What I See is What You Don't Get: The Effects of (Not) Seeing Emoji Rendering Differences across Platforms

HANNAH MILLER HILLBERG, GroupLens Research, University of Minnesota, USA
ZACHARY LEVONIAN, GroupLens Research, University of Minnesota, USA
DANIEL KLUVER, GroupLens Research, University of Minnesota, USA
LOREN TERVEEN, GroupLens Research, University of Minnesota, USA
BRENT HECHT, Northwestern University, USA

Emoji are popular in digital communication, but they are rendered differently on different viewing platforms (e.g., iOS, Android). It is unknown how many people are aware that emoji have multiple renderings, or whether they would change their emoji-bearing messages if they could see how these messages render on recipients' devices. We developed software to expose the multi-rendering nature of emoji and explored whether this increased visibility would affect how people communicate with emoji. Through a survey of 710 Twitter users who recently posted an emoji-bearing tweet, we found that at least 25% of respondents were unaware that the emoji they posted could appear differently to their followers. Additionally, after being shown how one of their tweets rendered across platforms, 20% of respondents reported that they would have edited or not sent the tweet. These statistics reflect millions of potentially regretful tweets shared per day because people cannot see emoji rendering differences across platforms. Our results motivate the development of tools that increase the visibility of emoji rendering differences across platforms, and we contribute our cross-platform emoji rendering software[1] to facilitate this effort.

CCS Concepts: • **Human-centered computing → Empirical studies in collaborative and social computing**

KEYWORDS

Emoji; computer-mediated communication; cross-platform; rendering; invisibility of system status

ACM Reference format:

Hannah Miller Hillberg, Zachary Levonian, Daniel Kluver, Loren Terveen and Brent Hecht. 2018. What I See is What You Don't Get: The Effects of (Not) Seeing Emoji Rendering Differences across Platforms. *Proc. ACM Hum.-Comput.* Interact. 2, CSCW, Article 124. (November 2018), 23 pages. https://doi.org/10.1145/3274393

1 INTRODUCTION

Emoji are "picture characters" (translation from Japanese) that have become commonplace in nearly all forms of text-based computer-mediated communication, including smartphone texting [31], social media sharing [17], and advertising [45]. Hundreds of millions of people likely interact with emoji on a daily basis, whether as authors, recipients, or both. As an indicator of the ubiquity of emoji, Oxford Dictionaries

This work is supported by the National Science Foundation Graduate Research Fellowship under grant 00039202 and by Northwestern University.
Author's addresses: Hannah Miller Hillberg, Zachary Levonian, Daniel Kluver, Loren Terveen, GroupLens Research, University of Minnesota, 200 Union Street SE, Minneapolis, MN 55455, US; Brent Hecht, Northwestern University, 2240 Campus Road, Evanston, IL 60208, US.
[1] https://github.umn.edu/emojilens/emoji-rendering-simulation

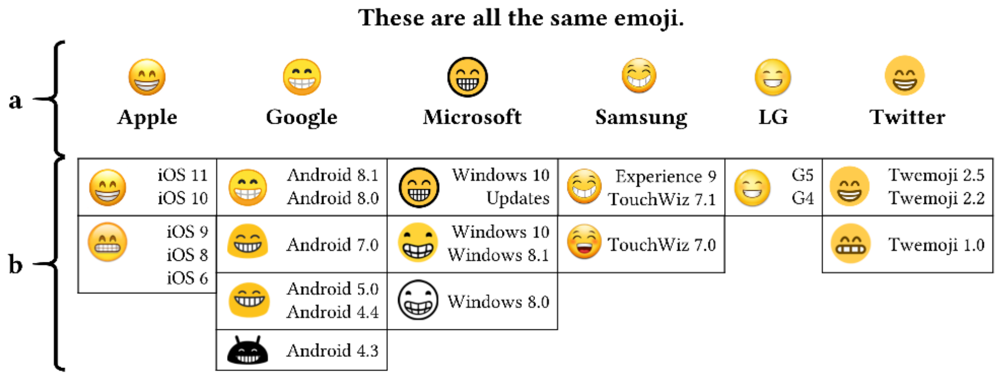

Fig. 1. (a) Vendor-specific renderings of the same "Beaming Face with Smiling Eyes" emoji (U+1F601) (b) Vendor-version specific renderings of the same emoji, in reverse chronological vertical order.

declared the "face with tears of joy" emoji ()[2] to be the 2015 "word of the year" [46], noting that "emoji have come to embody a core aspect of living in a digital world that is visually driven, emotionally expressive, and obsessively immediate" [46].

The Unicode Consortium provides a worldwide text-encoding standard for emoji characters just as it does for more traditional characters (e.g., Roman alphabet letters, numbers, Chinese characters) [47]). The Unicode standard provides a *code point* (or sequence of code points) and a name for each emoji character, but it is unlikely that people recognize emoji characters by these identifiers (i.e., is not usually described as either the "Beaming Face with Smiling Eyes" emoji or "U+1F601"). Rather, as picture characters, emoji convey their meaning through their graphics.

Graphics for emoji characters, however, are *not* standardized by the Unicode Consortium. Instead, the appearance of an emoji character is *rendered* by a font. Critically, emoji fonts are largely specific to individual technological *vendors*. This means that emoji look different on devices or applications from different vendors (e.g. Apple, Google; see Fig. 1a). In other words, when communicating with emoji, a receiver will see different emoji renderings than the sender if they are using devices from different vendors. Emojipedia, an emoji reference website, currently tracks 12 vendors that have their own emoji fonts [48].

Vendor emoji differences, however, only describe one part of the complexity of the emoji rendering ecosystem. Additional complexity is added by the fact that vendors update their emoji fonts over time along with their operating systems or applications. As such, emoji fonts are *vendor-version* specific, not just vendor-specific (see Fig. 1b). For example, a sender with an Android phone using version 8.1 of the OS would see a different rendering of the emoji in Figure 1.1b than a recipient with an Android phone using version 7.0, even though both of these devices use an operating system from Google.

We use the term *platform* to refer to a device or application using a specific vendor-version configuration (i.e., emoji font), and an emoji *rendering* is the graphic of an emoji character from the emoji font of the platform used to produce it. Emojipedia currently tracks over 50 vendor-version configurations (i.e., emoji fonts), which means any given emoji character may have 50 different renderings (though some may no longer be in use, e.g., if every person has updated their device(s) from a given vendor past a given version[3]).

Prior research has found this cross-platform *multi-rendering nature of emoji* is associated with serious potential for miscommunication [33]. This potential is due to people varying from person to person in their interpretations of different renderings of the same emoji character. It is unknown, however, whether people

[2] The emoji renderings included in the text are Apple's renderings, unless otherwise specified.
[3] This adds yet another dimension of complexity, because it is very difficult to determine which vendor-version configurations are in use (including new versions getting released).

are even aware that emoji have multiple renderings, since, while communicating, people only see the emoji renderings specific to the platform they are currently using. Critically, it is also unclear if users would change their communication behavior if they could see how their messages render on other platforms. That is, we know from prior work that people vary in their interpretations of different renderings of emoji, but we do not know if they perceive the differences between renderings to be large enough to make a difference in the context of their own communication.

The goal of this paper is to further our understanding of the real-world implications of the multi-rendering nature of emoji by addressing these open questions. Specifically, we first sought to explore whether people are aware of this characteristic of emoji. We next investigated whether showing emoji rendering differences across platforms would affect communication decisions. The stakes of these questions are significant. We know that billions of messages containing emoji are sent every day [40], and it is reasonable to expect that a non-trivial percentage of these messages are viewed on different platforms than they are sent.

In order to accomplish this paper's goal, we needed a way to show emoji rendering differences across platforms to people in the context of their own messages. However, no tool to do this currently exists. As such, we developed emoji rendering software that parses emoji from input text and accurately displays how the text would render on a wide variety of platforms. We embedded this software in an online survey deployed on Twitter so that we could use participants' own emoji-bearing tweets to expose the multi-rendering nature of emoji. We showed participants what their tweets look like across platforms and we inquired as to whether they would have changed their tweets if they had known how they appeared to followers using other platforms.

Our results provide strong evidence that surfacing the multi-rendering nature of emoji would have meaningful effects on real-world text communication. At least 25% of our 710 survey respondents were not aware that emoji have multiple renderings. This suggests that a substantial proportion of people do not know that the emoji renderings they see are not always the same as the renderings their communication partners see. Additionally, 20% of our respondents indicated that they would have edited or not sent their emoji-bearing tweet if they had known how the tweet rendered on other platforms. When we generalize to the population of all tweets that include emoji, this result means that millions of potentially regretful tweets are likely shared every day because people currently cannot see emoji rendering differences across platforms.

Our results motivate the need for new technology to better support people as they communicate with emoji. This need is exacerbated by the fact that intellectual property concerns and branding incentives will likely ensure that emoji rendering differences across platforms will persist into the foreseeable future [20]. As such, we propose building tools that provide an emoji "preview" function similar to that in our survey. These tools would give people the awareness and visibility they currently lack while communicating with emoji. To inform the design of these tools, we articulate some of the most important design considerations that emerge from our results. We are also releasing our emoji rendering software with this paper, as it can serve as a core engine for these tools.

2 RELATED WORK

Interest in emoji in the computing research community has increased dramatically in the past few years. Computing researchers have focused on topics ranging from functions of emoji [1,15,24,26,34] to emoji usage patterns [2,12,28] to sentiment classification and text understanding [3,18,21,35]. In this section, we highlight the threads of this literature that most motivated the present research. We also cover critical background information from other domains.

