

Emoticon Style: Interpreting Differences in Emoticons Across Cultures

Jaram Park

Graduate School of
Culture Technology, KAIST
jaram.park@kaist.ac.kr

Clay Fink

Johns Hopkins University
Applied Physics Laboratory
clayton.fink@jhuapl.edu

Vladimir Barash

Morningside Analytics
vlad@morningside-analytics.com

Meeyoung Cha

Graduate School of
Culture Technology, KAIST
meeyoungcha@kaist.edu

Abstract

Emoticons are a key aspect of text-based communication, and are the equivalent of nonverbal cues to the medium of online chat, forums, and social media like Twitter. As emoticons become more widespread in computer mediated communication, a vocabulary of different symbols with subtle emotional distinctions emerges especially across different cultures. In this paper, we investigate the semantic, cultural, and social aspects of emoticon usage on Twitter and show that emoticons are not limited to conveying a specific emotion or used as jokes, but rather are socio-cultural norms, whose meaning can vary depending on the identity of the speaker. We also demonstrate how these norms propagate through the Twitter @-reply network. We confirm our results on a large-scale dataset of over one billion Tweets from different time periods and countries.

Introduction

The most important thing in communication is hearing what isn't said. —Peter Drucker

Body language and facial expressions can sometimes tell more about what one is trying to express than what one actually says in face-to-face interactions. Changes in vocal intonation can serve a similar purpose in exclusively spoken communications. Such cues take up an estimated 93% of everyday communication (Mehrabian 1971) and help people better communicate complex emotions like humor, doubt, and sarcasm. In text-based communication, however, these cues are not present and their absence can result in misunderstanding and confusion. The growth in computer-mediated communications has led to the use of conventions where emotion or affect is referenced pictorially using alphanumeric, punctuations, or other characters. These symbolic representations are commonly referred to as *emoticons* (Walther and D'addario 2001).

The origin of emoticons is of dispute—especially for the basic smiley :)—but most studies suggest that they appeared in the early 1980s and have since gained massive popularity (Derks, Bos, and Grumbkow 2007). Emoticons, like

nonverbal cues, help people interpret the nuance of meaning, the attitude of a conversational partner, and the level of emotion not captured by language elements alone (Lo 2008; Gajadhar and Green 2005). With the advent of mobile communications, the use of emoticons has become an everyday practice for people throughout the world. Interestingly, the emoticons used by people vary by geography and culture. Easterners, for example employ a vertical style like ^^, while westerners employ a horizontal style like :-). This difference may be due to cultural reasons since easterners are known to interpret facial expressions from the eyes, while westerners favor the mouth (Yuki, Maddux, and Masuda 2007; Mai et al. 2011; Jack et al. 2012).

In this paper, we study emoticon usage on Twitter based on complete data of tweets from the period 2006 through 2009, the first three years of this now ubiquitous microblogging platform. We focus on the macro-level trend first, and examine what emoticons are popular and how they vary stylistically. Next, viewing emoticon usage as a social norm, we study how emoticons differ across cultural boundaries defined by geography and language. We then contrast what affect categories are associated with emoticons across countries. Moving from the macro-level to the level of user-to-user interactions, we use the Twitter @-reply graph to investigate the propagation processes of particular emoticons over social links, and study how the diffusion characteristics of emoticons can help us understand which emoticons have broader appeal in a new cultural setting.

We make several interesting findings:

1. Emoticons are generally used in positive and light context, and as a result tweets containing extremely angry or anxious sentiment rarely accompanied emoticons.
2. Users continuously expand the meanings of emoticons by adopting variants of the normative forms such as :) with pictorial representations of facial features such as winks ;) , forehead =:), and nose :-). These variants are sometimes associated with different kinds of affect than their normative forms.
3. While geography matters in determining the emoticon style, language has a higher impact. In the Philippines and Indonesia, where English is in common usage along with

local languages, users utilized horizontal style emoticons as in predominantly English speaking countries.

4. European users were multi-cultural in terms of emoticon usage with both vertical and horizontal styles being employed in tweets.
5. While popular emoticons like :) and :(are adopted either spontaneously or through sources outside of Twitter, less popular ones like :P, ^^ and T_T had higher chance of diffusion through the Twitter’s @-reply friendship relationship. This diffusion may be due to influence, and occurs almost entirely within cultural boundaries.

The remainder of this paper is as follows. We start by briefly reviewing the relevant literature. We then describe the Twitter dataset and our method for extracting emoticons from tweets. The next section presents the basic analysis of emoticons in Twitter, followed by cultural boundaries and diffusion processes of emoticons. Finally we discuss implications of findings and conclude.

Related Work

Emoticons are a crucial part of computer-mediated communication (Walther and D’addario 2001). Previous work confirmed that users reading text messages with emoticons are significantly better at interpreting the precise meaning of the author than those reading messages without emoticons (Lo 2008; Gajadhar and Green 2005). Emoticons are known to be used more frequently in socio-emotional contexts than in task-oriented contexts (Derks, Bos, and Grumbkow 2007). Recent work similarly demonstrated that Twitter users are more likely to use emoticons when conversing with others than when posting status updates (Schnoebelen 2012). Emotional valence has been shown to match well, as positive emoticons were used more in positive contexts and negative emoticons, more in negative contexts. However, emoticon usage decreased when people felt extreme emotions of anger or guilt, showing a tendency to drop emoticons for emotionally intense situations (Kato, Kato, and Scott 2009).

