existing models to include location-aware technologies and also to understand how people are using these apps and how to better design them.

In this paper, we present a case study exploration of Yik Yak, a popular anonymous location-aware social interaction app. We examine content ("yaks") collected from 35 U.S. university campuses using quantitative and qualitative methods, and also interviewed 12 U.S. undergraduate university students who use Yik Yak. Results suggest that people use language to invoke a shared local identity,

but that these references may be implicit or explicit, and geographically or temporally local. People also draw on their identity as app users in voting for content they deem appropriate or inappropriate.

2   Background

To explore the interplay of social identities on anonymous location-filtered social apps, we focused on three issues. 1) what people use these apps for and how uses reflect social identities, 2) how location and language reveal attributes of identity, and 3) response to content via moderation systems.

2.1   What Are People Doing?

Our first question concerns how people use anonymous location-filtered social applications and how these uses reflect social identities. As noted above, we know that anonymity can cause disinhibition that has both positive (Rains, 2014; Turkle, 1995) and negative (Binns, 2013; E. Whittaker & Kowalski, 2015) effects. From a SIDE perspective, the key issue is the extent to which one feels a member of a group or community, versus feeling like a disconnected individual (Reicher et al., 1995). On location-filtered apps, however, there are two potentially pro-social identities at play: app user and local community member.

Based on prior work, we would expect these identities to manifest themselves in different ways in content. First, as in conventional online interaction, the disconnected individual identity will be evident in posts that exploit anonymity to engage in negative or purely selfish behaviors. This can be particularly evident in that being anonymous with a geographically local audience means that specific individuals can be identified, as discussed by Blackwell et al. (2015). On apps like Yik Yak, this can take the form of posters inadvertently revealing identifying information about themselves, or identifying specific individuals, as in cyber-bullying attacks that name or harass specific people (Binns, 2013; E. Whittaker & Kowalski, 2015) or when photos are used to label individuals as "sluts" when they are seen on location-aware hookup applications like Grindr in the gay community (Blackwell et al., 2015).

Second, the local identity should be evident in posts that reference or reach out to the community. We know from studies of location-based dating applications that location can affect the type of content users are willing to share (Birnholtz, Fitzpatrick, Handel, & Brubaker, 2014; Blackwell et al., 2015). Prior studies of online interaction with nearby people have also shown evidence that people seek information about local topics (Birnholtz, Merola, & Paul, 2015), coordinate social encounters, (Sutko & de Souza e Silva, 2011) or reach out for and provide help in crises (Vieweg, Hughes, Starbird, & Palen, 2010).

Third, the identity as an application user should be evident in content that references other users or posts, and/or is consistent with the norms of the app (Litt, 2012).

To explore how these social identities are manifest in the use of location-aware social apps, we first wondered about the types of content that people post. We asked:

RQ1: What types of content do people post on location-filtered anonymous social apps? How does this content reflect users' different social identities?

2.2   How Does Location Affect Language?

To further examine how users' social identities play out on apps like Yik Yak, we also wondered if there was evidence of these identities in the language used in posts. Litt (2012), in her work on how users assess and imagine audience for their social media content, discusses how the social context of interactions can affect the nature of content. One might, for example, post a detailed knitting question on a knitting forum using terms and information that would be inappropriate for a more general platform like Facebook. With location-filtered applications, social context becomes more nuanced because the online and offline social contexts are blended. That is, posts are essentially visible to anybody nearby, but the intended audience may be a subset of that audience, such as students at a local university. We would expect to see evidence of this in two ways.

First, the assumption of a shared local community identity with one's audience should result in finer-grained location terms likely to be understood only by community members (Fussell & Krauss, 1992), such as specific buildings or landmarks (Birnholtz et al., 2014). These may be useful in coordination of events (e.g., "meet by the big rock"), identification of individuals (e.g., "were you at the central gym?") or seeking local information (e.g., "what happened on University Ave. at 8pm?") (Birnholtz et al., 2015). People have also been observed to draw on locally relevant location language in self-presentation (Clark, 1996; Pennycook, 2010). This would serve to make the conversation intelligible only to those who share relevant common ground (Clark, 1996). That is, location dictates who can view the messages, but location-specific language restricts who can understand them (Ames & Naaman, 2007).