2.1 Emoji-Related Miscommunication

Prior to emoji, Walther and D'Addario [42] found that participants varied little in their sentiment interpretations of the emoticons ":-)", ":-(" and ";-)". According to more recent research, however, this is not the case for emoji. Miller et al. [33] examined how much people vary in their interpretations of emoji and found that this variability can be extensive in terms of both sentiment and semantics. Tigwell and Flatla [41] extended Miller et al.'s research to consider sentiment along two dimensions instead of one, finding similar results. Not only did these efforts expose that emoji have potential for miscommunication, they also explicitly examined the potential for miscommunication associated with the multi-rendering nature of emoji by comparing people's interpretations across renderings.

Both Miller et al. [33] and Tigwell and Flatla [41] studied interpretation of standalone emoji. However, the most common use case of emoji involves emoji characters being embedded in surrounding text [31]. Hypothesizing that the presence of text would reduce the observed variability of emoji interpretations, Miller and colleagues followed up their work with similar methods exploring the role of text in emoji-related miscommunication [32]. Ultimately, they found little support for the notion that emoji become less ambiguous when viewed with accompanying text.

Despite the progress made by the above literature, it remains unknown whether people are even aware that emoji have multiple renderings. If they are aware, it is also unknown whether people perceive these different renderings as sufficiently distinct to change their messaging behavior. We address these questions in this research.

The importance of understanding user behavior around emoji rendering differences across platforms is bolstered by legal and economic factors that make it unlikely that cross-platform communication will disappear anytime in the foreseeable future [20]. Specifically, vendors are incentivized to create a distinct style for their emoji renderings as a way to build brand loyalty [20], as well as to take advantage of opportunities to incorporate their brand into their emoji renderings [5]. Also, emoji renderings may be protected by intellectual property rights, including copyright, trademark, design patents, and publicity rights [20]. These factors prevent vendors from using or adopting each other's emoji renderings, thereby generating large barriers to the convergence of emoji fonts. Indeed, recent events show evidence of increasing divergence [6]: Slack, a popular communication platform for group collaboration, recently switched from rendering emoji using a single emoji font (Apple's) to rendering emoji natively (i.e., using the viewing platform's emoji font) [7]. That is, Slack went from being a within-platform communication setting (using a single emoji font) to a cross-platform setting, and this was likely due to the legal and economic factors mentioned above.

These legal and economic factors mean that any forces that are pushing towards convergence tend to be incremental, e.g., encouraging people to upgrade their operating systems to reduce inter-version communication and redesigning renderings to mitigate worst-case cross-platform issues (e.g., [4]). One exception occurs when vendors with their own emoji font also have their own social network, chat application, or other communication application that is used across platforms. In a few of these unique situations, vendors have chosen to override native platform renderings with their font's renderings within their communication application, thus making the application a within-platform communication setting. For instance, Facebook does this in Facebook Messenger. Twitter also partially implemented this approach after our research was completed, but only for a portion of its users. Specifically, for users of phones with older Android operating systems, Twitter recently began replacing the platform-native emoji font with Twitter's emoji font [8] (previously only used in the web client). This change reduces the number of cross-platform contexts on Twitter, but an enormous amount still occurs even after this change (e.g., all communication between iOS and Android, newer Android and Twitter's renderings).

2.2 Technology to Support Emoji Use

Though emoji have been shown to have potential for miscommunication, few resources exist to assist people in managing this potential. Some emoji resources have been produced in the research community, but these resources are intended to assist machines in automatic text processing rather than to assist people while they are communicating. Researchers such as Liu et al. [27] and Novak et al. [35] have developed classifiers of emoji sentiment by labelling emoji with the sentiment of the surrounding text. Both projects found instances of emoji being associated with different, and occasionally opposite, sentiment labels. Wijeratne et al. provide a similar resource *EmojiNet* [49] but for semantics [43,44]. This "sense inventory" also associates multiple meanings with individual emoji [43,44]. Importantly, these resources do not differentiate emoji renderings, only emoji characters. Wijeratne et al. [44] recognize this crucial facet of emoji, mentioning that the senses identified in their inventory for a given emoji may be more or less associated with its different renderings. In fact, Wijeratne et al. [44] explored this hypothesis for a subset of 40 emoji, ultimately finding that this was true for the majority.

The best resources that currently exist for people to consult regarding the multiple renderings of emoji are websites that maintain repositories of renderings associated with each emoji character for some set of vendors (e.g., [48–51]). Most of these websites, including the full official Unicode emoji list, do not maintain historical versions of renderings, even though many such older versions are still in use. Likewise, many such websites are outdated, given that individual vendors act independently and update relatively frequently, in some cases every few months (e.g., Microsoft updated emoji in August 2016, April 2017, October 2017, and April 2018 [52]). To our knowledge, Emojipedia.org is the most comprehensive inventory of emoji, maintaining information from most platform vendors as well as most historical versions. However, if one is using Emojipedia to look up what a given communication will look like across platforms, one can only do so out of context. Additionally, doing this look-up for each emoji in each message is of course an excessive burden given the number of emoji-bearing messages that are sent by people each day.

Overall, the lack of technology to support people as they communicate with emoji means that they are almost always "flying blind" when it comes to managing the multi-rendering nature of emoji, if they even know to manage it at all.

2.3 Invisibility of System Status

In human-computer interaction, invisibility of system status is considered a significant design flaw [53] and occurs when some computation (resulting in a change of state) happens that is not immediately apparent to users. Cross-platform emoji rendering creates large-scale system status invisibility issues in computer-mediated communication. However, cross-platform emoji rendering is not alone in this regard: this is also true of complex communication-mediating algorithms, like those used to curate news feeds. Indeed, in analogy to our research here but within the algorithmic curation space, Eslami et al. [19] studied whether Facebook users were aware that their news feed was algorithmically curated. To conduct this study, they built a system called *FeedVis* to show their participants the difference between their algorithmically curated feed and their unadulterated feed. They then used this tool to better understand users' news feed preferences.

Our research stands on the shoulders of the FeedVis project. FeedVis, which is in a different application domain but targets the same visibility of system status issues, directly inspired our research questions and our primary approach: like Eslami et al. [19], we wrote software to expose a process in computer-mediated communication that was previously invisible and used that software as a probe into people's perceptions and desires with regard to the relevant communication domain.

3 STUDY DESIGN

In this research, we operationalized our informal questions discussed above into the following formal research questions (which were motivated by those asked by Eslami et al. [19]):

1. How aware are people of the multi-rendering nature of emoji?
2. How do people evaluate the transformation of their communications when shown other platforms' renderings? Would people prefer to change their communication, given the opportunity?

Our primary goal in developing our study design was to maximize both participation and ecological validity. We wanted our results to reflect a large sample of people, so we developed an online survey[4] to collect our data. Regarding ecological validity, rather than having stock examples of emoji usage and/or selecting a limited sample of emoji characters (out of 2,666) for the survey, we wanted to use participants' own communications containing emoji.

[4] An example path through the (branching) survey is included as a PDF in our supplementary materials.

Given this goal, we decided to use Twitter as our recruitment platform because (1) emoji are very common on Twitter [35] and (2) Twitter's APIs provide automatic access to a large volume of authentic emoji use cases and their associated users. Below, we detail specifically how we recruited participants and incorporated their emoji-bearing communications (i.e., tweets) in the survey. This study was reviewed and deemed exempt from further review by the University of Minnesota's Institutional Review Board.

3.1 Twitter Recruitment

We used a recruitment approach inspired by Kariryaa et al. [25] in which potential participants whose tweets meet desired criteria (i.e., containing emoji, in our case) are detected through the Twitter Streaming API (which returns a small random sample of all public tweets). These potential participants are then targeted (as a "tailored audience") with an advertisement that requests their participation via the Twitter ads platform.

While the core of our approach comes from Kariryaa et al. [25], this approach resulted in more time delay and less throughput than was desirable for our study. We wanted the tweets seen in the survey to be as recent as possible for ecological validity purposes, so we transformed the Kariryaa et al. approach into an iterative daily cycle. Specifically, for each day during the study period, we collected potential participants for the day, then uploaded the list of participants as a tailored audience, then advertised to this tailored audience the following day, and then repeated the process. With respect to throughput, the Streaming API is an efficient means for finding tweets that match a criterion, but only when that criterion corresponds to a filter in the Streaming API. Unfortunately, there is no filter for emoji. As such, we had to instead turn to Twitter's Search API. For each emoji, we used this API to search for up to 1,000 tweets containing that emoji per day. We chose this threshold so that all 2,666 queries (corresponding to each emoji) could finish overnight, given our daily workflow. We also filtered our search so that each returned tweet met the following criteria:

- Tweet must be in English.
- Tweet cannot be a retweet (otherwise it would not be the participant's original content).
- Tweet cannot contain media. Our survey was designed to only support text tweets.
- Tweet cannot contain user mentions. The visualization of the tweet across all platforms would be less relevant if the participant was targeting specific people. (Also, it would not be possible to infer with high accuracy the platforms of those specific people, given, e.g., vendor versions and Android variants—see below for more.)
- Tweet must be sourced from within Twitter (preventing automated tweets and spam as we were targeting tweets written by people).

To gain a sense of the proportion of tweets that satisfy our filter criteria, we collected all tweets returned by the Streaming API for one week during the time of the study. From this dataset, we were able to estimate that about 1.76% of tweets satisfy all of these criteria. Of these filtered tweets, approximately 38.7% bear emoji.