Several studies focused on emoticon usages across different cultures. One study comparing Japanese and American emoticons based on e-mail data found that the functions of emoticons were different (Markman and Oshima 2007). American emoticons primarily established punctuation, signature, and closing of a sentence, while Japanese emoticons often had more complex shapes, mimicking offline facial expressions. Another study based on SMS data confirmed that emoticon usages varied by the gender of users, while no relationship was found across different strengths of social ties (Tossell et al. 2012).

When considering emoticons as a social norm, it is crucial to consider the effects of influence and homophily. If we assume that the usage of a particular emoticon (or a set of emoticons) falls under the culture—what is called or “beliefs, attitudes and behaviors”—then influence is the mechanism whereby one’s peers influence one’s culture, and homophily is the mechanism whereby one’s culture over time affects one’s choice of peers. Foundational studies of influence (Axelrod 1986) and homophily (Lazarsfeld and R.K.Merton

1954) analyze these mechanisms on their own, but more recently researchers have begun to look at how influence and homophily interact (Axelrod 1997), especially in the context of social media to lead to complex effects on the behavior and link patterns of individuals.

In this paper, we focus on influence as the mechanism of interest when it comes to the adoption of emoticons. We use a sequential adoption model that is simpler than the dynamic matched sample approach in (Aral, Muchink, and Sundararajan 2009). Our formal approach is motivated by models studying the effects of triadic closure in Twitter (Romero et al. 2011).

Methodology

Twitter Data We use a corpus of the Twitter data in (Cha et al. 2010) from 2006 to 2009, which contains information about 54 million users and all of their public posts. Since we are interested in studying cultural differences, we tried to include both eastern and western countries that had significant Twitter populations. We classified users and their tweets into our list of countries based on their geo-location (Kulshrestha et al. 2012) and excluded some countries such as Brazil due to difficulty in processing their language. Table 1 shows the list of countries that we focus on in this study and their population proportion within the Twitter corpus.

Country	Language	Culture	Population
US	English	Western	57.74%
UK	English	Western	7.33%
Canada	English	Western	3.91%
Australia	English	Western	2.62%
Germany	German	Western	2.12%
Indonesia	English	Eastern	1.46%
Japan	Japanese	Eastern	1.45%
Netherlands	Dutch	Western	1.16%
Philippines	English	Eastern	0.97%
France	French	Western	0.83%
Italy	Italian	Western	0.65%
Spain	Spanish	Western	0.62%
Mexico	Spanish	Western	0.52%
Singapore	English	Eastern	0.48%
South Korea	Korean	Eastern	0.30%

Total data analyzed: 10 million users and 1.1 billion tweets

Table 1: List of countries and their data studied

Extracting Emoticons We limited our focus to only those emoticons expressing human facial cues, and compiled a list of candidate emoticons from a number of sources including the Wikipedia¹. Based on the compiled list, we constructed regular expressions to search our dataset.

As mentioned earlier, eastern and western countries employed different emoticon styles, as highlighted in Table 2. The horizontal style, popularly used in western countries, emphasizes the mouth for expressing emotion and commonly uses the colon sign (:) for the eyes. Different mouth shapes are used to express affect (e.g., positive, negative) and meaning (e.g., happy, sad, surprise). In contrast, the vertical style,

¹http://en.wikipedia.org/wiki/List_of_emoticons

Style	Normative form	Affect	Meaning	Variant examples
Horizontal (expression based on the mouth shape)	:)	positive	happy	wink ;))
	:(negative	sad	mouth :() :(((
	:o	neutral	surprise	nose :-) :- (:- [
	:P	positive	tongue sticking out	tear :' (:*(
	:D	positive	laugh	forehead or hair ;:(=:-)
Vertical (expression based on the eye shape)	^^	positive	happy	chin (^^)
	T_T	negative	sad	mouth ^__^ T__T
	@@	neutral	surprise	nose ^.^ ^.^ T.T
	--	negative	absent-minded	sweat ^^; --;;
	o.o	positive	curious, amazing	eyebrow --^

Table 2: Two different styles of emoticons: horizontal (popular in western countries) and vertical (popular in eastern countries).

popularly used in eastern countries, emphasizes the eyes for expressing emotion. The underscore character (_) is commonly used for the mouth, while various characters are used for the eye shapes to capture affect and meaning. The following characters were used for the mouth and eye shapes in the regular expressions:

Mouth variants: () { } D P p b o O O X # _ -
 Eye variants: : ; ^ T @ - o O X x + = ; ;

In addition to these basic facial cues, we captured the variants of each normative form in the regular expression, which we discuss in more detail in a later section.

Inferring Affect From Tweets In order to quantitatively measure what kinds of affect are associated with a given emoticon, we used LIWC (Linguistic Inquiry and Word Count) (Tausczik and Pennebaker 2010), which is a text analysis program that counts words in various psychological categories. LIWC supports many languages including English, French, Italian, and Spanish, all of which appear in our data.

