Do You Smile with Your Nose? Stylistic Variation in Twitter Emoticons

Tyler Schnoebelen*

1 Introduction

On the surface, emoticons seem to convey emotional stances, so we expect smiles and frowns to be used differently from one another. But there are other systematic patterns of variation in emoticons that are not easily described by terms like "friendly" or "sad." Consider the twenty-eight most common emoticons in American English tweets in Table 1; even in smiling, people vary in their use of eyes, mouth shape, face direction, and whether or not they represent a nose in the face. And we see these in groups other than smiles—in variants of frowns and winks, for example. This paper focuses on nose and non-nose users, demonstrating that the variants correspond to different types of users, tweeting with different vocabularies.

Shorthand	Emo	Count	Percent	Shorthand	Emo	Count	Percent
smile	:)	1,496,585	39.6%	**frownnose**	:-(	27,561	0.7%
wink	;)	397,745	10.5%	**smileapos**	:')	23,994	0.6%
frown	:(	312,769	8.3%	**dworry**	D:	23,901	0.6%
bigsmile	:D	281,907	7.5%	**smilebrac**	:]	21,030	0.6%
smilenose	:-)	183,131	4.9%	**eqeyesbigsmile**	=D	20,785	0.6%
tongue	:P	169,417	4.5%	**slantnose**	:-/	19,176	0.5%
rsmile	(:	155,571	4.1%	**eqeyesbrac**	=]	17,504	0.5%
slant	:/	126,640	3.4%	**winktongue**	;P	17,460	0.5%
xeyesbigsmile	XD	114,862	3.0%	**tonguenose**	:-P	16,263	0.4%
eqeyessmile	=)	79,054	2.1%	**frownapos**	:'(	15,964	0.4%
winknose	;-)	70,618	1.9%	**bigsmilenose**	:-D	15,679	0.4%
omouth	:O	60,822	1.6%	**eqeyesslant**	=/	15,241	0.4%
winkbigsmile	;D	34,907	0.9%	**eqeyestongue**	=P	14,055	0.4%
doublesmile	:))	28,614	0.8%	**eqeyesfrown**	=(	13,919	0.4%

Table 1: From the American English Twitter corpus being studied in this chapter: counts / percentages of tweets with the top 28 emoticons.

Emoticons were first proposed in order to guide affective interpretations (in particular jokes).[2] From their modest beginnings in Carnegie Mellon, they have achieved broad penetration in computer mediated communication. For example, I find that 9.7% of American English tweets include at least one emoticon.

Beyond understanding a set of widespread linguistic resources, I advance the claim that emoticons are preserving part of what happens in actual speech. So they are not idiosyncratic, unusual devices; looking at their usage is instructive for how other affective linguistic resources are used. By encoding meanings that are usually part of the speech signal, we can see what sorts of meanings are indispensable. Intonation is a rich source of affective meaning (Bolinger 1989, Chang 1958, Fonagy and Magdics 1963, Frick 1985, Lieberman and Michaels 1962). To some degree it's possible to communicate prosody with punctuation (*?!?!*, ...) and what I call affective lengthening

*This research has been enhanced through conversations with Eric Acton, David Bamman, Penny Eckert, Roey Gafter, Kate Geenberg, Steve Guilliams, Dan Jurafsky, Boris Khmelnitskiy, Robin Melnick, Kyuwon Moon, Christopher Potts, John Rickford, Ron Shigeta, Rebecca Starr, Kyle Wild, and Lal Zimman.

(soooo). And of course the contours of pitch are not the only thing one has access to in face-to-face conversations: we also have the dynamic contours of face and body movements. Emoticons offer stylized representations of what gets lost when you switch to a text-only medium, like intonation and facial expressions. But emoticons are not simply representations of internal emotional states. They are more interactive in nature, positioning authors and audiences around propositions. The meaning of a given emoticon goes beyond its affective stance. For example, emoticons have variants that have greater or lesser affinities to standard language. Researchers who are interested in style, stance, affect, computer mediated communication, variation, context, and sentiment analysis will find ample grist for their mills in the present paper and related work (Schnoebelen 2011, Schnoebelen 2012).