To recruit participants from among the Twitter users whose tweets were returned by our Search API queries, we set up an ad campaign on Twitter designed to maximize clicks to our survey link. Advertisements on Twitter are just tweets that are "promoted," so creating an advertisement is simply a matter of creating a tweet for Twitter to surface in users' feeds. Fig. 2 shows our promoted tweet. As described above, we created new tailored audiences from the potential participants we collected daily. However, due to initially low response rates (and Twitter processing delays), we targeted our tailored audiences up to 10 days since they were collected. We additionally specified audience targeting criteria so that our ad was restricted to people over the age of 18 and who spoke English.

We advertised for a period of two weeks in the spring of 2018. Over this period, our advertisement received 1,316,460 impressions (views) and 6,838 link clicks (0.52% click-through rate). The cost was an average of $0.41 per click. Of the clicks, 1,066 started the survey by providing consent and their Twitter username. 712 went on to complete the survey. The mean time from tweet to survey completion was approximately 4.5 days. Overall, we spent just under $4 per completed survey (but, as Kariryaa et al. [25] also note, this sum is paid to Twitter, not to participants).

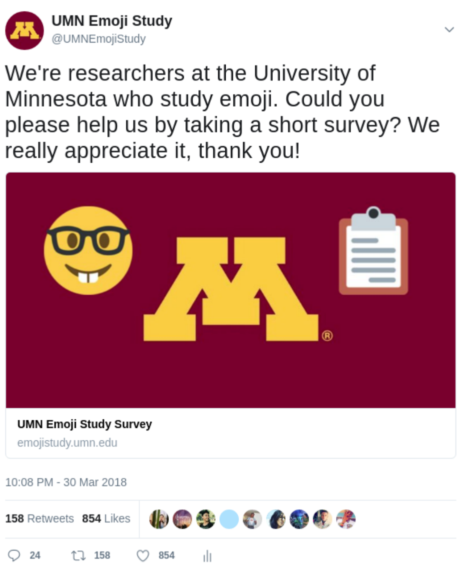

Fig. 2. Study Advertisement Tweet

Table 1. Tweet Sources

Tweet Source	N	%
Twitter for iPhone/iPad/Mac	422	59.4
Twitter for Android	269	37.9%
Twitter Web Client	19	2.7%

Table 2. Vendors of Participants' Devices

Vendor	N	%
Apple	439	61.8%
Google (Android)	144	20.3%
Samsung	111	15.6%
Microsoft	6	0.8%
LG	4	0.6%
HTC	1	0.1%
Unknown	5	0.7%

3.2 Participants

We included 710 participants' surveys in our study. We removed two completed surveys: one for irrelevant open-ended responses and one due to a small bug in our survey related to the specific device the participant was using. Of the 710 participants, 512 were female, 182 male, and 16 indicated they identified as gender non-binary (Table A5 in the Appendix). Though the Twitter user population is disproportionately male [54], this gender distribution is somewhat expected given that women use emoji more often than men [12]. Our participants also skewed young: 75% were between 18 and 25 years old (see Table A6 in the Appendix). This is also expected given both the populations of emoji users [55] and of Twitter users [38].

To gain a broad understanding of the geography of our respondents, we followed prior work and used the time zones attached to our respondents' Twitter accounts as a low-fidelity location indicator (e.g., [22,29,36]). From this data, it is clear that, as expected from our filters (see above), the vast majority of our respondents come from English-speaking countries. We observed that the plurality of our respondents had US/Canada time zones (e.g. "Pacific Time (US & Canada)"), and the most prominent non-US/Canada time zone was "London."

Twitter provides some indication of the "source" of each tweet in its API responses, where source is defined as the "utility used to post the tweet." Table 1 shows the source breakdown for tweets in our survey: 59.4% of the tweets came from Twitter for iPhone, iPad or Mac, 37.9% from Twitter for Android and 2.7% from the Twitter Web Client. Emoji renderings on Android devices are fragmented by manufacturer, but the source data given by Twitter does not capture manufacturer data. To gauge which vendors were represented in the devices respondents used to take our survey, we showed each respondent an emoji rendered natively (using the respondent's device's emoji font) and asked the respondent to choose the emoji rendering seen from a list of the emoji's renderings. From the answers to this question (Table 2), we estimate that the Twitter for Android devices are mostly split between Google's (Android) and Samsung's emoji renderings, with a very small percentage using either LG's renderings or HTC's renderings.

3.3 Emoji Rendering Software

Before describing the results of our survey, we first describe the software we built to simulate rendering emoji-bearing messages on different platforms. This software was used to implement the central component of our survey: asking participants whether they would edit their tweet after they could see how it appeared for followers using different platforms.

As explained above, emoji are rendered by vendor-version specific fonts. Even though Twitter has its own emoji font—Twemoji—Twitter applications largely render emoji natively using the device's emoji font (aside from the exception mentioned in Section 2.1). Twemoji are used in the web client, i.e., when viewing Twitter in a browser. Thus, since emoji-bearing tweets are often viewed on a wide variety of platforms, they are also viewed with a wide variety of renderings. When developing our emoji rendering software, we limited our renderings to vendors and their associated versions that were likely to be active on Twitter at the time of our survey (see Table A7 in the Appendix).

The approach we take in our emoji rendering software is largely straightforward but effective: our software parses out emoji characters in emoji-bearing input text and then outputs a list of HTML snippets that show how the message would render on each platform (in our sample of those that are active on Twitter). Each HTML snippet includes the original text, but with the emoji character(s) replaced by emoji rendering graphic(s) (hard-coded) to show how the emoji would render on the given platform.

To implement this approach, we first populated a database of emoji characters, vendors, vendor-versions, and renderings. We did this using a combination of data from the Unicode Technical Standard (UTS) for emoji [11,56] and from Emojipedia [14,48]. To render an emoji-bearing tweet across platforms, we first used an emoji regular expression [10] to parse the emoji from the text. Then, for each vendor-version, we replaced the emoji character(s) with that vendor-version's rendering(s) and output this information as HTML.

A significant challenge in executing the above otherwise-straightforward approach centered around a particularly important and moderately-common edge case: not all vendors' versions support every emoji character, meaning that a vendor-version does not always have a rendering for a given character. This is especially an issue for older versions that do not have renderings for newer characters, but also frequently occurs when platforms implement recently-released characters at different times (the Unicode Consortium adds new characters on an annual basis). In the interest of ecological validity, when a vendor-version does

Table 3. Examples of Emoji Code Points

Emoji	**Code Points and Constituent Emoji**
Beaming Face with Smiling Eyes	U+1F601
Clapping Hands: Medium-Dark Skin Tone	U+1F44F U+1F3FE [Clapping Hands] [Medium-Dark Skin Tone]
United States	U+1F1FA U+1F1F8 [Regional Indicator Symbol Letter U] [Regional Indicator Symbol Letter S]
Family: Man, Woman, Girl, Boy	U+1F468 U+200D U+1F469 U+200D U+1F467 U+200D U+1F466 [Man] [ZWJ*] [Woman] [ZWJ*] [Girl] [ZWJ*] [Boy]
Blond-Haired Woman	U+1F471 U+200D U+2640 U+FE0F [Person: Blond Hair] [ZWJ*] [Female Sign] [Variation Selector-16**]

* The Zero-Width Joiner (ZWJ) character indicates that surrounding characters should be joined into a single glyph if available [52].

** The Variation Selector-16 character, also known as the emoji variation selector, indicates that the preceding character be presented as an emoji (for characters that can also be presented as text, e.g., the Female Sign).

not have a rendering for a given character, our software carefully adheres to the exact rules of the UTS [56]. In some cases, this means rendering an "unsupported character" (□). However, in other cases, the behavior specified in the UTS is more complex. In particular, some emoji characters are encoded in the Unicode standard by multiple code points, including skin-tone modified emoji, flags, families and gendered emoji (e.g., see Table 3). According to the UTS, if an unsupported emoji character is composed of multiple code points, then its component code points should be rendered individually in sequence [56] (e.g., a family might be rendered as a string of its constituent members). We implemented this approach in our rendering software. For example, refer to Fig. 3 to see how our software rendered the emoji in Table 3.

3.4 Descriptive Statistics Regarding Emoji in Study

Each participant was shown one of her/his tweets in the survey, so we had a total of 710 tweets in our study. Of these tweets, 451 contained a single unique emoji character (either once or repeated), and 259 contained at least two different emoji characters. Across all 710 tweets, there were 1,488 total appearances of 583 unique emoji characters. Using Emojipedia's broad emoji categories [48], 164 emoji in the study were "smileys & people," 86 were "animals & nature," 34 were "food & drink," 24 were "activity," 29 were "travel & places," 47 were "objects," 63 were "symbols," 20 were "flags," and 116 were not categorized, 109 of which were skin-tone modified emoji.

As we would expect, emoji appearances in our sample followed a rough power law distribution: most emoji appeared in one (n=322) or two (n=121) tweets, with only 10 appearing in 10 or more tweets. See our supplementary material for the complete table of emoji in our sample, or Table A8 in Appendix A for a subset of this table containing the emoji that appeared most frequently in our sample.

Not surprisingly, the most popular emoji in our sample also are among the most popular emoji in general. Overall, though there are 2,666 total emoji characters, we estimate that the 583 emoji in our sample account for approximately 89% of all emoji appearances in tweets. This estimate is based on the distribution of emoji usage in the random sample of emoji-bearing tweets that we collected via the Twitter Streaming API as described above.