Second, coarse-grained location descriptors (e.g., university or neighborhood names) could be used in multiple ways. On the one hand, these could invoke a collective identity (e.g., "University X will

crush University Y in football!") among people who share the local identity (Postmes et al., 2001). On the other hand, these could also be used by disconnected individuals who do not share the local identity to draw attention to their outsider status, such as a University Y student posting "University X sucks!" at University X. To explore how language use reflects social identities in different locations, we asked:*RQ2: Is language use distinct in different geographic locations? Is there evidence of language being used to restrict audience or invoke shared identity?*

2.3   Norms and Feedback: Engagement with Content

A key aspect of any community, whether online or offline, is a shared understanding of norms for socially appropriate behavior. Bernstein et al. (2011), for example, showed evidence of normatively acceptable offensive and provocative content on 4chan. Other communities have employed many mechanisms to encourage and enforce respectful and pro-social behaviors. Studies of Wikipedia have shown how social processes there have been formalized into a series of roles, and how participation in these roles changes over time (Bryant, Forte, & Bruckman, 2005; Butler, Joyce, & Pike, 2008). Some sites also feature mechanisms by which users participate in community moderation. Popular examples include Slashdot and Reddit, which allow users to upvote and downvote content posted by others in order to separate high and low quality contributions (Lampe & Resnick, 2004). Even Facebook's newsfeed algorithm, which governs the visibility of content there, can be seen as a moderation mechanism, albeit one that many people do not clearly understand (Eslami et al., 2015). Importantly such mechanisms provide users with ongoing feedback as they learn to participate in the community (Lampe & Johnston, 2005).

Location-filtered social apps often have moderation or voting features as well, and we wondered what type of content is likely to be well received or rejected. As with posting content, this could also reflect both offline community and app-specific identities and norms. We asked:

RQ3: What identities do people seem to draw on, and what impacts response to and engagement with content on location-filtered social apps?

3   Research Context and Methods

Building on prior studies that have examined single platforms such as Facebook (Joinson, 2008; Lampe, Ellison, & Steinfield, 2006) or Twitter (Chen, 2011), we used a multi-method case study approach to analyze university student behavior on Yik Yak, a popular anonymous location-filtered social app. When opened, the Yik Yak app displays short messages (< 201 characters) called "yaks" posted by others nearby. Each yak can be upvoted or downvoted; votes determine the yak's score (+1 per upvote/-1 per downvote). Yaks that score below -4 are automatically deleted; there is no maximum score. Users can report offensive yaks, which are deleted if they are deemed to violate the app's content rules ("Yik Yak Support - Questions and Concerns," n.d.). Users can also post thread-style comments on yaks, with comments also subject to the character limit and voting.

Posting and commenting are anonymous in that no persistent identifying traits are shared by the system. It is possible for yaks to contain a screen name or "handle," but these need not be unique and do not persist from yak to yak. Handles are not frequently used, appearing in only 8% of yaks we collected. Privately, however, Yik Yak users do have a persistent identity linked to their mobile device. Users have a "Yakarma" score, visible only to them, that accumulates according to an algorithm that seems to be based on the number of yaks and comments posted, and points received. We did not have access to Yakarma scores, so do not discuss these further.

3.1   Quantitative Methods

Our primary data consisted of 1,985,318 unique yaks collected from 35 US university campuses between September 2014 and February 2015. We selected these campuses because they vary in size and region, and serve primarily residential (i.e., not commuter) student populations where Yik Yak is popular. The mean number of yaks per day per campus ranged from 51.4 to 921 (M=665.5; SD=507.0). We focused on students because Yik Yak is known to be popular on campuses and much attention has focused on student use. In collecting data, we were conscious of ethical and privacy concerns in gathering quasi-public data from social platforms (Zimmer, 2010). In this case, we considered the data quasi-public in that yaks would be visible to anybody who downloaded and opened the Yik Yak application in the locations studied (or who used the 'peek' feature to observe yaks at another location) during our data-gathering period. We stored yaks only on a password-protected server and plan to store them only for the duration of this project. We report no plausibly identifying information.

We gathered data via a Python script that made requests to open Yik Yak HTTP endpoints. The script ran once per hour each day during the collection period. On each run, the script collected the most

recent 100 yaks posted near each location. For each yak, we stored the message text, location, post time, handle (if any) and score, as well as comments and their scores. When the script encountered a yak we had previously stored, we created a new database entry with the updated score. In score analyses, we use the most recent values. Running once per hour allowed us to capture most activity, as we observed over 90% of yaks at least twice, indicating that most yaks are visible for at least an hour.

3.1.1   Linguistic Measures & Analysis

To analyze language use across locations we used tf-idf (term frequency-inverse document frequency) scores. tf-idf is a common information retrieval algorithm that increments the score for words appearing frequently in each document, but decrements scores for those words appearing frequently across multiple documents (Salton & Buckley, 1988). We constructed a vocabulary of 20 words or phrases specific to each campus by treating the set of yaks from each campus as a discrete document, and measuring the tf-idf score of each unigram (single words) and bigram (pairs of words) within each document. For any given campus, unigrams and bigrams with high scores are unique to that campus. Some terms occurred frequently on a specific campus, but only briefly (e.g., the name of an athletics rival just before or after a big game). We re-calculated tf-idf scores grouping yaks both by campus and month of posting, extracting the top 20 terms for each campus/month pair. We then produced a list of words persistently unique to each campus and a second list of words common only in one month.