2 Data

The corpus at the heart of this paper was collected by David Bamman for six months between January and June, 2011 (Bamman, Eisenstein, and Schnoebelen 2012). These tweets were collected to be representative of American English speech so only tweets whose authors were geolocated in America were used. While code-switching was allowed, full-time non-English tweeters were filtered out.[3]

There is an important convention in Twitter, which is to direct another user's attention to a tweet you append "@" to their account name. Since we were interested in real Twitter users, not single tweeters or corporations, we restricted the corpus to accounts that had sent messages @ someone and had been @'ed back by that person (this is a better measure than following/followers of a social network since it requires actual interaction). We restricted ourselves to people who had at least four but no more than 100 of these "mutual @ing" friends. Taken together, these measures protected us from spammers as well as mega- and non-users. The result was 38,927,633 tweets. The data was tagged with the CMU Twitter POS tagger (Gimpel et al. 2010).

Everything that received the emoticon tag was considered and the top 28 were taken. The corpus was then filtered to the 3,775,174 tweets that included at least one of the 28 emoticons. These tweets are made up of 18,559 word_pos pairs (13,411 unique words; 21,891,914 total tokens).

3 Variation by Authors

There are a total of 102,304 authors in our American English data. Let's consider the range of inter- and intra-tweeter variation. There are 43,962 people who use 10 or more emoticons. The most widely used emoticon is :). There are 21,124 people who have 10 or more uses of it and of these there are actually 17,073 who use :) only and no other emoticons at all.

One way of understanding emoticon variation is to look at what happens to people who are doing something other than the mainstream. There are 4,539 users who have no :)'s at all, but 10 or more other emoticons. In the four cases I show below, there are actually over 2.5 times as many users as we would've guessed just using raw probabilities from counts-per-user.[4]

- :-) 983 users with 10+ uses
- (: 606 users with 10+ uses
- =) 371 users with 10+ uses
- ;-) 357 users with 10+ uses

[2]Emoticons—the inventor prefers "smileys"—were apparently first proposed in order to mark jokes in bulletin board discussions:

19-Sep-82 11:44 Scott E Fahlman :-)
From: Scott E Fahlman <Fahlman at Cmu-20c>
I propose that the following character sequence for joke markers: :-)
Read it sideways. Actually, it is probably more economical to mark
things that are NOT jokes, given current trends. For this, use :-(

See also http://www.cs.cmu.edu/~sef/Orig-Smiley.htm and http://www.cs.cmu.edu/~sef/sefSmiley.htm.

[3]All authors had to use at least 50 of the 1,000 most common words in the US sample overall (predominately English terms).

These are the leading alternatives to the standard smile—adding a nose, putting the mouth first, changing the eyes to equal signs, or simply abandoning smiling for noseful winking. In later sections, we'll look at word patterns and also find more evidence that noses are different, so let's take a little more time with the per-user behavior. There are 23,773 tweeters who use at least 10 of :) and/or :-). As mentioned above, there are 21,123 people who have 10 or more :)'s. 17,731 of them have zero :-)'s. Only 804 people use :-) 20% of the time or more.

By contrast, there are 2,388 people who have 10 or more :-)'s. 938 of these people have zero noseless :)'s. 696 of the frequent nose users do use plain smile :)'s 20% of the time or more. Another way of figuring it is that we would've expected 1,147 people to use both 10+ :)'s and 10+ :-)'s. Instead we only get 513 such people.

Such constraints do not hold for all emoticons. When we look at people who have at least 10 tokens of :) and 10 tokens of other emoticons, we see that users of :) are also quite happy to be using a number of other emoticons, presumably to express different kinds of affective stances. Each of the following had at least twice as many actual users as we would've expected.

- ;) 3,739 users with 10+ uses (we would've guessed there were 1,505 such users)
- :(2,929 users with 10+ uses (we would've guessed 1,033)
- :D 2,785 users with 10+ uses (we would've guessed 1,167)
- :P 1,628 users with 10+ uses (we would've guessed 655)
- :/ 1,135 users with 10+ uses (we would've guessed 411)

In the next section, I shift to patterns between words and emoticons. But for the moment what I'd like to impress upon you is that there is variation in emoticon use, but there seem to be two types. The first is about affective stance, of which there is much to be said, but no room to say it in the present paper. The second type seems to be a kind of stylistic, almost aesthetic issue of what kinds of eyes and mouths you want to use, which way the emoticon should face, and whether or not it should have a nose.

