

How to Do Things without Words: Modeling Semantic Drift of Emoji

Eyal Arviv and Oren Tsur

Software and Information Systems Engineering

Ben Gurion University of the Negev

eyalar@post.bgu.ac.il orentsur@bgu.ac.il

Abstract

Emoji have become a significant part of our informal textual communication. Previous work, addressing the societal and linguistic functions of emoji, overlooked the relation between the semantics and the visual variations of the symbols. In this paper we model and analyze the semantic drift of emoji and discuss the features that may be contributing to the drift, some are unique to emoji and some are more general. Specifically, we explore the relations between graphical changes and semantic changes.

Warning: Some readers may find some of the words and emoji interpretations offensive.

1 Introduction

Emoji have become a significant part of our informal textual communication. One of the major linguistic functions served by emoji¹ is the re-introduction of physical-like gestures in textual messaging (Evans, 2017; McCulloch, 2019). While Emoji is not a ‘natural language’ in the typical sense, it is a linguistic phenomena that can be addressed from a (computational-) linguistic perspective. Moreover, the unique nature and inherent limitations of the use of emoji, provide a distinctive perspective for understanding various linguistic phenomena such as lexicization, gesturing, ambiguity and semantic drift. In this paper we focus on temporal distributional semantics of emoji. To the best of our knowledge, this is the first work to address semantic drift of emoji with respect to graphical variation across time and platform.

The Emoji Echo-system The Unicode Consortium started adding Emoji to the Unicode standard at 2010, assigning a code to a short description².

¹The term ‘emoji’ refers to the set of codified pictograph (the ‘language’), to a specific icon (singular), and to a set of icons (plural).

²The complete list of codes and descriptions can be found at: <https://unicode.org/emoji/charts/full-emoji-list.html>.

For example, U+1F602 is described as *face with tears of joy*, U+1F9D0 as *face with monocle*, and the U+1F36A code is simply defined as *cookie*. The graphic realization of the description is left for the discretion of the designers of each specific platform. For example, the graphic interpretation of the *cookie* code (U+1F36A) ranges from variations of single chocolate-chip cookie to a pair of crackers: 🍪, 🍪, and 🍪 (iOS 13.2, Twitter’s Twemoji 12.2, and Samsung Touchwiz 7.1, respectively).

The exact meaning of an emoji depends on the social and textual contexts in which it is used – the smiley emoji, for instance, can be used to show empathy, express amusement or to ease tension. The pistol emoji can be used to convey frustration or amusement (“you are killing me”), bragging (“I killed it”) or a proper threat (“I WILL kill you”) – a disturbing illocutionary act (Austin, 1975; Salgueiro, 2010), performed without words. Consequently, the accurate interpretation of an emoji may have significant legal implications, possibly landing a defendant accused of murder a life in jail, if the emoji is considered a threat – proving premeditation of the murder (Goldman, 2018). Similarly, if the semantics of the 🍆 have drifted from a vegetable to a sexual reference (which is, indeed, the case), sending this emoji to a colleague could be perceived as a sexual misconduct.

Semantic Drift Language keeps changing. Semantic drift, often referred to as semantic shift or semantic change, is the process in which words change their meaning over time (Hock and Joseph, 2009). Classic examples include ‘gay’, originally meaning ‘happy’ or ‘cheerful’, now used to describe a sexual orientation; and ‘awful’, originally meaning majestic, then used as ‘horrible’ before gaining the opposite meaning – ‘wonderful’ (Hamilton et al., 2016b).

Broadly speaking, a semantic shift is a grass-root process – a result of the way the speaker community use the language, rather than a shift forced by

a regulatory body. That is, an institution (e.g., a corporation) may develop an ‘anti-virus’ (referring to a computer virus, rather than an infectious agent that replicates within a host organism) but speakers may not adopt the new use of the word. Organized social groups can press for a change in the ways words should/not be used (e.g., negro, retard, or guys) but these efforts are not institutionalized. Standardization enforced by language regulatory bodies tend to fail, as famously illustrated by the evolution of the French negation (Jespersen, 1917; Dahl, 1979; Hock and Joseph, 2009).