4 RESULTS

In this section, we present the results from our survey, which consisted of closed-form, structured questions as well as optional, open-ended questions that inquired as to participants' reasoning behind their closed-form responses. (See our supplementary material for an example path through the survey.) We primarily report descriptive quantitative statistics emerging from our structured questions. We additionally share insights from reading participants' open-ended responses to shed light on possible explanations behind our quantitative results.

4.1 RQ1: Awareness

We assessed participants' prior awareness of the multi-rendering nature of emoji with two questions. First, we showed the participant's tweet (rendered natively or with Twemoji, if the tweet was sourced from the Twitter web client) and asked, "Do you think that this tweet will appear exactly this way to everyone who views it?" Our intention was to assess natural recall, i.e. whether participants were already aware of emoji's multi-rendering nature and had it at the top of their minds when engaging with emoji-bearing messages. We found that 47% of participants (n=334) chose "Yes," they thought the tweet would appear exactly the same way to everyone who views it. This means that 47% of participants were either unaware that emoji look different on different platforms or did not recall this fact in the context of their emoji-bearing message.

In the second question, we showed the participant's tweet and asked more explicitly, "Did you know that the emoji in your tweet will appear differently to other users on Twitter? For example, your tweet will appear as the following on the associated devices / operating systems:" and then showed the renderings of their tweet across platforms (as in Fig. 3). For this question, 25% of participants (n=178) chose "No, I did not know this." The difference between this 25% number and the 47% number from the first question can likely be interpreted as, at least in part, a manifestation of the expected effects of recognition versus recall; once

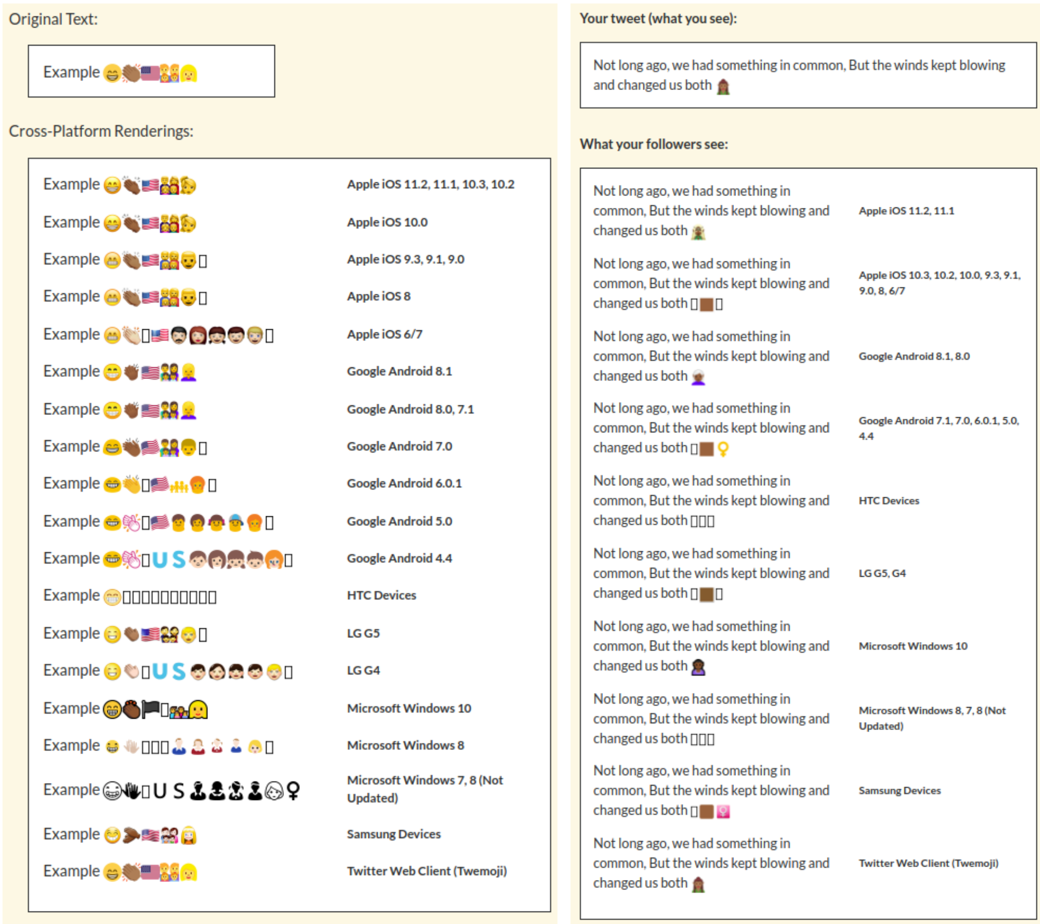

Fig. 3. Rendering tweets across platforms. The figure on the left shows the emoji from Table 3 rendered across platforms by our emoji rendering software. The figure on the right shows the view participants saw in the survey of their tweet rendered across platforms.

prompted, some participants likely had an "oh yeah!" reaction. Some portion of this difference may also be due to observer-expectancy effects, which would not manifest in the wild.

Putting these results together, the 25% result from our second question can be considered a lower-bound estimate of the percentage of emoji-using Twitter users that are not aware of the multi-rendering nature of emoji. The 47% result from our first question can be considered an upper-bound estimate. Regardless of the precise true value, it is clear that *a significant proportion of Twitter users that communicate using emoji are unaware that their emoji-bearing messages likely look different to many of their followers*. To put this into context, with over 300 million active Twitter users, this group of people likely contains tens of millions of people, and this does not account for the much larger group of people who use emoji on platforms other than Twitter.

For those in the 25% group, we provided a page in the survey that explained the multi-rendering nature of emoji and asked this group of participants about their reaction to learning this information. Looking at participant responses, some people found it moderately "interesting" or indicated that they were "surprised." The following participant quotes reflect these sentiments:

> "Interesting. Didn't know how many different ways it was viewed"

"I am mildly surprised there are so many types!"

Additionally, some participants were more than just surprised and expressed shock and/or worry:

"I feel completely blindsided. And amazed too. I feel that it is extremely important to be aware of this because we use this platform to communicate, and if the emojis that we use are not expressed in the manner we thought it would, that might lead to misinterpretation of our statements"

"Well, I was pretty bothered. What if some people misunderstood what I tweeted or posted because of the different renders of emojis? OMG "[5]

Likewise, we also see some clearly negative responses including "*That's annoying*," "*Not happy*," "*Kinda sucks*" and "*disappointed*." Finally, the opposite end of the spectrum is also represented in participant responses. Some found it "unsurprising" or were "indifferent." Examples of these sentiments include:

"Not very surprised but this is helpful"

"Almost indifferent, I'm sorry for who doesn't have an iPhone"

"interesting, but it is what it is; indifferent"

For those that were previously aware of the multi-rendering nature of emoji, we were curious about how they learned about it. Given the options of "Personal observation," "Someone else told you," "You read about it (e.g., in an article)," and "Other (fill in)," the vast majority (472 out of 532 aware participants) indicated that they became aware of the multi-rendering nature of emoji via personal observation. Among the other options, 16 participants (3%) became aware from someone else telling them and 27 participants (5%) became aware from reading about the multi-rendering nature of emoji. From examining participant open-ended explanations, the personal observation path consists of people who use multiple devices, who have had different devices in the past, who have seen or compared with friends' devices, and who have made inferences from seeing unsupported characters. These latter two cases are likely associated with instances of misconstrual.

In summary, with respect to RQ1, we found that at least a quarter of our participants were not aware that emoji appear (i.e., render) differently on different devices. Upon learning this information, some participants were surprised, shocked, and/or worried. For those that knew about this property of emoji, overwhelmingly it was due to personal observation.

4.2 RQ2: Effect on Communication Choices

After capturing our data related to prior awareness, we again showed participants the renderings of their tweet across platforms, now to capture whether this would have any effect on participants' communication behavior. To first get a broad sense of the potential effect of seeing emoji rendering differences across platforms, we asked participants, "Do you think your followers' versions of the tweet convey the same message you intended to send with your tweet?" Participants could choose between the following options:

- **Yes**, I think my followers' versions convey the same message.
- I think **some** of my followers' versions convey the same message, some do not.
- **No**, I think my followers' versions do not convey the same message.

Overall, the majority (60%) of participants reported that all of the tweet renderings conveyed the same message, but a large minority 38% felt that some of them did not and 2% felt that all of the tweet renderings did not convey the same message. Some of the open responses we received from people who were among

[5] This participant was using a Samsung phone, so this is Samsung's rendering.

the 40% of participants for whom cross-platform emoji rendering affected the perceived meaning of their message include:

> "The emojis below are mad. Mine was meant as irritated. If I wanted it to be mad, I would've put "" or something sarcastic like "". I hope the one reading this is using a Samsung to see my point."[6]

> "Some of them are really ugly. My message is "I'm kinda pissed and mad at nothing -_- so imma just sit here stone face." The ones that don't show the defined features of the stone face (hooded eyes, eyebrows, nose, flat lips) simply does not convey MY message and possibly paints another image."

Our survey next moved from interpretation to directly asking about participants' communication behavior. Specifically, following the above question, we asked: "If you had known that this is how your tweet would look to your audience, would you have sent it as-is?" Fifty-nine participants (8%) selected "No."

When asked, "How would you edit your tweet knowing this is how it looks to your audience?" participants responded as reported in Table 5. Table 5 shows that, knowing how the tweet would look across platforms, 18% of respondents (128) would have preferred to edit their tweet. These participants were relatively evenly split between choosing that they would edit the text, add more emoji, replace the emoji with another, and remove the emoji altogether.