To further analyze linguistic behavior, we used LIWC, or Linguistic Inquiry and Word Count, which is a text analysis package that counts certain word types in documents (Pennebaker, Booth, & Francis, 2007). In typical use, LIWC results are continuous values representing word counts relative to document length. As yaks are short, we used a binary indicator of whether words in a category appeared in a yak.

3.2   Content Analysis

To assess content types, we iteratively developed a coding scheme to categorize the perceived poster intent for each yak within a randomly selected subset of our yak data set. Our eight-category (Table 1) coding scheme was developed by repeatedly reading and discussing yaks, and coding subsets of data until we could reliably and mutually exclusively code virtually all yaks. Two coders independently coded small groups of yaks, with discussion in between, until their agreement rate was consistently above 80%. After this training phase, an additional 1850 yaks were coded. Of these, 350 were coded by both coders (agreement rate = 79.5%). The remaining 1500 were coded by one of the two coders.

3.3   Interviews

To preliminarily validate our results, we interviewed 12 undergraduates (8 female; age 18-21) at our university who self-identified as regular Yik Yak users. Participants were recruited using fliers on campus, email lists, and posts on Yik Yak. Interviews followed a semi-structured protocol consisting of: 1) general questions about their use of Yik Yak, and 2) showing participants 15 local yaks and, for each one, asking them to discuss reactions, whether they would vote them up or down, if they would comment, and why they had these responses. Interviews were audio-recorded and the interviewer took detailed notes, including key quotations. These notes, augmented after the interviews by consulting recordings, were used to identify themes and points of intersection (or departure) with our content analyses. Given that these interviews were not the primary focus of this study, and we did not recruit more participants for a full interview study, we present these findings primarily to help validate our other results, rather than attempt to present in depth qualitative results.

4   Results

4.1   What do people use Yik Yak for?

Our first question concerned content on Yik Yak. Table 1 shows the breakdown across the eight categories in our coding scheme, which was non-equal (χ^2 = 1126.5, p < .001).

In considering the different social identities we mentioned earlier, we first looked for evidence of posts that clearly drew only on an individual identity and did not seem to reflect community identity or community norms at all. However, this type of content did not merit a category in our coding scheme because we saw virtually no evidence of this.

Much more common, across most categories, were yaks that seemed to draw on the local community identity by referencing experiences that the audience was likely to understand or find amusing (e.g., in the 'observations' category: "Bus 36 smells like weed to the max!!!!," or "A guy literally just pulled out a water bottle of whiskey and took a swig in the middle of class. Dead week gets to all of us I suppose"). Another way people showed evidence of drawing on the local shared identity was to presume

the community was there in asking questions or seeking advice. In the "information/advice" category, for example, one yak asked "Would joining a fraternity be worth it?" This question is specific to the campus where it was posted, as campus social scenes may vary substantially. Moreover, asking a location-filtered audience including strangers may yield a more diverse pool of responses (i.e., from those he does not know) than the poster would get from a social network site. In this example, the poster may receive responses from people who have already joined fraternities, who are likely older.

Next, we wondered whether people were also influenced by their social identity as Yik Yak users. While we explore this further in our interview data later, it became clear in reviewing yaks that Yik Yak is frequently used to share or seek information about topics that may be sensitive or personal. This is evidenced in that the "personal admissions" and "information/advice" categories were among the most popular. In examining content, moreover, it became clear that it was normatively acceptable on Yik Yak to discuss sensitive topics that one might not otherwise discuss with strangers on either named social networks or face-to-face on a college campus. Some yaks in the "information/advice" category posed questions like, "How do you tell a guy you want him without saying 'let's fuck'?".

The same was true for personal admissions, which made up the largest percentage of coded yaks. Many of these show users dealing with sensitive issues such as health and body image issues (e.g., "Trying to eat healthy when you haven't been for a long time is so hard. It's literally an addiction. I just want pizza and wings and burgers and bacon and subway and pizza") and relationship issues (e.g., "Been hanging out with a guy who I know I won't marry and doesn't have the same standards as me... But I don't feel like I can do any better so I continue to cuddle and kiss and hang out with him.").