4 Clustering by Word Cooccurrence

Human beings are good at picking out 2 or 3 dimensions from a pile of data, but here we have 28 emoticon dimensions, so we have to find a different way to find the ways the data cluster. In this paper, I will focus on hierarchical cluster analysis, although in other work I have presented results from factor analysis and topic modeling (Schnoebelen 2012). In this section, I demonstrate that by expanding to over 13,000 words, we can uncover the relationship between emoticons and the key axes along which they vary from one another.

4.1 Data Transformations

I begin with the data described in Section 2, that is 3,775,174 American English tweets with one of the 28 emoticons (18,559 word_pos pairs made up of 13,411 unique words; 21,891,914 total tokens). We'd like to know which words go with which emoticons but have that restricted to those with a statistically significant relationship.

Each word_pos-per-emoticon was treated as being part of a 2x2 contingency table, using a tool provided by Carlson, Heckerman, and Shani (2009). This tool reports both p-values and q-values, the latter of which has to do with false discovery rates.[5] In the present data, q-values were always more conservative. That is, the word_pos x emoticons that failed to reach < 0.05 significance using p-values were a perfect subset of those that fail to reach < 0.05 significance using q-

[4]Here's an example of the math: of our 43,962 users with more than 10 emoticons, there are 816 who use ten or more ;-). There are 4,539 users with more than 10 emoticons but no :) tokens. So if there was nothing conditioning emoticon patterns, we would expect (816/43.962)*4,539=84 people who never used :) but did use more than 10 ;-). Instead we get 4.2 times that—357 users.

[5]See Benjamini & Hochberg 1995 for more about False Discovery Rates; the expected proportion of false positives among all significant hypotheses. See Storey 2002, 2003 for further development & the notion of a q-value, which gives a Bayesian measure of significance in terms of (positive) False Discovery Rates.

values. For that reason, all word_pos x emoticons that had a q-value > 0.05 were seen as having no special relationship between the word and the emoticon. They are coded as “1”, while all significant values were coded as their “observed/expected” (OE) ratio. Since clustering techniques are sensitive to spread, all word_pos x emoticons with an OE > 3.0 were reduced to 3.0.

Since we are interested in clustering, we only consider word_pos’s that have over- or under-representation with at least two emoticons. That is, word_pos’s that have 27 or 28 1’s don’t tell us about how the words and emoticons cluster, so they are left to the side. The analysis consists of 8,913 word_pos pairs (6,909 unique words).

4.2 Hierarchical Clustering

There are a number of hierarchical clustering techniques; this research used agglomerative hierarchical clustering. This method starts with individual points and fuses like points together one-by-one. Once a fusion is made, it’s done and the new fused point is available for further fusing if there’s another point nearby. By joining “like points” one-by-one, the algorithm ultimately ends up with one big point, which can show the hierarchy that went into building it. For the clustering in the present study, similarity/difference was calculated using Euclidean distance and with Ward’s method. Ward’s method minimizes variance so that at each step, the points that are merged are the ones that “do the least damage” to the existing clusters.

Agglomerative hierarchical clustering shows a structure, but that doesn’t mean we can trust every division equally. One way to see which groups come out most strongly is to calculate p-values for each cluster using multiscale bootstrap resampling. This tests the hypothesis “this cluster doesn’t exist” to see whether it is accepted or rejected. Clusters that don’t achieve significance with bootstrapping may just be caused by sampling error. Here I use 6,000 bootstraps (all standard errors are below 0.010). The results are shown in Figure 1.

The clusters supported by bootstrapping include (i) happy noses, (ii) other happy faces, with subclusters around eye shape and mouth shape, (iii) a third happy cluster with tongues and big grins, with some subclusters, (iv) noseless unhappy faces, (v) other happy with equal-sign eyes forming a further subcluster.

There are several things to observe about this clustering, the first and perhaps greatest is that the basic division is between pleasant and unpleasant emotions: smiles, winks, and tongues versus frowns and slanted mouths. The division between positive and negative affect really is basic enough to the enterprise that if our methods didn’t recover this distinction, we would be skeptical that they had any worth at all.