Adding together the results for participants who indicated that they would edit the tweet with those that would not have sent the tweet, 20% of tweets (144) would have been edited or not sent had the authors seen how it would look across platforms. Note that a small portion (n=16) of the participants who said they would not send their tweet as-is also selected that they would not edit their tweet. For some of these participants, multi-rendering issues likely caused their tweet to be beyond repair (other causes for this could include confusion about the question or not seeing the "Other" option).

Because we had to use the Twitter Search API instead of Streaming API (see Twitter Recruitment section above), the distribution of emoji in our sample may differ somewhat from that of the population of emoji-bearing tweets. Although we observed that the most popular emoji on Twitter are also very common in our sample (see above), we wanted to formally correct for any sample-population discrepancies on this front by performing a stratified analysis of our data. Stratification is a method for adjusting the sampling weights of groups or *strata* within the data to account for these types of potential biases [39]. In other words, performing this analysis results in an estimate that more accurately reflects what one would expect from a true random sample of the population (with respect to emoji characters, in our case).

We stratified our data by the unique combinations of emoji contained in tweets in our sample (rather than simply by each unique emoji), because some tweets contained more than one emoji character and strata cannot overlap. We estimated the relative population weight of each emoji combination by searching for tweets containing each emoji combination (using the Twitter Search API) and computing the relative popularity of that emoji combination from all the tweets searched across all of our emoji combinations. After performing the stratification with these parameters, our proportion estimate went from 20% (of emoji-

Table 4. Would Edit Tweet Responses

How would you edit your tweet knowing this is how it looks to your audience?	N	%
I would not edit my tweet.	582	82.0%
I would edit the text in my tweet.	30	4.2%
I would add another emoji to my tweet.	32	4.5%
I would replace the emoji with another in my tweet.	32	4.5%
I would remove the emoji from my tweet.	29	4.1%
Other	5	0.7%

[6] We used Samsung's renderings in this quote.

bearing tweets *in our sample*) to 17.7% (of *all* emoji-bearing tweets). In other words, with this adjustment, we estimate that 17.7% of all emoji-bearing tweets would be edited or not sent as-is if the authors could have seen the emoji rendering differences across platforms.

These findings indicate that *emoji rendering differences across platforms represent a truly substantial problem for computer-mediated text communication.* To put this 17.7% figure into more context, we collected an unfiltered sample of over 7.7 million tweets from the Streaming API in July 2018, finding that approximately 16.3% of all tweets contain emoji. Since there are approximately 500 million tweets shared per day [57] and approximately 16.3% contain emoji, our 17.7% estimate suggests that *there are over 14 million tweets shared per day that would not have been sent or would have been edited if the authors could have seen the emoji rendering differences across platforms.* Even if our estimate only strictly applies to our very narrowly filtered population representing just 0.6% of all tweets[7], we would still estimate that there are over 530,000 such potentially regretful tweets shared per day. However, since emoji also are used in the rest of the tweet population (i.e., those in other languages than English, those with media, etc.), we expect that our observed effect applies more broadly than our highly filtered context. Of course, all of the above also only applies to emoji use in Twitter, which represents a small fraction of overall emoji use.

4.3 Factors Associated with Communication Choices

One question that emerges from our survey results is why some people were concerned about how their message rendered across platforms while others were not. We hypothesized that some factors behind this variation may include (1) characteristics of the emoji contained in the message, (2) [lack of] platform support for the emoji, and (3) the role that the emoji plays in the message and the overall purpose of the message. We also hypothesized that (4) prior awareness of the multi-rendering nature of emoji may have affected participants' communication choices. To explore these hypotheses, we performed straightforward univariate statistical tests. We also examined participant open-ended responses for evidence of whether any of these factors affected their choices.

4.3.1 Emoji-Related Factors

Emoji characters range from faces and gestures to basic objects, e.g., a trophy, a basketball, a plane. Since facial expressions are nuanced and complex [23] whereas visual object recognition is simpler, we hypothesized that "face" emoji would be more liable to meaningful changes across platforms than other types of emoji. Indeed, several participants supported this hypothesis in their open responses. With respect to object recognition, one participant wrote "*It's just a trophy*" and another wrote "*the emojis are still a train and a basketball.*" On the other hand, with respect to facial emoji, one participant reported "*I might be aware when I'm using other smiley emojis because some of them look really ugly in other devices…*"

The Unicode Consortium provides an emoji categorization [50] that includes several "face" categories (e.g., "face-positive," "cat-face," etc.). Using this categorization, we determined which tweets in our study contained face emoji. We observed that 24.0% of tweets that contained face emoji would be edited or not sent compared to 17.6% of tweets that did not contain any face emoji. Though not significant at the $p < 0.05$ level ($\chi^2(1,N=710) = 3.83$, $p = 0.05$), the results marginally suggest the expected trend: participants were more likely to indicate that they would edit or not send their tweet if it contained a face emoji.

Another emoji-related factor we hypothesized might be playing a role in our results is that, as was found by Miller et al. [33], some specific emoji characters are more likely to cause cross-platform problems (e.g., "beaming face with smiling eyes,"[8] U+1F601). Accordingly, we hypothesized that these more "ambiguous" [33] emoji would be of more concern for our participants. To investigate, we used Miller et al.'s [33] data for the sentiment ambiguity of 22 emoji characters. We reduced our data to tweets that only contained one of these 22 emoji characters (N=33), and we associated each tweet with its emoji's sentiment ambiguity. We

[7] Recall that our sample contains only English, original (not retweeted), media-less, Twitter-sourced (i.e., people-written, non-spam or machine-generated) tweets containing a subset of all possible emoji. We observed that 1.76% of tweets satisfy these filter criteria, and 38.7% of these filtered tweets bear emoji (see Section 3.1). Our results reflect 89% of all emoji (see Section 3.4), so, strictly speaking, altogether our sample represents 0.6% of the tweet population.

[8] Note: The Unicode has changed the name of this emoji since the time of the work of Miller et al. [33]

observed that the ambiguity of contained emoji was greater for tweets that would be edited or not sent (median ambiguity score[9] = 2.22) than for those that would not be edited (median ambiguity score = 1.84), but this difference was not significant (W(N=33) = 31.5, p = 0.10)[10]. However, given our limited sample size for this test, this may be a viable hypothesis to examine in a larger study in the future.

4.3.2 Platform-Related Factors

We hypothesized that "component" (e.g., the family emoji rendered as the man, woman, girl, and boy emoji individually) or "unsupported" (e.g., □) versions of emoji characters may play a significant role in whether or not participants indicated that they would edit or not send a tweet. Indeed, this issue has been cited as the primary potential benefit in Twitter's recent decision to use Twemoji for older Android OSes [8], and participants explicitly mentioned that this issue was problematic, e.g.,

"Some people wouldn't even see the emoji. Just empty boxes."

"If I'm texting someone and include an emoji that they can't see, the message may be taken a different way. I usually use emojis to lighten up a message and make it a little less serious so if they can't see it, it might change the way they read the text"

However, we did not find support for this hypothesis that component or unsupported versions of the emoji in a tweet would influence the decision to edit or not send it. We tracked whether participants were shown "component" or "unsupported" versions of the emoji in their tweet. Though these versions were more common for those that would edit or not send (80.3% vs. 78.6% for those that would not edit), this effect was not significant (χ^2(1,N=710) = 0.08, p = 0.78).

Somewhat unexpectedly, we saw two additional platform-related factors in participants' open-ended responses. First, several participants indicated they have a degree of "platform pride," meaning that they mentioned only really caring about the platform they use, e.g.,

"iPhone emojis are the best emojis. Everything else is just an ugly ripoff."

"Everybody knows it's the iphone emojis that are most popular and usually automatically interpret it as such."

Second, although this is very likely inaccurate, some participants felt that most of the people who would see their tweets use the same platform, e.g., *"Everyone has iphones"* and *"Because the majority of people I know have iPhones and iOS."*

4.3.3 Context-Related Factors

The specific role an emoji character plays in a tweet is also likely to be an important factor in the decision to edit or not send the tweet. For example, is the emoji used to add new meaning to the text? To reinforce the meaning of the text? To replace text? To change the tone of the text?

We found evidence of diverse emoji roles in participant responses:

"I choose emoji to supplement, rather than convey a message"

"I usually use emojis to lighten up a message and make it a little less serious"

"The emoji was only there to make it look a little attractive."

---

[9] This ambiguity score represents the average pairwise distance between people's sentiment interpretations of the emoji on a scale from -5 (Strongly Negative) to 5 (Strongly Positive), so the higher the ambiguity score, the higher the ambiguity [33].

[10] We use nonparametric Wilcoxon rank sum test [30] because the ambiguity measure from Miller et al. [33] does not follow a normal distribution.

For the purpose of this study, we captured a more general, self-reported measure of the role an emoji was playing in a tweet. Specifically, we asked participants to indicate the degree to which they agreed that their tweet needs the emoji to convey what they meant. We used a Likert scale from Strongly Disagree (-2) to Strongly Agree (2). With this information, we do not know the specific role the emoji was playing in the text, but we at least have a broad estimate of the importance of this role.

Using Wilcoxon rank sum test [30], we found that self-reported emoji importance was greater for tweets that would be edited or not sent (mean = 0.87 on Likert scale) than for those that would not be sent (mean = 0.56 on Likert scale) (W(N=710) = 35,133; p = 0.01). As we would expect, the more important the participant believes the emoji is to the tweet, the more likely the participant was to prefer to edit or not send the tweet.