Intent	Count (Percentage)	Example	Median Score
Personal Admission	508 (34.8%)	"I got so drunk that I pissed my sheets last night. Other than that the party was great"	7.0
Observation	310 (21.2%)	"Just watched two grounds keepers walk under the Century Tree together, daaaaaaaamn"	9.5
Information/Advice	253 (17.3%)	"Im depressed and need to go for a walk to clear my head. Where can I go (on campus)?"	4.0
Opinion	188 (12.9%)	"These seniors are an inspiration no matter what team you're rooting for"	6.0
Venting/Complaining	125 (8.6%)	"its to that point in the semester that anything my roommate does is so freaking annoying."	8.0
Invitation	43 (2.9%)	"Who wants to go ice skating this Friday before the parties begin?"	1.0
Favor	20 (1.4%)	"Went too hard last night and lost my shoe somewhere between lax house and rich:/ it's sparkly and gold if anyone's seen it...."	2.0
Joke	12 (0.8%)	"I'm not slurring, I'm just talking in cursive"	16.5

Table 1. Counts and Examples of Yak Categories

4.2   Language and the Local Audience

Our second question concerned the location-specific use of language and how this reflected social identities. To explore language use across and within campuses, we examined the phrases with high tf-idf scores (see above) each month, on each campus. Of all yaks, 511,546 (25.8%) contained at least one word or phrase from our location specific lists. In turn, 173,963 yaks (9% of all yaks) contained a phrase that occurred in only one monthly list. We also compared phrases that were repeated for at least 3 months to those that were not. We realized in making these comparisons that we could characterize local references on two dimensions: implicit vs. explicit, and geographically local vs. temporally local.

Explicit references use language that directly identifies local entities. Examples include a school name, buildings on campus, and references to specific classes. Implicit references depend on local, shared understanding. The word "midterms" by itself, for example, is not specific to any campus. Rather,

as different campuses have midterms at different times, the word carries local meaning during these times because one can assume that others nearby will understand.

Geographically local references include landmarks, buildings, city names, and other references that refer to specific locations. Temporally local references refer to socially-defined entities that are salient during specific times, such as a class name like "ECON 310," which is relevant when that class is offered.

We first looked for explicit references specific to each campus that would be primarily understood by the local audience, but did not reference the institution by name. For example, Ballantine Hall, which was high on the tf-idf list every month, is a building at Indiana University that apparently has difficult-to-climb staircases referenced in many yaks (e.g., "The worst kind of torture is being sick and having to walk up the Ballantine stairs. Breathing quietly is not an option." and "12:30 at night, still tired from going up 6 Ballantine flights of stairs.") While outsiders can easily find this information (as we did), the intent here is to draw on Indiana student identity via the shared experience of difficult stair climbing.

In comparing the monthly lists, we also saw temporally local references, as in the case of the phrase "ECON 310" that was only common in one month on one campus. Posters of these yaks seemed to be addressing other students in this class, whom they may or may not know, to share anxiety or to draw sympathy from those who had taken the class before. In this way, the reference identifies and renders salient another group identity, that of a student in a particular course.

Implicit temporal references were also evident in our data, such as the phrase "financial aid," which appeared frequently only during one month on one campus. While this phrase is not specific to a university it was used primarily when financial aid deposits were due and students were anticipating payments (e.g., "YO when are we seeing that financial aid though" and "Anybody know when financial aid drops") or celebrating when they arrived (e.g., "Getting a lot more than I thought I was going to back from financial aid. Hell yeah."). This also reflects a topic that may be uncomfortable for students to discuss face-to-face, as socio-economic status can be a friction point on U.S. college campuses.

We also wondered about more coarse-grained geographical references such as institution names. In 29 of the 35 universities in our data set, the institution name was in the high tf-idf list of words for every month. As we noted earlier, we expected these names to be used to either invoke shared identity, or for outsiders to draw attention to their outsider status. Often these positively invoked shared identity (e.g., "I like to walk around campus early in the morning when there is no one around. It's really beautiful. I live with my family. I love you FSU.") though sometimes these were more negative descriptions in which the poster sought show distance from the campus identity (e.g., "Cornell is a haven for type A overachievers. You need to get excellent grades, have awesome internships, be in a 'top' frat/sorority, lead a bunch of clubs, and still feel like you're not good enough").

4.3   How Do People Engage with this Content?

Our final question concerned content engagement via comments and voting. The majority of yaks (57%) in our data set received no comments. On the other hand, only 7% of yaks we examined have a score of 0. We therefore focus our analysis on voting rather than commenting. Median scores across our yak categories ranged from 1-16.5. Results from a logistic regression modeling the relationship between the likelihood of a yak receiving a high score (which we defined as greater than one standard deviation above the campus mean) and the use of words from different lexical categories are shown in Table 2.