At the beginning of the chapter, I examined emoticon variation at the author-level. The present analysis uses a different lens: words, not authors. The author-level analysis only looked at authors who used emoticons at least ten times. The word-level analysis here uses a different set of assumptions, considering a different part of the data. Nevertheless, the findings are similar. For example, noses don’t go with non-noses and eyes, mouths, and face-direction seem to matter, too. That is, noses consistently pattern separately from non-nose variants—not just for smiles but for frowns, slants, winks, and tongues, too.[6] Equal eyes pattern with noses and not with the more prevalent non-nose versions, with the exception of =] and :] which pattern together. Simple ;) and :P pattern separately (though the patterns for other wink/tongue variants are a little harder to describe).

[6]Note that some mobile phones require users to include a nose if they want the emoticon to be rendered graphically (as a straight-facing yellow/green smiling face, for example, rather than as punctuation on its side). This is not controlled for in this study, but should pose no problem. If you make the (odd) assumption that mobile phone use and/or desire for yellow faces instead of punctuation are randomly distributed across social categories, then there is decidedly no effect. If you make the (more reasonable) assumption that such things are not randomly distributed, then it is simply part and parcel of what I am talking about: the meaning of an emoticon is partly who uses it. Part of "who" includes traditional categories like gender, age, and race, but also things like personality and purchasing decisions; all of which are wrapped up in each other.