How are emoticons used in Twitter?

We first present the overall emoticon usage patterns.

Emoticon Usage Table 3 displays the number of tweets, mentions, and retweets from our Twitter corpus. In total 7% of all tweets contained at least one emoticon, where nearly half of them (52%) were used in mentions of others appearing with the @username mark, while only few of them (4%) were used in retweets. This means that emoticons were more popularly used in conversations than in information propagation.

	#Tweets	#Retweets	#Mentions
Non-emoticon	1,624,968,457	52,501,839	507,177,878
Emoticon	130,957,062	2,369,449	67,656,408
Total	1,755,925,519	54,871,288	574,834,286

Table 3: Summary of dataset

The fraction of emoticon tweets starts at nearly zero during the first few months after Twitter’s launch, then since June 2007 increases slowly to reach 4%–8% of all tweets (Figure 1). Its usage remains rather steady from 2009, possibly




Figure 1: Emoticon usage over time

because emoticons are not specific to Twitter (i.e., have existed since the 1980s) and also may be because it is the very rate of emoticon usage in online conversations in general. For comparison, we also show the fraction of mention tweets over the same time period, which is a Twitter-specific convention and it shows a rapid trend of adoption from 0% to 34% over the years. While 23% of all users posted at least one tweet with emoticons, 80% of the heavy users posting more than 100 tweets had used emoticons. This implies that emoticon usage is more prevalent for heavy users. Having seen the overall trend, we next focus on the different styles, variants, and context of emoticons.

Emoticon Style As we discussed earlier in Table 2, there are two kinds of emoticon styles: vertical and horizontal. This division is based on which facial part carries the meaning. Because there lacks a systematic division of the two types, we propose to investigate the different emoticon styles as summarized in Table 2. Emoticons can be in either a normative form or a variation of that form, where the normative form for the horizontal style has a colon (:) as the eyes and one mouth. All other changes to this normative form can be considered variants. For the vertical style, the normative form is defined by the shape of the eyes and by default does not contain mouth. We allow the mouth to appear in the normative form as in (T.T), in cases where the normative form without a mouth (TT) has an ambiguous meaning.

Based on these definitions, we captured a total of 15,059 different kinds of emoticons from the Twitter corpus. Their popularity distribution was heavy-tailed as shown in Figure 2(a), so that a small fraction of emoticons had a disproportionately large share of all usages. Only 523 emoticons




Figure 2: Popularity distribution of emotions


appeared more than 1,000 times and 76% of all emoticons appeared fewer than 10 times.

The most popular emoticon is the horizontal smiley :), which appeared in 46 million tweets. Most emoticons in the top 10 list are horizontal styles except for ^^, indicating a natural bias towards the US and other English-speaking countries. Therefore, we try to address this limitation in the next section by delving into emoticons used by people from different countries, including countries where English is not the dominant language.

The top 10 emoticons took up about 43% of all emoticon tweets. As the top list in Figure 2(b) shows, not only the normative forms, but diverse variants like a wink with a nose, ;-), were also used popularly in tweets. Low ranked emoticons were mostly variants of the normative forms.

Variants Variants of the normative forms depicted facial features such as nose, tears, hair, chins, and eyebrows (see Table 2). A common variant for both horizontal and vertical styles was the lengthening of mouth like :)) and T__T. This may be related to a phenomenon, where people lengthen words to emphasize their sentiment, as in “cooooooooooooooooooolllll” (Brody and Diakopoulos 2011). Emoticons may have evolved to incorporate this convention, where people repeat the mouth to indicate a stronger affect while denoting the same meaning.

The sweat drop variant in the vertical style was popular in Japan and South Korea. The particular variant expressed feelings of shyness, embarrassment, confusion, or shock (for example, ^^; or ^_^; and -_-;). The sweat marks are thought to have originated from Japanese anime, where characters in anime often exhibit large sweat drops on their heads or beside their eyes in embarrassing moments as depicted in Figure 3.



(a) Smiley (^_^;) (b) Absent-minded (-_-;)

Figure 3: Anime characters with a sweat drop. Images from (a) Pocket Monsters and (b) Wagamama Fairy: Mirumo de Pon!

Context Given the wide range of variants, we sought to investigate how their meanings differed from the normative



Figure 4: Word clouds of the representative emoticons

forms. In particular, we wanted to know to what extent the meaning of a normative form like :) changed with a nose :-;) or a wink ;) . To investigate this, we focused on the affect categories defined by LIWC² for different words. We randomly selected 10,000 tweets containing each emoticon and extracted words that appear in the affect category. A sample of 10,000 emoticon tweets ensures a margin of 1% or less error with a 95% confidence interval. Figure 4 shows the top 50 co-appearing word stems associated with each emoticon, where word stems such as ‘amaz’ and ‘funn’ mean ‘amazing’ and ‘funny’ in tweets. The size of each word stem is adjusted to reflect its frequency.