Users appear to draw on different identities when deciding what to upvote, and these votes serve different purposes. For example, categories related to social support, such as personal admission and venting, have high median scores (7 and 8, respectively), suggesting that upvoting is one way users offer support. Some venting and personal admission yaks that received higher than average scores include "Nothing like a shitty day to make me wish I had a girlfriend. I just want to cuddle up with someone and forget everything else," and "I would rather lay out in this blizzard all night than be stuck in the room with my roommate for 2 fucking days." High scores here seem to indicate identifying with these feelings.

A theme that came up in our interviews was rewarding yaks about topics students could relate to. As one student put it, "Some of the ones that I like the most, that turned out to be the most popular, were just like general college sentiments. Classes suck... I'm not getting enough sleep. And it's like, I can relate to that." This also is similar to the idea of temporal locality discussed above. We would not expect to see these yaks over holiday or summer breaks. It is not enough to be relatable; time matters as well. We see evidence for this in our regression results. The use of location specific language increases the odds of receiving a high score by 39.1%, the largest effect in our model. While the quote above describes "general college sentiments," these results suggest that users also value content that is locally relevant.

Similarly, users seem to value content that invokes a collective identity. While the use of first person singular personal pronouns is associated with a slight increase in the likelihood of receiving a high

score, we see larger effects in the use of first person plural personal pronouns and words from the LIWC category Social Processes. For example, the yak "So you think [NAME] is trying her best to make us hate chemistry so that she can have a better chance at a Nobel prize?" received a high score. This provides an interesting example of how the poster did not post a message such as "I dislike this class," but rather directly addressed a local audience of others in that class, invoking a shared identity.

Jokes offer a counterexample to treating upvotes as a support mechanism. Jokes, which make up only 1% of our coded yaks, have the highest median score (16.5). Examples include "What do you call Batman when he skips church? Christian Bale." and "Just found my prosthetic leg finallyyyy! It's always the last place you hop..." A related theme from our interviews was the importance of cleverness. Almost all students interviewed said they would upvote the yak, "The hardest part about riding my razor scooter through campus is dodging all the panties being thrown at me." Students used words like unique, funny, and clever to describe it. Location also plays an important role here, as students also mentioned finding this funny because they had seen students riding scooters on the campus where it was posted.

All of this suggests that different types of yaks receiving upvotes can be seen as students drawing on different social identities. In upvoting venting and personal admission yaks they relate to, students draw on their offline, local identity and offer support to others like them. Rewarding funny yaks has less to do with this offline identity and more to do with reinforcing content they enjoy seeing on Yik Yak. For example, one student said she would downvote the yak, "guys - what do you think of a girl gives you a handjob?" not because of the content, but rather because "Yik Yak's not like, a question and answer platform." The data support this in that the categories with the lowest scores included favor (2), information/advice (4), and invitation (1). We also found that these categories were no more likely than the others to receive comments, suggesting that they simply did not receive much engagement.

	β (SE)	Odds Ratio
Intercept	-2.744 (.008)***	0.064
Location Specific Words	0.330 (.009)***	1.391
Negative Affect	0.022 (.003)***	1.022
Positive Affect	0.074 (.003)***	1.077
First Person Singular Personal Pronouns	0.061 (.002)***	1.062
First Person Plural Personal Pronouns	0.114 (.006)***	1.121
Social Processes	0.124 (.001)***	1.132

Table 2. Logistic Regression Results of Yak Receiving a High Score

5   Discussion

5.1   Implications for Theory

One implication of our results concerns audience. Engaging with an audience that is simultaneously geographically local and online challenges ideas about audience. Litt (2012) discusses the notion of different shared online contexts having distinct social norms and providing some shared understanding of the type of content that is appropriate for an audience. Here we saw evidence that people's posts were likely impacted by both the "Yik Yak" app context and the geographically local context.

From a local standpoint, yaks about exams, events and a building with difficult staircases suggest that people were trying to engage with local others. At the same time, however, the distribution of posts across categories and our interviews suggest that people did not use Yik Yak to coordinate with others or meet strangers. That is, that they saw the app as a sort of overlay on geographic space that had its own norms and expectations. This additional layer factored into both content and reactions to it in our results. All social apps add additional contexts for interaction; some of these (e.g., Facebook) blend our online and offline identities. Apps like Yik Yak add another potential dimension to this complexity by factoring in identities associated with organizations and location. For example, Naaman, Boase, & Lai (2010) included "information sharing" in their classification of tweets. What they saw in this category was information relevant to other Twitter users all over, whereas our users who shared information focused on the local audience, such as reporting a local bike thief or asking about places on campus.