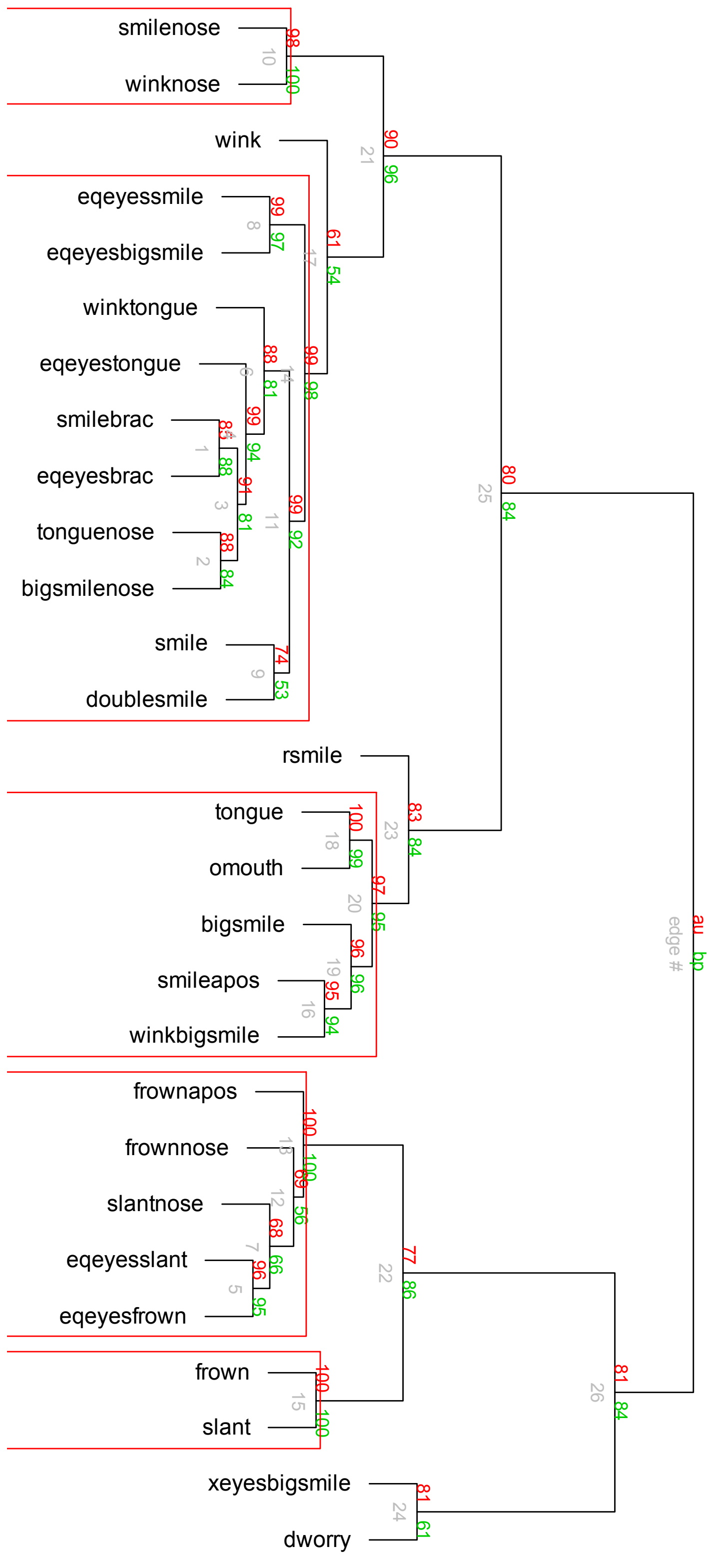

Figure 1: Cluster dendrograph with approximately unbiased and bootstrap probability p-values. Red boxes and scores above 0.95 indicate clusters that are significant using resampling methods. Distance is Euclidean. The method is “Ward.”

5 What does Presence/Absence of a Nose Mean?

We have found that noses and non-noses cluster separately whether we analyze them from the author-perspective or the co-occurring-word perspective. But surely :) and :-) mean the same thing, don't they?

In Schnoebelen 2012, I report results from restricting clustering to how the emoticons pattern with "emotion" words and find that the important affective dimensions are characterized by positive/negative, flirting, and by teasing. Those findings show that when we restrict the analysis to emotion terms, noses do not make a particularly important contribution. Yet when we use all words (as in the present analysis), we see that the nose and non-nose variants clearly pattern differently. In this section, I pursue the correlates of nose presence/absence.

5.1 Reduction due to Frequency

We might suspect that people who use emoticons a lot don't use noses as much. This is true. People who use emoticons in 250 or more tweets, use noseless variants of smiles/winks/tongues/frowns/slants more than those who don't use emoticons very often: people with 250+ emoticons use noses 6.2% of the time, while people who use 10–15 emoticons use them 14.8% of the time (p = 1.623e-28 by t-test).

This is also true if we restrict ourselves to users who have some variation—that is, people who *do* use nose-variants between 10–90% of the time. Of these people, the ones who use emoticons a lot—250+ times—use noses 33.4% of the time, but the people who use only 10–15 emoticons use noses 41.2% of the time (p = 0.00352 by t-test).

5.2 Functional Reduction

Twitter imposes a 140-character limit per tweet. So perhaps noses are really just thrown out to make room. If this is so, then we'd expect that tweets with :) should be longer than tweets with :-) because more people were feeling the pressure to save a character.

But this isn't what happens. If we take 30,000 random :) tweets and 30,000 random :-) tweets and assess the number of characters, we find that people who use noses are writing *more*, not less. The average number of characters in a tweet with :-) is 68.2, while it is 65.5 for :) (p = 2.96e-21 by t-test). Another way to say this is that people who leave off the noses are shortening a lot of other things, too.

5.3 Stylistic Differences

One of the most interesting length phenomena on Twitter is "expressive lengthening," for example *sooo, hahahaaa, heeeey, yayyyy, lollll, yummm*.

There are 467 words like this in our data set. If our finding that noseless users are shortening other things holds, then they really shouldn't be lengthening these. But they are. The average OE value for no-noses is 1.06—that is, non-noses like to go with expressive lengthening. The average OE for noses is 0.90—that is, noseless emoticons don't go with expressive lengthening.[7]

Wikipedia offers a set of 4,424 misspellings, including a shortlist of the most common and a longer list of the most difficult.[8] 49 of the "common" words are attested in our corpus (*tommorow, and, no one, thru*). Across smiles, frowns, winks, tongues, slants, and big smiles, the average OE for no-noses is 1.15. That is, no-nosers are unusually attracted to these. The average OE for noses is 0.968. That is, nosers are slightly constrained against these (mis)spellings.[9] Of the 1,016 "tough" words (*acquired, atheist, Connecticut, definitely*), the average OE for no-noses is 1.00 ("as expected"), but the average OE for noses is 1.12 (more than expected). In other words, no-nosers like these harder-to-spell words and spell them correctly.[10] Finally, there are 26 contrac-

[7] The difference is significant by a two-tailed t-test (p = 8.859e-32).

[8] http://en.wikipedia.org/wiki/Wikipedia:Lists_of_common_misspellings/For_machines

[9] A difference that is significant by t-test (p = 1.46e-07).