The general purpose of the overall tweet is also likely to be important for the decision of whether or not someone would edit or not send the tweet. Mainly, how critical is it that the message be understood correctly? Who is it intended for? What would be the implications if it were misunderstood? Participants also provided comments related to these considerations. For example,

> "Sending a tweet that's not addressed to anyone particularly is just a message you put out it's not that important to me if people on other platforms see it differently."

> "I just like to tweet what I want and feel at the time I don't really pay much attention as to what my followers would say or think when I tweet."

From these assertions, emoji rendering differences across platforms appear to be less of a concern for those that use Twitter as a low-stakes communication platform. To at least partially test this hypothesis, we examined whether there was a relationship between a participant's number of followers and whether they chose to not send or edit the tweet. However, using a Wilcoxon rank sum test [30], we did not detect a significant relationship.

4.3.4 Prior Awareness

Finally, we hypothesized that respondents would be more likely to be affected by seeing emoji rendering differences across platforms if they were not previously aware of the multi-rendering nature of emoji. Indeed, while 15.8% of respondents who were previously aware would have edited or not sent their tweet, this number is 32.6% for people who indicated they were not previously aware of rendering differences. We found this relationship to be significant (χ^2(1,N=710) = 22.48, p < .001).

5 DISCUSSION

The top-level result from our survey is that emoji's potential for miscommunication that has been identified in prior work is having demonstrable, real-world effects on people's communication. A large minority of our respondents were not aware that emoji render differently across platforms, and being informed of this incited worry and frustration for some of them. We also observed that 8% of tweets in our sample *would not have been sent* had the Twitter users known how those tweets would render on viewers' platforms. Similarly, 18% of the tweets in our study would have been edited if the sender had visibility into the various potential ways the tweet would render. Indeed, our results suggest that hundreds of thousands if not millions of such potentially regretful emoji-bearing tweets are shared per day because the authors cannot see the emoji rendering differences across platforms.

Additionally, this estimate represents a very conservative lower-bound for the real-world effect of people not being able to see emoji rendering differences across platforms. Twitter is just one of many applications that support emoji communication across platforms; others include Instagram (nearly half of all text on Instagram bear emoji [17]), Slack (8 million active daily users across 500,000 organizations [37]), chat applications like Google Hangouts, and SMS text messaging (22 billion messages are sent every day worldwide [9]). Indeed, given the increasing prevalence of emoji-bearing communication, it is not unreasonable to expect that the effect observed in this study applies to a non-trivial percentage of all computer-mediated text communication.

5.1 The Need for Cross-Platform Emoji Communication Tools

A straightforward solution to the problems we identified would be for all platforms to converge on emoji rendering designs or to standardize to one emoji font. Though individual vendors may be able to take steps towards convergence (e.g., by getting users to update their devices or developing single-emoji-font communication applications that are used across platforms), this suggestion for complete convergence across vendors is unfortunately at odds with the twin forces of intellectual property law and branding incentives [20]. Thus, this is likely not a tractable solution for the foreseeable future.

As an alternative, a clear implication of our results is the need for new technologies to assist people with emoji communication in cross-platform environments. These technologies will likely all have the same core need as we had with our survey: to be able to simulate how an emoji-bearing message looks on a variety of platforms. As such, in order to facilitate the development of these technologies, we are releasing the source code for our rendering software.[11]

We can imagine many different instantiations of tools that use our rendering software to help users understand how their messages will appear to recipients. For instance, one could develop a Slack plugin that implements emoji previews, a third-party Twitter application that offers a similar functionality, or a web browser extension that surfaces the output of our rendering software for Gmail users who wish to include emoji in their emails.

These types of tools could also help to continue the trajectory of research related to emoji rendering differences across platforms begun by Miller et al. in 2016 [33]. Specifically, by logging the behavior of users of these tools, one could observe in an even more ecologically-valid fashion than our survey how users interact with the multi-rendering nature of emoji. While our survey asked people to reflect on their own real messages, this log data would allow researchers to observe this reflection in-situ. Indeed, building one of these tools is a subject of immediate future work for our team.

Relatedly, our results regarding factors associated with editing or not sending an emoji-bearing message suggest means by which future cross-platform emoji communication tools can be made more intelligent. For instance, one could imagine our hypothetical Slack plug-in from above popping up a warning message when a Slack message that is about to be sent is particularly vulnerable to cross-platform issues, but staying silent by default in other cases.

We identified several factors that may be relevant to this prediction task, e.g., whether or not an emoji character contained paralinguistic cues, the ambiguity of the emoji character, and certain contextual properties of the emoji-bearing message. However, we do not know how these factors interact, and some factors could mediate the others. As such, implementing a feature that predicts whether a given emoji-bearing message is problematic will likely require the training of a model to understand patterns in a relatively complex decision space. To do so, much more data than was provided by our survey will be necessary. The log data recorded by one of the suggested tools above (or a similar tool) could likely take significant steps towards accomplishing this goal, if deployed to enough people.

An additional obstacle in developing such intelligent tools is that some of the relevant factors are challenging to capture and/or quantify at scale, e.g., the ambiguity of an emoji character or the importance of an emoji to a message. More work is necessary to develop more robust and scalable measures for these factors, though some relevant work is under way. Several research efforts have contributed to understanding the different possible functions of emoji in text [15,24], and some have started trying to detect such functions automatically [34]. Our results suggest that these lines of work will be useful in predicting when emoji rendering differences across platforms will have an effect.

One question is whether the family of tools suggested above would be of interest to the group of people who do not seem concerned about emoji rendering differences across platforms, e.g., survey respondents who previously knew about the multi-rendering nature of emoji but are still choosing to communicate with emoji (as evidenced by their tweet). The data from our survey indicates that some of these people indeed do not care about miscommunication risks because of, e.g., "platform pride." However, it is much more likely

[11] https://github.umn.edu/emojilens/emoji-rendering-simulation

that these people do perceive risk, but they have decided that the risk does not outweigh the benefit of the ability to communicate with emoji. Some participants alluded to this tradeoff explicitly:

> "Your tweeps don't necessarily have the same phone as you so you try to choose the emoji that convey your thoughts or feelings as best you can with the choices you have"

> "I cannot predict exactly what each emoji will look like on each device so just have to keep using ones relevant to my device."

This perspective suggests that tools that surface emoji rendering differences across platforms would also be of significant utility to those that do not on the surface seem to be concerned about these differences: such tools would enable the weighing of risk versus benefit on a per-use basis. Further, these tools would make this easy compared to the next best existing alternative: looking up each emoji character one wants to use on a website like Emojipedia. For instance, one participant wrote, "*Sometimes I wonder how it would appear on other devices, but it's too much of a hassle to check all the time so I just roll with it.*"

Similarly, our survey results suggest that such tools would also be useful in cases in which emoji rendering differences are perceived to be of limited risk. Many of our participants indicated that they would send the tweet as-is after seeing the different renderings across platforms. Though some of these participants belong to the group of people that generally aren't concerned, it is likely that at least some of these participants are concerned but just not about the single tweet we showed them. This suggests that tools surfacing emoji renderings can provide a useful service regardless of the perceived risk of the rendering differences. If the differences are perceived to be risky, one can take appropriate action to edit or not send; if the differences are not perceived to be risky, one can take comfort in the decision to send as-is.

5.2 Platform Awareness Tools

Our results also highlight the need for new tools to encourage what one might call *platform awareness*. Some participants assumed that emoji rendering differences across platforms were not applicable to them because they believed everyone in their audience used devices from the same vendor. This perception was likely incorrect in almost every case (and substantially so). However, there is currently no way for a person to assess the platform distribution in their audience short of contacting each potential recipient to ask which platform he or she is using (not to mention which version). A similar risk exists (but in the reverse direction) in believing that all platform-versions are represented in a given audience. Naturally, this is especially the case when considering direct communication (e.g., SMS, direct messages) rather than broadcasting (e.g., standard Twitter).

Unfortunately, implementing client-side accurate platform detection is a nontrivial technical challenge. In the case of Twitter, Twitter has all the required information internally, but the company does not make this information available through its APIs (even the "source" information discussed above is far from sufficient as it does not provide information about all platforms nor versions). Outside of Twitter, the challenge becomes even more difficult. One way to infer platforms or devices is to analyze the User-Agent string from a browser (HTTP) request. However, this is difficult without a service specializing in these techniques (with an expansive database of learned User-Agent strings) like Device Atlas [13,16]. Also, using this approach would require all audience members to make a browser request of some sort from all of their devices.

One intriguing possibility is to scale the approach we took for our survey and use emoji themselves to disambiguate platforms. A given emoji rendering reflects a vendor-version combination. Thus, by rendering an emoji natively and then asking which from a list of renderings is the emoji being shown, the user can implicitly select the vendor and version being used to view the emoji. However, this would only be practical for certain types of applications and it would be necessary to determine and maintain a list of maximally informative emoji renderings.

5.3 Future Work and Limitations

This research produced the first empirical information on (1) general awareness of the multi-rendering nature of emoji and (2) the behavior of users in response to exposure of emoji rendering differences across platforms. However, because this is early research in this space, our work also raised a number of key follow-on questions and has a few limitations that represent important directions of future work.

First, the aforementioned recent update by Twitter that renders Twemoji in place of older Android emoji fonts moderately reduces the forward-looking ecological validity of our study because it essentially deactivates the emoji fonts of some of the vendor-version configurations that were included in our survey. However, this does not apply to the many other enormously popular applications that support emoji communication across platforms (e.g., e-mail, SMS, Instagram), and all other cross-platform contexts on Twitter will remain (including future versions). Also, although this update is primarily being discussed as a means to reduce unsupported versions of emoji being seen [8], we did not detect evidence that seeing unsupported versions of emoji affected the decision of whether or not to edit a message. That is, we do not have evidence to suggest that our results would differ if fewer unsupported versions were shown.