This intersection and overlapping of contexts is further complicated by anonymity. Where people interacting on sites with persistent identities (or in face-to-face local contexts) are often concerned about

impression management and future credibility, Yik Yak users do not have these concerns. As such, people could admit to common failings, such as poor academic performance, and reap the possible benefits of support (via upvotes) without being accountable or even visible to the people providing the support. In some ways this resembles prior online support forums or discussions (Rains, 2014; White & Dorman, 2001) but these were often characterized by persistent pseudonymous identity and revolved around longer-term support. On Yik Yak, identities and content are ephemeral. One could theoretically post about a failing grade, receive support via upvotes and then essentially forget the experience without fear that friends will ask about it or that future employers will know about the failure.

5.1.1   Language, Location and Identity

We also found that people engaged the local audience using unique local language. This may not seem surprising in that it happens in other interaction forums as well, but it is interesting here in that, unlike in-person co-located discourse, messages posted to Yik Yak are visible to all those nearby. This public visibility can affect nearby users' social identities. For example, an outsider reading yaks that make local references she does not understand may be primed to feel very differently than would somebody who understands the references. From a SIDE standpoint (Postmes et al., 2001), this could provoke very different posting behavior such as negativity instead of positive posts, and this merits further exploration as we seek to better understand problematic behavior on anonymous apps.

In addition, we found local references occurring at varying levels of specificity, each of which can serve different purposes. Some yaks named the institution or local area, for example, but these tended to invoke broadly shared identity or distance the poster from the institution (e.g., the disparaging post about Cornell). More common were yaks that used implicit language to draw on shared knowledge of the location, or used explicit language that was sufficiently specific (e.g., the stairs at Indiana) to exclude outsiders. Distinguishing these types of references helps provide insight into the different ways users use location specific language in platforms like Yik Yak, and what their effect on social identity may be.

Another implication is that location-based filtering of content differs in important ways from location-awareness, which has been the focus of much prior discussion of location (e.g., Guha & Birnholtz, 2013; Licoppe, 2015; Lindqvist et al., 2011). Location-awareness apps often focus on social coordination or using location information to manage impressions. Location-based filtering, on the other hand, uses location to provide shared social context for interactions with a local audience mostly about things other than location or coordination. We saw people use and draw on location information in nuanced ways that allowed them to demonstrate insider status or invoke shared social identity, but they rarely used location to discuss actual locations or coordination. This suggests that, in exploring location-based filtering, researchers must think broadly about the many ways that location provides shared context and the multiple ways this can be drawn on. Our notions of implicit vs. explicit and temporal vs. geographic provide a starting point for a framework in this area.

5.2   Implications for Design

Our results also have implications for designing mechanisms for location-based content filtering. One is that designers consider designs that prime local community identities. This could be achieved, for example, via school colors or logos, or other locally relevant visuals, which could prompt people to think about their association with the local community as they post and vote. SIDE suggests this should prime pro-social behavior for many users, though not for visitors or outsiders.

Another possibility is to distinguish yaks from local insiders vs. visitors or outsiders. For example, on a college campus, users who authenticate with a valid .edu email address could have content they post highlighted in some way. This could affect the type of content they post, the type of content the non-authenticated users post, and how these different groups interact. The effects of such a design merit further exploration in understanding how social identities are primed by design and drawn on.

Another suggestion is to identify and allow users to browse content by "trending" temporally relevant topics. We found users make both geographically and temporally local references. Identifying temporally relevant yaks provides another way to prime community identity. Using our example above, allowing users to view all "financial aid" yaks from the day would allow those students affected by financial aid news to find information, commiserate, or celebrate with other local students.

5.3   Limitations and Future Work

As with any study, this work has limitations that urge interpretation with caution. One key limitation is that we are examining posts and responses to them, but have only preliminary information (from 12 interviews) about people's intent. Additional research is needed to more broadly validate our insights and interpretations, but we believe our data support the claims we present here. Second, we studied publicly

visible elements of Yik Yak, meaning that we missed posts immediately reported or deleted by users or system administrators. It is possible that these unseen elements affect voting and posting behavior in ways that we have not accounted for. Given that we saw over 90% of Yaks at least twice, however, we believe we have observed a substantial fraction of user behavior. Third, we focused on Yik Yak use on university campuses and interviewed students at only one campus. General demographic statistics are not available for Yik Yak users, so we cannot claim a representative sample of users or that our findings will generalize more broadly, but informal evidence suggests that the app is popular in this group.