[10] The difference is significant by t-test (p = 2.142e-93).

tions that we can test—words like *wasn't/wasnt*.[11] People who use noses avoid dropping the apostrophe (average OE of 0.709) while noseless people drop the apostrophe (average OE of 1.170).[12]

In fact, people who use noses don't seem to abbreviate, misspell, or make typos as often. So the noseless people have longer messages despite the fact that the non-nose users are spelling things out and avoiding expressive lengthening.

There is another dimension of "standard" language we can try out: taboo words. I assemble 155 taboo words and curse words from a variety of sources—words like *f*ck, shit, jizz, damn, hell, skank, fricken*. The average OE of both groups suggests they like these words, but the rate is much higher for non-nose users: 1.14 for them compared to 1.026 for nose users.[13]

5.4 Twitter Celebrities

One way that people use Twitter is to mention and message celebrities using the same @username behavior they do for friends and acquaintances. That gives us the opportunity to distinguish nose-users from non-users by which public figures they are talking to/about. In order to limit our scope to major public figures, we can turn to Twitaholic, which lists the 1,000 most-followed Twitter accounts overall.[14] Of these people, 96 have interesting emoticon patterns in our data.[15]

The first distinction to make is that non-nose users are more positively inclined towards celebrities. The average OE for celebrities in terms of :), :D, ;), and :P is 1.0711, while the comparable OE for :-), :-D, ;-), and :-P is 0.9324. Non-nose use of positive emoticons goes with celebrities while nose use is constrained against occurring with celebrities.[16] When we take a look at negative emoticons, we find that noseless :(and :/ are constrained with celebrities; the average OE is 0.8892. The average OE for :-(and :-/ is 1.0120, a slight bump in favor of @'ing celebrities with negative noseful emoticons.[17]

It may be useful to distinguish which celebrities are treated the most differently between nose and non-nose groups. Looking at the differences between nose and non-nose variants of the positive emoticons, the following celebrities are especially associated with non-nose use:

- @jennettemccurdy—actress/singer known for her Nickelodeon sitcom, *iCarly* (19 y/o)
- @justinbieber—singer originally discovered on YouTube (17 y/o)
- @arianagrande—actress/singer/dancer best known for a role in the Nickolodeon sitcom *Victorious* (18 y/o)
- @jonasbrothers—musical band of brothers made famous on the Disney Channel (ages 19, 22, 24)
- @msrebeccablack—singer made famous by her (dreadful but slickly produced) YouTube hit, "Friday" (14 y/o)
- @officialjaden—son of Will Smith and Jada Pinkett Smith, raps and acts (the recent remake of *The Karate Kid*, for example; age 13)
- @selenagomez—actress/singer made famous for the Disney Channel's *Wizards of Waverly Place* (age 19)
- @mileycyrus—singer/actress who gained fame in Disney's *Hannah Montana* show (19 y/o)
- @jasminevillegas—singer, part of Justin Bieber's world tour (18 y/o)
- @chrisbrown—singer famous for hit single "Run It!" (22 y/o)

---

[11] We have to leave out *I'm*, *I'll*, and *we'll* because dropping the apostrophe results in them just being different words. By contrast, there is no word *werent.*

[12] Significant by *t*-test (p = 3.199e-12).

[13] Significant by *t*-test (p = 1.48e-05).

[14] http://twitaholic.com/

[15] In the entire dataset, there are 1,671 different people who are @-mentioned. 381 of them have more than 100 @-mentions.

[16] The difference is significant by *t*-test (p = 9.45e-05).

[17] The difference is significant by *t*-test (p = 1.66e-03).

The chief celebrity for noses-not-non-noses is @pepeaguilar, a Texan singer (age 43) but there are fewer other ones we can make a distinction for. Among the celebrities that are at-chance for non-noses but over-represented for noses are @mashable (a technology news site), @jessemccartney (singer and soap actor, age 24), @craigyferg (comedian and late-night talkshow host, age 49), @aplusk (actor Ashton Kutcher, age 34), and @jlo (actress/singer, age 42).[18]

By now it's probably clear what's going on: non-nose users are younger than nose users. They keep up with a younger set of celebrities (sending them positive, not negative vibes); they use more taboo words, more expressive lengthening, more non-standard spellings, and they use emoticons a lot more, too.