Second, this research targeted only English-speaking Twitter users (and even excluded retweets and tweets with media). However, as noted above, English-speaking Twitter users represent only a small minority of the overall emoji-using population. People who speak many different languages and communicate with many different platforms (e.g., Instagram, SMS, WhatsApp) include emoji in their messages. The results of our study strongly call for similar studies to be done with these other groups of emoji users.

Relatedly, while our survey provided a broad understanding of the real-world awareness and effects of (not) seeing emoji rendering differences across platforms, other methodological approaches could shed more light on these phenomena. In particular, our results call for in-depth qualitative interview work with a limited number of participants to identify themes in what might be causing the results we observed above. The open-ended responses in our survey highlight preliminary possibilities, but our initial themes should be verified and explored more rigorously. Similarly, as discussed above, while our survey provided increased ecological validity by having users consider their own tweets (instead of, e.g., hypothetical sample tweets), assessing user behavior in the context of sending a specific message would increase this validity (e.g., via log data from a tool).

Because of limitations in platform awareness technology, our survey was not able to inform respondents about the platforms on which their tweet was actually viewed. Instead, respondents saw their tweets on the platforms that are active across all of Twitter. As noted above, only Twitter has the information necessary to determine specific platform viewing information for every tweet. If platform awareness technology improves, it could be useful to replicate our work with actual per-respondent platform distribution information.

Lastly, as noted by Kariryaa et al. [25], Twitter's ad placement algorithm is rather opaque and prevents us from gaining a good understanding of the representativeness of our sample. While our geographic information matched expectations and we collected some demographic information ourselves that matched rough expectations for emoji use, it is possible that our sample is skewed in ways that we did not detect.

6 CONCLUSION

In this paper, we advanced the line of research that seeks to understand the impact of the multi-rendering nature of emoji. Using a survey of over 700 people deployed via Twitter's ad platform, we found that a substantial proportion of emoji-using Twitter users (at least 25%) were not even aware that their emoji look different to people using different platforms. We also saw that 18% of Twitter users indicated that they would edit or not send a recent emoji-bearing tweet had they known how that tweet rendered on other platforms. We discussed the implications of our results for new cross-platform emoji tools and our understanding of cross-platform emoji communication. We also released our cross-platform emoji rendering software that we used in our survey, which should facilitate the development of these tools.

ACKNOWLEDGMENTS

We thank the Twitter users who participated in our study. We are also grateful to our anonymous reviewers for their comments and suggestions that led us to strengthen our paper. This work was supported in part by the National Science Foundation Graduate Research Fellowship under grant 00039202 and Northwestern University.

A Appendix

Table A5. Participant Gender

Gender	Twitter Impressions	Link Clicks	%	Participants	%
Female	808,861	4,347	63.6%	512	72.1%
Male	490,397	2,394	35.0%	182	25.6%
Other	17,202	97	1.4%	16	2.3%

Table A6. Participant Age

Age Group	N	%
18-25	534	75.2%
26-35	115	16.2%
36-45	32	4.5%
46-55	19	2.7%
56+	10	1.4%

Table A7. Platforms in Emoji Rendering Software

Vendor	Version	Release Date	Vendor	Version	Release Date
Apple	iOS 11.2	12/2/2017	Google	Android 8.1	12/5/2017
	iOS 11.1	10/31/2017		Android 8.0	8/21/2017
	iOS 10.3	3/27/2017		Android 7.1	10/20/2016
	iOS 10.2	12/12/2016		Android 7.0	8/22/2016
	iOS 10.0	9/13/2016		Android 6.0.1	12/7/2015
	iOS 9.3	3/21/2016		Android 5.0	11/3/2014
	iOS 9.1	10/27/2015		Android 4.4	10/31/2013
	iOS 9.0	9/9/2015	Microsoft	Windows 10	10/17/2017
	iOS 8.3	4/8/2015		Windows 8.1	10/17/2013
	iOS 6.0	9/19/2012		Windows 8.0	10/26/2012
HTC	Sense 8	4/12/2016	Samsung	Experience 8.5	9/15/2017
LG	G5	4/1/2016	Twitter	Twemoji 2.5	2/22/2018
	G4	5/18/2015			

Vendor-version configurations that we assessed were likely to be active on Twitter at the time of the survey.

Table A8. Most Frequent Emoji in Participants' Tweets

Code Point(s)	Emoji Name	Emojipedia Category	Appearances		
			Total	Tweet	Alone
U+1F495	Two Hearts	symbols	25	14	2
U+1F602	Face With Tears of Joy	people	22	15	1
U+2764 U+FE0F	Red Heart	symbols	20	17	0
U+1F60D	Smiling Face With Heart-Eyes	people	17	12	1
U+1F62D	Loudly Crying Face	people	15	10	1
U+1F49E	Revolving Hearts	symbols	15	10	1
U+1F496	Sparkling Heart	symbols	15	9	3
U+1F60A	Smiling Face With Smiling Eyes	people	13	11	1
U+1F497	Growing Heart	symbols	13	7	1
U+1F629	Weary Face	people	12	10	3
U+1F483 U+1F3FD	Woman Dancing: Medium Skin Tone	None	12	3	3
U+1F3B6	Musical Notes	symbols	11	11	5
U+1F631	Face Screaming in Fear	people	11	8	4
U+1F44F U+1F3FC	Clapping Hands: Medium-Light Skin Tone	None	11	6	1

Includes the 14 most frequently occurring emoji by total number of appearances in survey respondents' tweets. Appearance counts are given for total occurrences (including repetitions in a given tweet), the number of unique tweets in which the emoji appeared, and the number of unique tweets in which only that emoji appeared (alone).

REFERENCES

[1] Fathiya Al Rashdi. 2015. Forms and Functions of Emojis in Whatsapp Interaction among Omanis. Georgetown University, Washington, DC, USA.

[2] Francesco Barbieri, German Kruszewski, Francesco Ronzano, and Horacio Saggion. 2016. How Cosmopolitan Are Emojis?: Exploring Emojis Usage and Meaning over Different Languages with Distributional Semantics. In *Proceedings of the 2016 ACM on Multimedia Conference* (MM '16), 531–535. https://doi.org/10.1145/2964284.2967278

[3] Francesco Barbieri, Francesco Ronzano, and Horacio Saggion. 2016. What does this Emoji Mean? A Vector Space Skip-Gram Model for Twitter Emojis. In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC 2016).*

[4] Rachel Been and Agustin Fonts. 2017. Redesigning Android Emoji. *Google Design*. Retrieved October 14, 2017 from https://medium.com/google-design/redesigning-android-emoji-cb22e3b51cc6

[5] Bianca Bosker. 2014. How Emoji Get Lost In Translation. *Huffington Post*. Retrieved from https://www.huffingtonpost.com/2014/06/27/emoji-meaning_n_5530638.html

[6] Jeremy Burge. 2018. Apple's Emoji Crackdown. *Emojipedia*. Retrieved July 8, 2018 from https://blog.emojipedia.org/apples-emoji-crackdown/

[7] Jeremy Burge. 2018. Slack Overhauls Emoji Support With One Catch. *Emojipedia*. Retrieved April 19, 2018 from https://blog.emojipedia.org/slack-overhauls-emoji-support/

[8] Jeremy Burge. 2018. Twitter Switches to Twemoji on Android. *Emojipedia*. Retrieved September 1, 2018 from https://blog.emojipedia.org/twitter-switches-to-twemoji-on-android/

[9] Kenneth Burke. 73 Texting Statistics That Answer All Your Questions. *Text Request*. Retrieved July 8, 2018 from https://www.textrequest.com/blog/texting-statistics-answer-questions/

[10] Mathias Bynens. 2018. emoji-regex: A regular expression to match all Emoji-only symbols as per the Unicode Standard. Retrieved from https://github.com/mathiasbynens/emoji-regex

[11] Mathias Bynens. 2018. *unicode-tr51: Emoji data extracted from Unicode Technical Report #51*. Retrieved from https://github.com/mathiasbynens/unicode-tr51

[12] Zhenpeng Chen, Xuan Lu, Sheng Shen, Wei Ai, Xuanzhe Liu, and Qiaozhu Mei. 2017. Through a Gender Lens: An Empirical Study of Emoji Usage over Large-Scale Android Users. In *WWW '18 Proceedings of the 2018 World Wide Web Conference* (WWW '18), 763–772. https://doi.org/10.1145/3178876.3186157

[13] Martin Clancy. 2018. User Agent parsing: how it works and how it can be used. *DeviceAtlas*. Retrieved September 3, 2018 from https://deviceatlas.com/blog/user-agent-parsing-how-it-works-and-how-it-can-be-used

[14] Ben Congdon. 2018. *python-emojipedia: :sunglasses: Emoji data from Emojipedia :tada:* Retrieved from https://github.com/bcongdon/python-emojipedia

[15] Henriette Cramer, Paloma de Juan, and Joel Tetreault. 2016. Sender-intended Functions of Emojis in US Messaging. In *Proceedings of the 18th International Conference on Human-Computer Interaction with Mobile Devices and Services* (MobileHCI '16), 504–509. https://doi.org/10.1145/2935334.2935370

[16] DeviceAtlas. Device Detection. Retrieved July 11, 2018 from https://deviceatlas.com/device-detection

[17] Thomas Dimson. 2015. Emojineering Part 1: Machine Learning for Emoji Trends. *Instagram Engineering Blog*. Retrieved from http://instagram-engineering.tumblr.com/post/117889701472/emojineering-part-1-machine-learning-for-emoji

[18] Ben Eisner, Tim Rocktäschel, Isabelle Augenstein, Matko Bošnjak, and Sebastian Riedel. 2016. emoji2vec: Learning Emoji Representations from their Description. In *Proceedings of the 4th International Workshop on Natural Language Processing for Social Media at EMNLP 2016.*

[19] Motahhare Eslami, Aimee Rickman, Kristen Vaccaro, Amirhossein Aleyasen, Andy Vuong, Karrie Karahalios, Kevin Hamilton, and Christian Sandvig. 2015. "I Always Assumed That I Wasn'T Really That Close to [Her]": Reasoning About Invisible Algorithms in News Feeds. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems* (CHI '15), 153–162. https://doi.org/10.1145/2702123.2702556

[20] Eric Goldman. 2017. *Surveying the Law of Emojis*. Social Science Research Network, Rochester, NY. Retrieved July 31, 2017 from https://papers.ssrn.com/abstract=2961060

[21] Fredrik Hallsmar and Jonas Palm. 2016. Multi-class Sentiment Classification on Twitter using an Emoji Training Heuristic. KTH Royal Institute of Technology.