At a glance we see that all six emoticons are used with both positive and negative affect words. In fact 10% of the word stems are common to all six emoticons. For example, ‘haha’ is the most frequently co-appeared word stem for most emoticons. This word is not only used with positive emoticons also with negative emoticons like T_T, :(and -_-;. Other positive words, e.g., friend, wow, sure, or funny, are also used with both affects emoticons. Likewise, negative words such as kill, damn, fail, and bitch are also used with positive affect emoticons. This contrast between the affect of the text and the nonverbal cues, i.e., emoticons, could be an indicator of

²<http://www.liwc.net/descriptiontable1.php>




Figure 5: Different emoticon usage by the country (a) Emoticon rates and (b) Multidimensional scaling for emoticon usage (square means English-speaking countries and circle means the other countries)

sarcasm. For example:

*"@bad_decisions I would have if I had any money, haha.
:(Maybe next time."*

This suggests that emoticons may help us analyze the sentiment of online text by more precisely capturing sarcasm and irony.

The Cultural Boundaries of Emoticons

We now discuss how emoticons are used differently across various countries. We calculated the emoticon usage rates based on the frequency of ten emoticons in Table 2. Figure 5(a) shows the extent to which countries differ in their rates of adopting horizontal and vertical emoticons. The yellow portion in each pie chart represents the percentage of the horizontal style, and the green portion represents that of the vertical style. English-speaking countries used horizontal style overwhelmingly. Korea most actively used vertical style for 74% of the time, while Japan used horizontal and vertical styles to a similar extent. Although Indonesia and Philippines are located in Asia, they showed similar patterns to other English-speaking countries. This means that language has stronger effect than geography, since English is in common use on Twitter in these countries. France, Germany, and other European countries had non-negligible fraction of vertical style adopters. Regardless of geography, smileys like :) and ^^ were the most frequently used horizontal and vertical emoticons, respectively.

Clustering Countries The types of emoticons that are popular in each country can be used to measure how similar a given pair of countries is. The MDS (Multi-Dimensional Scaling) in Figure 5(b) shows the distance of countries according to similarity of emoticon usage rates of each country. We calculated the Euclidean distance between pairs of countries based on their emoticon rates. If the countries had similar usage patterns of emoticons, they would have shorter distance in the resulting graph.

On the left top side of the graph, English-speaking countries are located close to each other, indicating that these countries are very similar in the usage rates of emoticons.




Figure 6: Comparison of the top two dominant languages used for the horizontal and vertical style emoticons in representative countries

Japan and Korea are on the right end of the graph, because these countries are different from the rest in their emoticon usage rates. Vertical styles such as ^^ and T_T were more popular than horizontal emoticons in Japan and Korea. Interestingly, France and Germany are more similar to these two countries than the other three Asian countries: Philippines, Indonesia, and Singapore. As a result, the clustering of countries is divided mainly by East Asia and the rest.

In order to determine whether the language of the speaker has any effect on choice of emoticon style (i.e., vertical or horizontal), we compared for each country the style of emoticons used and the dominant language of those tweets. We randomly selected 10,000 tweets containing the emoticons in Table 2 and distinguished the language of those tweets by using language detection library in Python.³ Figure 6 shows the two most dominant languages of tweets with the horizontal style emoticons and the vertical style emoticons, respectively. For many countries (including those not included in the figure), the predominant language was the same for tweets with both horizontal emoticons, and vertical emoticons. Most people used their native language, and English was the second most popular language.

However, Korea, Japan, and Netherlands show a different language pattern. While Korea and Japan most actively