Measurement	**Corresponding variant**
More frequent use of emoticons	No nose
Expressive lengthening	No nose
Common misspellings	No nose
Contractions without apostrophes	No nose
Taboo words	No nose
Young celebrities	No nose
Longer tweets	Nose
Correctly spelled "difficult words"	Nose

Table 2: Summary of characteristics distinguishing noses from non-noses among American English tweets.

Another way of putting this is that non-noses orient to the less standard and noses to the more standard. If this is part of what's happening then there is an oppositional aspect. Historically, emoticons with noses came first—that means that they themselves were "standard" for a while. Given an orientation to non-standard, it is inevitable for them to be changed—for example, the elimination of the nose (or the replacement of colons/parentheses). Had the first emoticons *not* had noses, we would expect people interested in non-standardness to add them, instead.

6 Concluding Thoughts

The previous section began with a question: nose/non-nose variants of an emoticon mean the same thing, don't they? The answer requires us to get a handle on "meaning" and this is tough to do because it is an emergent property of social relations, not something that an object or a symbol has in and of itself. A sonogram or a bouquet of roses are meaningful because there are patients, doctors, lovers, and florists to give them meaning. What we usually mean by "meaning" is an interpretation that is shared by people we're familiar with using familiar interpretive schemes. Our inquiry is a lot more tractable if we shift from a concern about "meaning" to asking what ranges of interpretations are conventionally hooked to particular linguistic resources and how the different interpretations are distributed across people and contexts. To look at use (and non-use) will always involve looking at both who and when. In the immediate case, people are distinguishable by their use or non-use of noses. And this carries over to how they use their emoticon of choice—they use them with different words. I have pursued nose-variation out of a commitment to defining the meaning of emoticons not only in terms of affect but also in terms of who and when and how.

[18]If you take off the @, you often get a proper name like "Oprah" or "ABC." If we look at words that are tagged as proper nouns—not @'s—but which are in the top 1,000 most followed tweets, there is more evidence to distinguish the nose users from the non-nose users. The nose users mention *google, lakers, cnn*, and *disney* a lot more, while the non-nose users mention *twitcam, ustream, youtube, facebook, etsy*, and *mtv* more.

References

Bamman, David, Jacob Eisenstein, and Tyler Schnoebelen. 2012. Gender in Twitter: Styles, Stances, and Social Networks. Ms.

Benjamini, Yoav, and Yosi Hochberg. 1995. Controlling the false discovery rate: a practical and powerful approach to multiple testing. *Journal of the Royal Statistical Society* 57:289–300.

Bolinger, Dwight. 1989. *Intonation and Its Uses: Melody in Grammar and Discourse*. Stanford: Stanford University Press.

Carlson, Jonathan, David Heckerman, and Guy Shani. 2009. Estimating False Discovery Rates for Contingency Tables. Microsoft Corporation TechReport.

Chang, Nieng-chuang. 1958. Tones and intonation in the Chengtu Dialect (Szechuan, China). *Phonetica* 2:59–85.

Fonagy, Ivan, and Klara Magdics. 1963. Emotional patterns in intonation and music. *Zeitschrift Für Phonetik* 16:293–326.

Frick, Robert. 1985. Communicating emotion: the role of prosodic features. *Psychological Bulletin* 97:412–429.

Gimpel, Kevin, Nathan Schneider, Brendan O'Connor, Dipanjan Das, Daniel Mills, Jacob Eisenstein, Michael Heilman, Dani Yogatama, Jeffrey Flanigan, and Noah Smith. 2010. Part-of-speech Tagging for Twitter: Annotation, Features, and Experiments. DTIC Document.

Lieberman, Philip, and Sheldon Michaels. 1962. Some aspects of fundamental frequency and envelope amplitude as related to the emotional content of speech. *The Journal of the Acoustical Society of America* 34:922–927.

Schnoebelen, Tyler. 2011. Affective Patterns Using Words and Emoticons in Twitter. Paper presented at NWAV 40, Georgetown University.

Schnoebelen, Tyler. 2012. Emotions Are Relational: Positioning and the Use of Affective Linguistic Resources. Doctoral dissertation, Stanford University.

Storey, John. 2002. A direct approach to false discovery rates. *Journal of the Royal Statistical Society* 64:479–498.

Storey, John. 2003. The positive false discovery rate: A Bayesian interpretation and the Q-value. *Annals of Statistics* 31:2013–2035.

Department of Linguistics
Stanford University
Stanford, CA 94305
tylers@stanford.edu