This exploratory study opens substantial opportunity for further study of location-filtered social apps. While our observations suggest that language understandable only to an ingroup may prime different feelings related to social identity, we cannot make causal claims. Experimental or log data analysis to further explore this question would be required. Additionally we have made limited use of automated analysis techniques for processing large numbers of yaks because we were initially unaware of how the app was being used or what might be fruitful topics to try to identify. Future work could report more completely on the relationships between such matters as topic, voting and language use by using coding schemes and human-coded examples such as ours as a starting point. Finally, while we focused broadly on social identity, we did not focus on aspects of identity such as gender, race, socio-economic class, and able-bodiedness, which may be addressed by future work.

## 6	Conclusion

We have presented an exploratory study of Yik Yak, an anonymous, location-aware social media application. Our findings indicate users engage local audience in the types of content they post, and they do so by using different types of location specific language. We also found evidence that suggests that users draw on different social identities when engaging with posted content. We believe this work provides an empirical understanding of location-based anonymous interaction apps, which is important given recent media coverage of these apps over cyber-bullying cases, and provides opportunities for both researchers and designers of these types of applications.

## 7	References

Allen, J. (2015, April 27). U-M, Yik Yak prod students to "pledge to be positive online." Retrieved May 19, 2015, from http://www.mlive.com/news/ann-arbor/index.ssf/2015/04/u-m_partners_with_yik_yak_enco.html

Ames, M., & Naaman, M. (2007). Why We Tag: Motivations for Annotation in Mobile and Online Media. In *Proc. CHI* (pp. 971–980). New York, NY, USA: ACM.

Barak, A., Boniel-Nissim, M., & Suler, J. (2008). Fostering Empowerment in Online Support Groups. *Comput. Hum. Behav.*, *24*(5), 1867–1883.

Bernstein, M., Monroy-Hernández, A., Harry, D., André, P., Panovich, K., & Vargas, G. (2011). 4chan and /b/: An Analysis of Anonymity and Ephemerality in a Large Online Community. In *Proc. ICWSM*

Binns, A. (2013). Facebook's Ugly Sisters: Anonymity and Abuse on Formspring and Ask.fm. *Media Education Research Journal*. *4(1)* 27.

Birnholtz, J., Fitzpatrick, C., Handel, M., & Brubaker, J. R. (2014). Identity, Identification and Identifiability: The Language of Self-presentation on a Location-based Mobile Dating App. In *Proc. MobileHCI* (pp. 3–12). New York, NY, USA: ACM.

Birnholtz, J., Merola, N. A. R., & Paul, A. (2015). "Is It Weird to Still Be a Virgin": Anonymous, Locally Targeted Questions on Facebook Confession Boards. In *Proc. CHI* (pp. 2613–2622). New York, NY, USA: ACM.

Blackwell, C., Birnholtz, J., & Abbott, C. (2015). Seeing and being seen: Co-situation and impression formation using Grindr, a location-aware gay dating app. *New Media & Society*, *17*(7), 1117–1136.

Bryant, S. L., Forte, A., & Bruckman, A. (2005). Becoming Wikipedian: transformation of participation in a collaborative online encyclopedia. In *Proc. SIGGROUP* (pp. 1–10). New York, NY, USA: ACM.

Butler, B., Joyce, E., & Pike, J. (2008). Don't look now, but we've created a bureaucracy: the nature and roles of policies and rules in wikipedia. In *Proc. CHI* (pp. 1101–1110). New York, NY, USA: ACM.

Chen, G. M. (2011). Tweet this: A uses and gratifications perspective on how active Twitter use gratifies a need to connect with others. *Computers in Human Behavior*, *27*(2), 755–762.

Clark, H. H. (1996). *Using language* (Vol. xi). New York, NY, US: Cambridge University Press.

Dabbish, L., Farzan, R., Kraut, R., & Postmes, T. (2012). Fresh Faces in the Crowd: Turnover, Identity, and Commitment in Online Groups. In *Proc. CSCW* (pp. 245–248). New York, NY, USA: ACM.

Eslami, M., Rickman, A., Vaccaro, K., Aleyasen, A., Vuong, A., Karahalios, K., … Sandvig, C. (2015). "I Always Assumed That I Wasn't Really That Close to [Her]": Reasoning About Invisible Algorithms in News Feeds. In *Proc. CHI* (pp. 153–162). New York, NY, USA: ACM.

Fussell, S. R., & Krauss, R. M. (1992). Coordination of knowledge in communication: Effects of speakers' assumptions about what others know. *J. of Personality and Social Psychology*, *62*(3), 378–391.

Guha, S., & Birnholtz, J. (2013). Can You See Me Now?: Location, Visibility and the Management of Impressions on Foursquare. In *Proc. MobileHCI* (pp. 183–192). New York, NY, USA: ACM.