³Guess-language 0.2: <http://tinyurl.com/guess-language>

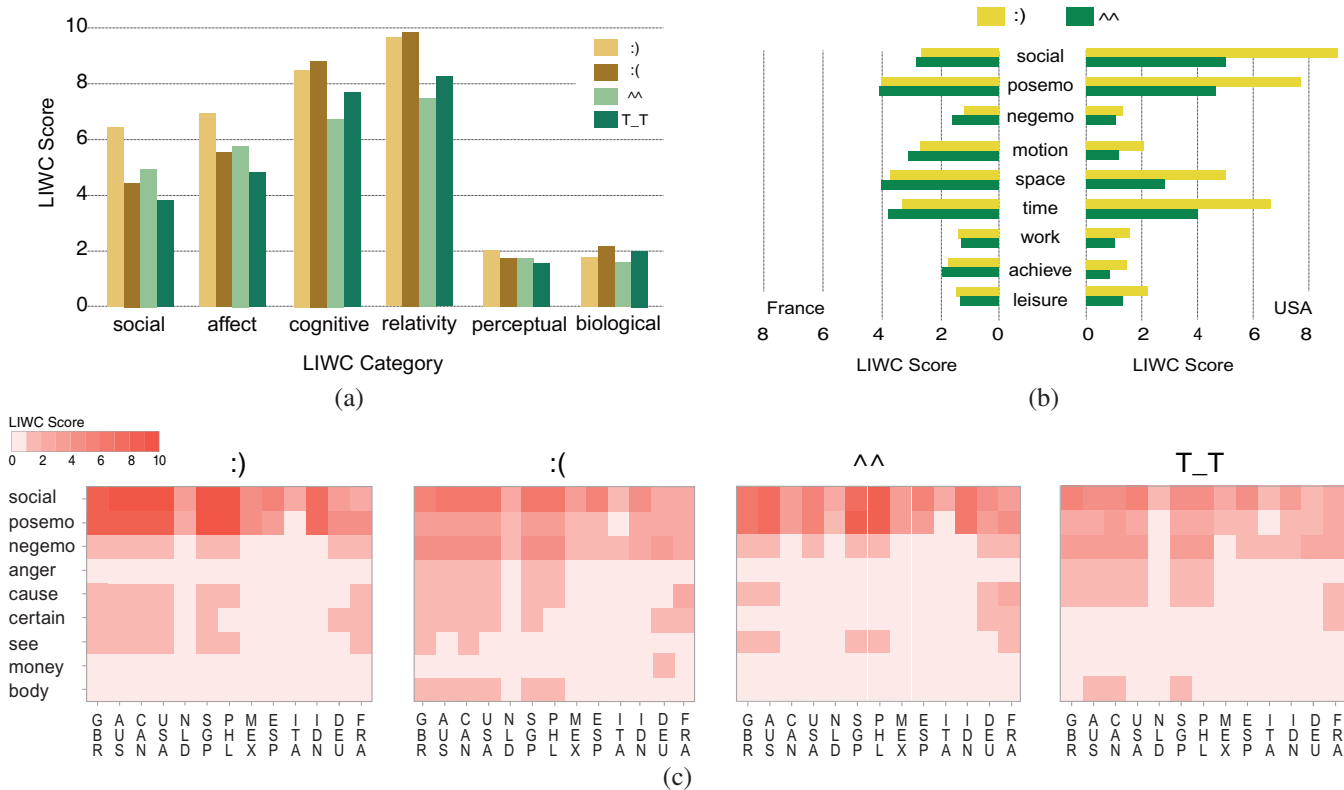


Figure 7: LIWC category scores of smiley and sad emoticons (a) Usage difference by emoticons (b) Difference usage with emoticons in horizontal and vertical styles of France and US users (c) The correlation between LIWC scores and emoticons

used horizontal emoticons in tweets written in Korean and Japanese, which are their respective mother tongues, they used vertical emoticons in tweets written in English. Twitter users in Netherlands most actively employed horizontal emoticons in tweets written in Dutch, their official language, but employed vertical emoticons in English tweets. This finding indicates that emoticon style is determined by the language of the speaker in some countries.

Words with Emoticons We next investigate the kinds of meaning each emoticon carries in different cultures. Rather than predefining their meaning, we inferred their meaning through psychological categories of LIWC. For each emoticon, we chose a random set of 10,000 tweets where it appears and examined how they are associated with various LIWC categories. We repeated this process for every country except for Japan and Korea, where LIWC libraries are non-existent. For the remaining 13 countries we compared the meaning of the same emoticon across different languages like Dutch, French, Spanish, Italian, and German. For comparison, we also chose a random set of 10,000 tweets from each country that do not contain any emoticon.

We make the following observations. First, regardless of whether a tweet contains an emoticon or not, tweets commonly contained words related to social, affect, inclusive, exclusive, space, and time categories in LIWC. However, tweets scarcely co-appeared (fewer than 1% of all tweets) with words in the anxiety, inhibition, home, money, reli-

gion, and death categories. The home, money, religion, and death categories belong to the ‘personal concerns’ higher-level category in LIWC. According to the results of analysis of variance, categories such as work, achieve, and leisure, which also belong to the personal concerns category, were used significantly more ($p < 0.001$). This result implies that people consider Twitter as a public space and avoid sensitive topics like religion and death.

Second, people used smiley emoticon like :) and ^^ when they posted tweets containing the words in the social (e.g., mate, talk, they, child), affect (e.g., happy, cried, abandon), and perceptual (e.g., observing, heard, feeling) categories in psychological process as seen in Figure 7(a).⁴ People used sad emoticon like :(and T_T when they posted tweets containing words in the cognitive (e.g., cause, know), relativity (e.g., area, bend, exit, stop), and biological (e.g., eat, blood, pain) categories. This trend was common across most countries.

Third, users from the same country even varied in the context for when they use the basic smiley in vertical and horizontal styles, as seen in Figure 7(b). We performed the paired t-test to figure out the difference of usage in both emoticon styles. The figure compares the LIWC scores of :) and ^^ in France and the US. The US shows a larger difference than France, yet both countries expressed words from different sets of LIWC categories when they used :) and

