

140 Characters Of F*ck, Sh!t, And @ss: How We Swear On Twitter

fastcoexist.com/3026596/140-characters-of-fck-sht-and-ss-how-we-swear-on-twitter

2/19/2014

At its best, Twitter is a place to find breaking news, thoughtful dialogue, and unexpected voices. At its worst, it's a forum for knee-jerk reactions, trolls, and harassment. At both of these extremes, as in life, there's bound to be some bad language along the way.

Just how much do people curse on Twitter? Who swears and what do they say?

Four mild-mannered researchers endeavored to find out, and [the paper they produced](#)—presented [this week](#) at the ACM Conference on Computer-Supported Cooperative Work & Social Computing—is a compendium of offensive language that rivals [the depraved, cuss-filled brilliance](#) of the recent film *The Wolf of Wall Street*.




Figure 2. Cumulative distribution of curse Words: The top 7 curse words cover more than 90% of all the curse word occurrences.

After examining a random one-month sample of 51 million English-language tweets from 14 million distinct user accounts, they came up with this conclusion: We curse a lot on Twitter, where our language is usually public, even more than we do in real life. Even more compelling, they discovered the underlying context of when and why cursing happens and who is cursing to whom.

"It's a sizable fraction of the words we use. On average, one tweet out of 13 tweets will contain at least one cursing word," says [Wenbo Wang](#), a PhD researcher at Wright State University who led the study. "Because of social media,

people don't see each other. They can say things they wouldn't say in the physical world." Other studies have found that 0.5 to 0.7% of words we say in the physical world are curses—on Twitter, the researchers found the rate to be 1.15%. Or as the paper reads, and as Wang was too polite to repeat during our phone interview:

The most popular curse word is fuck, which covers 34.73% of all the curse word occurrences, followed by shit (15.04%), ass (14.48%), bitch (10.34%), nigga (9.68%), hell (4.46%), whore (1.82%), dick (1.67%), piss (1.53%), and pussy (1.16%).

The findings are interesting for anyone who uses Twitter, but for the team, all affiliated with [Ohio Center of Excellence in Knowledge-enabled Computing](#), the paper will fold into work with broader societal implications related to mental health, verbal abuse, online harassment, and gender differences in online communications.




Figure 6. Cursing Volume in Different Days of Week

"Social content is extremely rich," says the center's director [Amit P. Sheth](#). "The cursing issue is an expression of sentiment and emotion...it's kind of a core issue of understanding the language." The center is working on developing automated tools that could flag issues of worrisome harassment on social media, especially in high school and college years, or could identify depressive disorders or disposition to violence. Creating filters for kids on social media is also another potential application.

Of course, timing and context is everything, and not all cursing is negative. For example, "I f*\$%-ing love you" could be music to the @ recipient's ears.

An attempt at "sentiment analysis," which has its limitations, revealed that while negative emotions won out in swearing tweets, love and playfulness (two friends saying "you whore") emerged as real signals in the data. Anger and sadness represented 22% and 17% of cursing tweets, whereas 7% seemed to express love. One in four of all tweets sampled that were classified as "angry" contained curse words.

Other insights in the paper involved timing, location, gender, and level of influence of the NSFW tweeters. People curse more and more as the day passes, reaching a peak at 12 a.m. to 1:30 a.m. before bedtime, and Mondays, Tuesdays, and Wednesdays contain the most curse words relative to tweet volume.

And as in real life, people who are in more relaxed environments like home or at a club—rather than in the office—are more likely to curse in a tweet, but the differences are smaller than in the face-to-face world (the researchers looked only at geo-located tweets for these). College and high school students aren't shy about cursing, however, even when they are at school. Men curse more than women, but both genders are more likely to curse when directly conversing on Twitter with someone of the same gender. As for social rank, celebrities in the top 1% of follower counts on Twitter get treated better than "middle ranked" Twitter users:

The cursing ratio among tweets received by the top 1% group, is the lowest across all recipient groups: these popular users receive a lot of friendly messages from their fans, e.g., "@Harry Styles follow me babe<3", "@NiallOfficial I can't sleep :("

The researchers do say there is room to improve their classification system, as the endeavor of even pinning down curse words isn't as straightforward as it seems. After removing spam, the team had to code each tweet as "cursing" or "non-cursing." That meant deciding what a curse is. For this, the four authors—none of whom are native English-speakers—compiled a lexicon of offensive words and asked two undergraduates (who else?) to assist in resolving ambiguities. They realized "gay" can both be used as a slur or in common descriptive speech. The lexicon also had to be modified to include all sorts of variations: " e.g., a55, @\$\$, \$h1t, b!tch, bi+ch, c0ck, f*ck, l3itch, p*ssy, and dik."

"I think our vocabularies have increased," says another one of the other authors, Lu Chen.

Never miss a story.