Joinson, A. N. (2008). Looking at, Looking Up or Keeping Up with People?: Motives and Use of Facebook. In *Proc. CHI* (pp. 1027–1036). New York, NY, USA: ACM.

Lampe, C., Ellison, N., & Steinfield, C. (2006). A Face(Book) in the Crowd: Social Searching vs. Social Browsing. In *Proc. CSCW* (pp. 167–170). New York, NY, USA: ACM.

Lampe, C., & Johnston, E. (2005). Follow the (Slash) Dot: Effects of Feedback on New Members in an Online Community. In *Proc. SIGGROUP* (pp. 11–20). New York, NY, USA: ACM.

Lampe, C., & Resnick, P. (2004). Slash(Dot) and Burn: Distributed Moderation in a Large Online Conversation Space. In *Proc. CHI* (pp. 543–550). New York, NY, USA: ACM.

Licoppe, C. (2015). Seams and folds, detours, and encounters with "pseudonymous strangers": Mobilities and urban encounters in public places in the age of locative media. In *i3 Working Papers Series, 15-SES-05*.

Lindqvist, J., Cranshaw, J., Wiese, J., Hong, J., & Zimmerman, J. (2011). I'm the Mayor of My House: Examining Why People Use Foursquare - a Social-driven Location Sharing Application. In *Proc. CHI* (pp. 2409–2418). New York, NY, USA: ACM.

Litt, E. (2012). Knock, Knock. Who's There? The Imagined Audience. *Journal of Broadcasting & Electronic Media*, *56*(3), 330–345.

Madden, M., Am, Lenhart, a, S, Cortesi, Gasser, U., … Beaton, M. (2013). Teens, Social Media, and Privacy. Retrieved from http://www.pewinternet.org/2013/05/21/teens-social-media-and-privacy/

Naaman, M., Boase, J., & Lai, C.-H. (2010). Is It Really About Me?: Message Content in Social Awareness Streams. In *Proc. CSCW* (pp. 189–192). New York, NY, USA: ACM.

Pennebaker, J. W., Booth, R. J., & Francis, M. E. (2007). Linguistic Inquiry and Word Count (LIWC2007). Retrieved from https://researchspace.auckland.ac.nz/handle/2292/18109

Pennycook, A. (2010). *Language as a Local Practice*. Routledge.

Postmes, T., Spears, R., Sakhel, K., & Groot, D. de. (2001). Social Influence in Computer-Mediated Communication: The Effects of Anonymity on Group Behavior. *Personality and Social Psychology Bulletin*, *27*(10), 1243–1254.

Rains, S. A. (2014). The implications of stigma and anonymity for self-disclosure in health blogs. *Health Communication*, *29*(1), 23–31.

Reicher, S. D., Spears, R., & Postmes, T. (1995). A Social Identity Model of Deindividuation Phenomena. *European Review of Social Psychology*, *6*(1), 161–198.

Salton, G., & Buckley, C. (1988). Term-weighting approaches in automatic text retrieval. *Information Processing & Management*, *24*(5), 513–523.

Smith, A. (2014). 6 new facts about Facebook. Retrieved from http://www.pewresearch.org/fact-tank/2014/02/03/6-new-facts-about-facebook/

Sutko, D. M., & Silva, A. de S. e. (2011). Location-aware mobile media and urban sociability. *New Media & Society*, *13*(5), 807–823.

Toch, E., & Levi, I. (2013). Locality and Privacy in People-nearby Applications. In *Proc. UbiComp* (pp. 539–548). New York, NY, USA: ACM.

Turkle, S. (1995). *Life on the Screen*. Simon and Schuster.

Vieweg, S., Hughes, A. L., Starbird, K., & Palen, L. (2010). Microblogging during two natural hazards events: what twitter may contribute to situational awareness. In *Proc. CHI* (pp. 1079–1088). ACM.

White, M., & Dorman, S. M. (2001). Receiving social support online: implications for health education. *Health Education Research*, *16*(6), 693–707.

Whittaker, E., & Kowalski, R. M. (2015). Cyberbullying Via Social Media. *J. of School Violence*, *14*(1), 11–29.

Whittaker, S., Terveen, L., Hill, W., & Cherny, L. (2003). The Dynamics of Mass Interaction. In *From Usenet to CoWebs* (pp. 79–91). Springer London.

Yik Yak Support - Questions and Concerns. (n.d.). Retrieved from http://support.yikyakapp.com/

Zimmer, M. (2010). "But the Data is Already Public": On the Ethics of Research in Facebook. *Ethics and Inf. Technol.*, *12*(4), 313–325.