⁴7(a) show aggregated results across all 13 countries

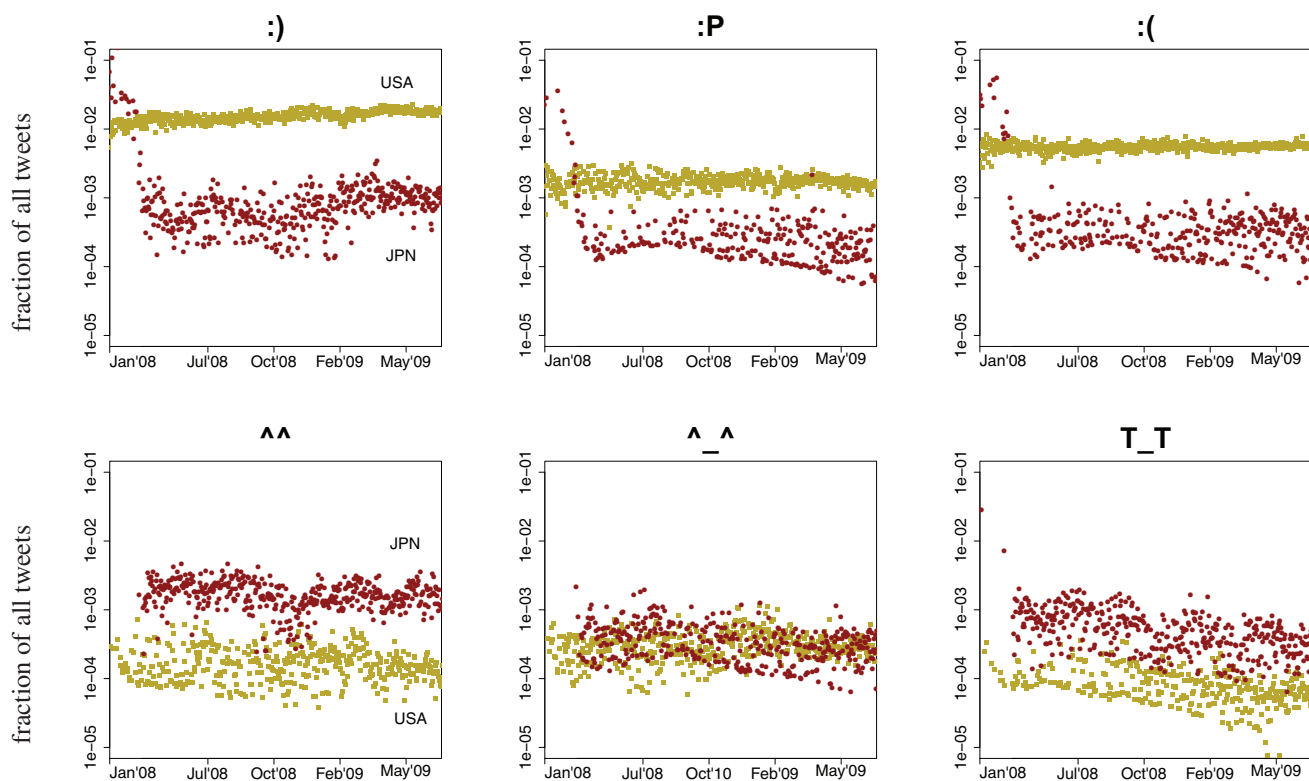


Figure 8: Fraction of emoticon-containing tweets per day, divided by total number of tweets per day from that country, for six representative emoticons. The yellow squares indicate the US and the red circles indicate Japan.

^^ emoticons ($p < 0.001$). This indicates that people employ the horizontal and vertical emoticons for different contexts although their upfront meaning is the same, i.e., smiley.

Fourth, the same emoticon was used for different contexts depending on the country. Figure 7(c) shows a heatmap for four main emoticon based on the occurrence of words in LIWC categories across the 13 countries. The map shows that the smiley emoticons like :) and ^^ are related to the words in see categories (e.g., view, seen), and the sad emoticons like :(and T_T are related to the words in the body (e.g., cheek, hands, spit) and anger (e.g., hate, kill, annoyed) categories. The map shows other interesting cultural differences. For instance, the money category came out dominant only in tweets containing :(emoticon in Germany (denoted DEU) and the Australia, Canada, and Singapore users expressed words related to body category with negative connotation, i.e., :(and T_T emoticons.

Temporal Dynamics of Emoticons

Next, it is natural to consider the temporal properties of emoticons and ask, does the usage of emoticons remain stable over time or does it change? To measure the temporal dynamics of emoticons, we examined the total number of tweets per day that contain a particular emoticon in the first 50M tweets of our data. To examine cultural aspects, we further broke down these plots by country (for simplicity, we only considered the US and Japan as representative countries of mouth-oriented

and eye-oriented emoticons). We also normalized the number of emoticon-containing tweets per day, by dividing it by the total number of tweets on that day from that country.

Figure 8 shows the fraction of all tweets per day that contain a particular emoticon, for six different emoticons. We usually consider ^_^ as a variant of ^^ and have the same meaning, but here we separate them to show an interesting trend in their usage over time.

The six emoticons are split into three groups by their dynamics: stable over time (:P, ^^, ^_^), increasing over time (:(, :), and decreasing over time (T_T). In addition, the emoticon rates for Japan fall sharply over the first three months after January 1, 2008 for some emoticons, probably reflecting the rapidly growing number of Japanese Twitter users over that time period. The y-axis of each graph is logarithmically scaled, so even a subtle trend indicates a profound change of an order or magnitude or more. After the first hundred days, changing volumes of tweets over time by country are not enough to account for the effects, since different emoticons exhibit different trends. The difference between :(and T_T is particularly intriguing, since the two emoticons represent the same meaning (sadness); the fact that :(increases and T_T decreases over time both for Japan and the US suggests that globally, :(is not only the more prevalent way of expressing sadness, but becomes ever more prevalent over time.