[22] Bo Han, Afshin Rahimi, Leon Derczynski, and Timothy Baldwin. 2016. Twitter Geolocation Prediction Shared Task of the 2016 Workshop on Noisy User-generated Text. *Proceedings of the 2nd Workshop on Noisy User-generated Text (WNUT)*: 213–217.

[23] Ran R. Hassin, Hillel Aviezer, and Shlomo Bentin. 2013. Inherently Ambiguous: Facial Expressions of Emotions, in Context. *Emotion Review* 5, 1: 60–65. https://doi.org/10.1177/1754073912451331

[24] Susan Herring and Ashley Dainas. 2017. "Nice Picture Comment!" Graphicons in Facebook Comment Threads. *Hawaii International Conference on System Sciences 2017 (HICSS-50).*

[25] Ankit Kariryaa, Isaac Johnson, Johannes Schöning, and Brent Hecht. 2018. Defining and Predicting the Localness of Volunteered Geographic Information using Ground Truth Data. In *Proceedings of the 36th Annual ACM Conference on Human Factors in Computing Systems* (CHI '18). https://doi.org/10.1145/3173574.3173839

[26] Ryan Kelly and Leon Watts. 2015. Characterising the inventive appropriation of emoji as relationally meaningful in mediated close personal relationships. *Experiences of Technology Appropriation: Unanticipated Users, Usage, Circumstances, and Design*.

[27] Kun-Lin Liu, Wu-Jun Li, and Minyi Guo. 2012. Emoticon Smoothed Language Models for Twitter Sentiment Analysis. In *AAAI'12 Proceedings of the 26th AAAI Conference on Artificial Intelligence*, 1678–1684.

[28] Xuan Lu, Wei Ai, Xuanzhe Liu, Qian Li, Ning Wang, Gang Huang, and Qiaozhu Mei. 2016. Learning from the Ubiquitous Language: An Empirical Analysis of Emoji Usage of Smartphone Users. In *Proceedings of the 2016 ACM International Joint Conference on Pervasive and Ubiquitous Computing* (UbiComp '16), 770–780. https://doi.org/10.1145/2971648.2971724

[29] Jalal Mahmud, Jeffrey Nichols, and Clemens Drews. 2014. Home Location Identification of Twitter Users. *ACM TIST* 5, 3: 1–21. https://doi.org/10.1145/2528548

[30] H. B. Mann and D. R. Whitney. 1947. On a Test of Whether one of Two Random Variables is Stochastically Larger than the Other. *The Annals of Mathematical Statistics* 18, 1: 50–60.

[31] Ben Medlock and Gretchen McCulloch. 2016. The Linguistic Secrets Found in Billions of Emoji. Retrieved October 24, 2016 from http://www.slideshare.net/SwiftKey/the-linguistic-secrets-found-in-billions-of-emoji-sxsw-2016-presentation-59956212

[32] Hannah J. Miller, Daniel Kluver, Jacob Thebault-Spieker, Loren G. Terveen, and Brent J. Hecht. 2017. Understanding Emoji Ambiguity in Context: The Role of Text in Emoji-Related Miscommunication. In *Proceedings of the 11th International AAAI Conference on Web and Social Media* (ICWSM '17), 152–161.

[33] Hannah J. Miller, Jacob Thebault-Spieker, Shuo Chang, Isaac Johnson, Loren Terveen, and Brent Hecht. 2016. "Blissfully Happy" or "Ready toFight": Varying Interpretations of Emoji. In *Proceedings of the 10th International AAAI Conference on Web and Social Media* (ICWSM '16), 259–268.

[34] Noa Na'aman, Hannah Provenza, and Orion Montoya. 2017. Varying Linguistic Purposes of Emoji in (Twitter) Context. *Proceedings of ACL 2017, Student Research Workshop*: 136–141.

[35] Petra Kralj Novak, Jasmina Smailović, Borut Sluban, and Igor Mozetič. 2015. Sentiment of emojis. *PloS one* 10, 12.

[36] Reid Priedhorsky. 2014. Inferring the Origin Locations of Tweets with Quantitative Confidence. *CSCW* 29. https://doi.org/10.1016/j.biotechadv.2011.08.021.Secreted

[37] Slack. About Us. *Slack*. Retrieved July 8, 2018 from https://slack.com/about

[38] Aaron Smith and Monica. 2018. Social Media Use in 2018. *Pew Research Center: Internet, Science & Tech*. Retrieved July 7, 2018 from http://www.pewinternet.org/2018/03/01/social-media-use-in-2018/

[39] T. M. F. Smith. 1991. Post-Stratification. *Journal of the Royal Statistical Society. Series D (The Statistician)* 40, 3: 315–323. https://doi.org/10.2307/2348284

[40] SwiftKey. 2015. Most-used emoji revealed: Americans love skulls, Brazilians love cats, the French love hearts. *SwiftKey Blog*. Retrieved from https://blog.swiftkey.com/americans-love-skulls-brazilians-love-cats-swiftkey-emoji-meanings-report/

[41] Garreth W. Tigwell and David R. Flatla. 2016. Oh That's What You Meant!: Reducing Emoji Misunderstanding. In *Proceedings of the 18th International Conference on Human-Computer Interaction with Mobile Devices and Services Adjunct* (MobileHCI '16), 859–866. https://doi.org/10.1145/2957265.2961844

[42] Joseph B. Walther and Kyle P. D'Addario. 2001. The impacts of emoticons on message interpretation in computer-mediated communication. *Social science computer review* 19, 3: 324–347.

[43] Sanjaya Wijeratne, Lakshika Balasuriya, Amit Sheth, and Derek Doran. 2016. EmojiNet: Building a Machine Readable Sense Inventory for Emoji. In *Social Informatics*, 527–541. https://doi.org/10.1007/978-3-319-47880-7_33

[44] Sanjaya Wijeratne, Lakshika Balasuriya, Amit Sheth, and Derek Doran. 2017. EmojiNet: An Open Service and API for Emoji Sense Discovery. In *Proceedings of the 11th International AAAI Conference on Web and Social Media* (ICWSM '17).

[45] 2015. #ChevyGoesEmoji. *media.gm.com*. Retrieved January 24, 2018 from https://media.gm.com/content/media/us/en/chevrolet/news.detail.html/content/Pages/news/us/en/2015/jun/0622-cruze-emoji.html

[46] Oxford's 2015 Word of the Year Is This Emoji. *Time*. Retrieved January 24, 2018 from http://time.com/4114886/oxford-word-of-the-year-2015-emoji/

[47] Unicode Consortium. Retrieved January 24, 2018 from http://unicode.org/consortium/consort.html

[48] Emojipedia. Retrieved January 14, 2017 from http://emojipedia.org/

[49] EmojiNet. Retrieved April 19, 2018 from http://emojinet.knoesis.org/home.php

[50] Full Emoji List, v11.0. Retrieved July 10, 2018 from https://unicode.org/emoji/charts/full-emoji-list.html

[51] 😍iEmoji.com - 👀Lookup, ✨Convert, and Get Emoji! 🔥. Retrieved September 3, 2018 from https://www.iemoji.com/

[52] Microsoft Windows 10 April 2018 Update Emoji List. Retrieved August 16, 2018 from https://emojipedia.org/microsoft/windows-10-april-2018-update/changed/

[53] 10 Heuristics for User Interface Design: Article by Jakob Nielsen. *Nielsen Norman Group*. Retrieved April 19, 2018 from https://www.nngroup.com/articles/ten-usability-heuristics/

[54] Global Twitter user distribution by gender 2018 | Statistic. *Statista*. Retrieved July 7, 2018 from https://www.statista.com/statistics/828092/distribution-of-users-on-twitter-worldwide-gender/

[55] Millennials Say GIFs, Emojis Communicate Their Thoughts Better Than Words. *Time*. Retrieved April 19, 2018 from http://time.com/4834112/millennials-gifs-emojis/

[56] UTS #51: Unicode Emoji. Retrieved January 24, 2018 from http://unicode.org/reports/tr51/

[57] Twitter Usage Statistics - Internet Live Stats. Retrieved July 11, 2018 from http://www.internetlivestats.com/twitter-statistics/

Received April 2018; revised July 2018; accepted September 2018