Finally, the relative rates of use between Japan and US mostly confirm Figure 5, with mouth-oriented emoticons

much more popular in the US than in Japan, while eye-oriented emoticons are the opposite. There is an interesting exception, however: ^ is much more prevalent in Japan than in the US, whereas ^_ is about equally prevalent in both countries. This finding shows that variants of the same emoticon can have widely different adoption rates by culture.

Diffusion of Emoticons

As we investigate the nature of emoticons as social norms in Twitter, we must consider the question: are emoticons socially transmissible, that is, do some of them diffuse through the social network of Twitter users? For any given emoticon E , we can imagine two theoretical scenarios: either E is adopted spontaneously by Twitter users, much like umbrellas are spontaneously “adopted” by people walking outside on a rainy day (without explicit diffusion), or E is transmitted from one Twitter user to another, as a rumor is transmitted in a social network. In the first scenario, E is not a social norm. In the second scenario, E is a social norm.

In practice, for any emoticon E some Twitter users will adopt it spontaneously (or else will be influenced to adopt it by sources outside their Twitter network, e.g. friends on another text-based service), while others will be influenced to adopt it by their Twitter friends. We develop a statistical test to handle this uncertainty: given the null hypothesis of spontaneous adoption, we determine for a given emoticon E whether we can confidently reject the null hypothesis and assume that E is a socially transmissible norm, or whether we cannot reject the null hypothesis and thus E is either transmitted outside Twitter, or adopted spontaneously.

In the following analysis, we define adoption of emoticon E as posting at least 3 tweets, each of which contains E . We define friendship directionally from user i to user j as i replying to j via the “@”-reply at least 2 times and j replying to i via the “@”-reply at least 2 times. In our analysis, we experimented with different posting thresholds in the range of 1 to 5, and with different reply thresholds in the range of 1 to 5. Choosing different thresholds did not qualitatively affect our results.

In performing our statistical test, we must be wary of confounding factors such as homophily and independent adoption. It is possible for two people to be friends on Twitter and to adopt an emoticon simultaneously due to external factors, much like two friends out on a walk may open their umbrellas at the same time, because it starts to rain. Furthermore, it is possible for two people to adopt the emoticon and then to become friends due to the principle of homophily, which states that similar people are more likely to become friends (similarity can be defined as using the same emoticon).

Being aware of these confounds, we set up the following definitions, inspired by related work in triadic closure on Twitter (Romero et al. 2011). We consider “adoption possibly due to influence” to be one of two series of events: either, Twitter user i who has not adopted emoticon E becomes friends with Twitter user j who has adopted E , and subsequently i adopts E ; or, Twitter user i who has not adopted emoticon E becomes friends with Twitter user j who has not adopted E , then j adopts E and finally i adopts E . This definition enforces that both i ’s friendship with j and j ’s

adoption of E are prior to i ’s adoption of E , a condition that is necessary for influence to be a possible cause of adoption.

Similarly, we can define “adoption not due to influence” to be the logical negation of “adoption possibly due to influence”, that is, the series of events where Twitter user i adopts E having no friends who have already adopted E . This definition sets up the condition wherein it is impossible for influence to have led to adoption.

The advantage of this pair of definitions is that it allows us to create a statistical test for influence being a cause of the adoption of E . In any particular case, determining the sufficient condition for influence to be a cause of adoption requires knowledge of i and j ’s mental states, which is unfeasible for large N . However, aggregated over thousands of cases, a pattern whereby the likelihood of “adoption due to influence” is statistically significantly higher than the likelihood of “adoption not due to influence”, we are armed with statistical evidence that influence has led to adoption.

The statistical test we use is as follows. First, we construct four sets:

$A(E)_I$ all users i who have adopted E possibly due to influence, by adopting E either after making friends with a j who had already adopted E or after their current friend j adopts E

$NA(E)_I$ all users i who have not adopted E despite being possibly exposed to influence, by not adopting E despite having a friend j who has adopted E

$A(E)_{NI}$ all users i who have adopted E not due to influence, by adopting E despite not having any friends who have already, at that moment, adopted E

$NA(E)_{NI}$ all users i who have not adopted E and were not exposed to influence, by not adopting E and having no friends who have already, at that moment, adopted E

Having constructed these four sets, we can determine the probability p_I of adopting possibly due to influence, and the probability p_{NI} of adopting not due to influence:

$$p_I = \frac{|A(E)_I|}{|A(E)_I| + |NA(E)_I|}$$

$$p_{NI} = \frac{|A(E)_{NI}|}{|A(E)_{NI}| + |NA(E)_{NI}|}$$

Finally, having the respective probabilities, we can ask whether p_I is statistically significantly higher than p_{NI} . The simplest test of this condition is a binomial test where the number of trials is $X_I = |A(E)_I| + |NA(E)_I|$, the observed number of successes is $|A(E)_I|$ and the success probability is p_{NI} . Table 4 lists p_I , p_{NI} , X_I , and the significance level that $p_I > p_{NI}$ if appropriate for the emoticons listed in Table 2.

As the table shows, for all the emoticons except for :) and :(we may reject the null hypothesis of spontaneous adoption. This finding is ecologically valid: :) and :(are very popular emoticons, so it is reasonable to assume that Twitter users either are exposed to them outside Twitter (in chat, forums) or come up with them spontaneously, because

Emoticon	p_I	p_{NI}	X_I	Significance
:)	0.086	0.095	36815	-
:(0.058	0.064	32309	-
:P / :p	0.049	0.029	21324	***
:D	0.050	0.033	21067	***
:o	0.010	0.003	3017	***
^^ / ^^	0.027	0.014	7530	***
T_T / TT	0.013	0.004	1924	***
--	0.010	0.004	2396	***
@_@ / @@@	0.012	0.001	859	***

Table 4: Diffusion of emoticons. Significance levels legend: “-” means not significant at the $p < .05$ level, “*” means significant at the $p < .05$ level, “**” means significant at the $p < .003$ level, and “***” means significant at the $p < .00007$ level.

they are easy to invent as they consist of only two symbols. In contrast, emoticons like T_T / TT or @_@ / @@@ are much less popular, so it is reasonable to assume that Twitter users are first exposed to them by their friends on Twitter, and are influenced to adopt them as a social norm, or by imitation.

Cross-Cultural Diffusion An interesting follow-up question to our diffusion analysis is, does diffusion happen across cultural boundaries? One would expect, given that emoticons like :) are much more common in the US and Europe than in Asian countries, while emoticons like ^^ / ^^ are much more common in Asian countries than in the US and Europe, that diffusion would flow mostly within cultural boundaries. Still, as the example in Figure 8 shows, sometimes two variants of the same emoticon are used in different cultures, and one culture adopts the other’s variant. Furthermore, given the widespread use of Twitter and lack of restriction by country boundaries, communication (and thus, potentially, diffusion) of norms from one country to another is possible.

We investigate diffusion across cultural boundaries by further breaking down the set $A(E)_I$ from the previous section into users who adopted across cultural boundaries, that is, all users i who adopted E either after becoming friends with a j who was already using E , or after their current friend j adopts E , provided i and j come from different cultural sets; and users who adopted along cultural boundaries, that is, the same adoption scenario, but provided that i and j come from the same cultural set. For the purposes of this analysis, we define a North American / European cultural set as the US, UK, France, Italy, and Germany, and an Asian cultural set as Japan, China, and Korea. The results are listed in Table 5. We note that the number of adoptions is so small because we are using only a subset of all countries, and because we do not have complete mapping of users to countries.

As the table shows, cross-cultural emoticon adoption is extremely rare; even though Table 4 demonstrates that many emoticons diffuse from one Twitter user to their friends, in most cases, these friends are from the same culture.

Conclusion

In this paper, we examined the use of emoticons on Twitter. Emoticon styles can be either horizontal or vertical, where

Emoticon	Same culture	Cross-cultural	Fraction
:)	177	5	0.028
:(81	0	0.000
:P / :p	72	2	0.028
:D	86	0	0.000
:o	1	0	0.000
^^ / ^^	42	1	0.024
T_T / TT	4	0	0.000
--	7	0	0.000
@_@ / @@@	0	0	0.000

Table 5: Diffusion of emoticons across cultural boundaries

horizontal style is known to be preferred by western countries, and the vertical style by eastern countries. This study finds that an important factor determining emoticon style is language rather than geography. Regardless of their inherent meaning, most emoticons co-appeared with both positive and negative affect words (e.g., haha, smile, kill, freak). Furthermore, the contexts and sentiments that were frequently associated with a given emoticon varied from one culture to another. Our finding confirms that facial expressions may not be universal (Jack et al. 2012); people from different cultures perceive and employ facial expressions in unique ways, as easterners smile and frown with their eyes, whereas westerners do so with their mouth. This was even true in the online world. Therefore one might want to consider the cultural background of one’s followers to communicate efficiently in online social networks.

We also find that many emoticons diffuse through the Twitter friendship network, which suggests that emoticons are used as social norms, and that Twitter users may influence their friends to adopt particular styles of emoticons especially for less popular variants. Given the culturally idiosyncratic relationship we found between emoticons and emotion words, and the extreme rarity of cross-cultural emoticon diffusion, it is possible that emoticons are evolving from a universal way of expressing faces in text to culturally-bounded emotional dialects, much as many natural languages have evolved from a common desire to communicate into culturally-mediated forms of expression and interaction. However, more research into the use of emoticons is necessary to support this hypothesis; we hope to pursue such research in future work.

Emoticons are a critical part of nonverbal communication, taking the place of body language and facial expressions in text-based media. Text analysis methods used everywhere from search to summarization to opinion extraction can leverage emoticons to extract subtler shades of meaning from tweets, blog posts, and online reviews alike. One practical implication of our research is that automated text methods should, if possible, model the cultural context of a particular expression and not simply assign a static meaning to an emoticon character. Cultural modeling can help tell the difference between “:)” used to express positive emotion, or “:.” used as a purely social marker, like hello. Another practical implication of our research is that emoticons seem to diffuse readily within cultural boundaries, but rarely across them; thus, a user from one culture might be completely un-