Emoji vs. Written Text Emoji is unique in the way it combines four elements, each is governed by a different mechanism: (i) a subtle contextual meaning (like many proper words in the language), (ii) a strict “spelling” (the Unicode code), (iii) a formal description (e.g., “face with tears of joy”) decided by a regulatory body (the Unicode Consortium), and (iv) a flexible visual standard decided by any number of a third-party organizations. Together, these elements restrict the ways in which emoji can be creatively manipulated by the users. While users can write ‘coz’ instead of the formal dictionary entry ‘because’, or ‘on da fon wit ma ma’ instead of “on the phone with my mother”, they cannot be playful with the unicode of a specific smiley from the ‘emoji dictionary’. Creativity is typically achieved by emoji collocations and by allusions inspired by visual features. The eggplant emoji 🍆, for example, is commonly used as a sexual reference to the male genitalia (replacing the 🍌, previously used for the same purpose).

2 Related Work

Emoji Echo-system and Representation The emoji echo-system, its cultural and historical roots, its evolution and the socio-linguistic function of the use of emoji are surveyed in (Evans, 2017; McCulloch, 2019). Cultural differences in the use of emoji are studied by (Chen et al., 2018; Guntuku et al., 2019; Li et al., 2019). The ambiguous interpretation of emoji is studied in (Miller et al., 2016), and temporal variations in usage and popularity are studied by (Ai et al., 2017; Barbieri et al., 2018).

Recent works use neural models for learning emoji representation in a way useful for downstream tasks. A skip-gram model was used to analyze the use of emoji on Instagram (Dimson, 2015) and Twitter (Barbieri et al., 2016). Emoji2vec (Eisner et al., 2016) demonstrates that learning the em-

beddings based on the formal description of the unicode yields better representation than learning the embeddings using the emoji unicode as a textual token. DeepMoji (Felbo et al., 2017) uses millions of tweets for training a multi-layer (Bi-LSTM with attention layers) network in learning embeddings for 64 popular emoji. The quality of the representation is reported across tasks and benchmarks (sentiment analysis, emotion detection, sarcasm detection). A CBOW model was used in a comparative study of the use of emoji on Twitter and Sina-Weibo (Guntuku et al., 2019).

All of the works cited above assume a static meaning of the emoji, although ambiguity or cultural differences are acknowledged. Few works do address changes in emoji usage. Temporal variations in usage and popularity of emoji are addressed by (Ai et al., 2017; Barbieri et al., 2018). To the best of our knowledge, the only work to address semantic change of emoji from a linguistic perspective is (Robertson et al., 2021). No prior work study the semantic change emoji with regards to visual differences and graphical changes across time and platforms.

Semantic drift Semantic change is studied by linguists for decades, see (Hock and Joseph, 2009) for a survey. The availability of large textual corpora gave rise to the computational study of semantic change. Early work used changes in word frequency in the span of two centuries, inferring grammatical and semantic shifts (Michel et al., 2011) and the changing relations between syntax and meaning (Goldberg and Orwant, 2013). The change in gender and racial stereotypes through the 20th century was modeled by (Garg et al., 2018).

Other works use word embeddings to capture semantic change by measuring the mean-shift (Kulkarni et al., 2015) or the Spearman correlation between time series (Hamilton et al., 2016b). The intersection between the sets of k nearest neighbors of a word w in different time spans is demonstrated to provide stable and interpretable results (Gonen et al., 2020).

The linguistic factors and processes that facilitate the semantic change include word frequency, polysemy, grammatical category, subjectification, and grammaticalization (Algeo, 1977, 1980; Brinton and Traugott, 2005; Xu and Kemp, 2015; Dubossarsky et al., 2015; Hamilton et al., 2016a; Dubossarsky et al., 2017), and cultural phenomena, such as the emergence of new technologies (Hamil-







ton et al., 2016a).

3 Data







Data Collection Using the Twitter streaming API we collected 1% (daily) of all English tweets containing emoji between May 2016 and April 2019. In total, our data contained over 209 million tweets posted by 37 million unique users. Table 1 presents the total number of tweets collected, as well as the tweets per each of the three main platforms used to tweet. We provide the numbers for each of these platforms since the visual realization of emoji may change between platforms and since the monthly subset per platform may impact the models learnt for each month and platform. The total number of users per platform exceeds the total number of users since many users tweet from more than one platform. For example, 1.9 million users use both an iOS device and the web app; 1.7 million use both an Android device and the web app, and 367K users use all three platforms.

	Tweets	Users
ALL	209M	37M
iOS	113M	20M
Android	56M	13M
Web	33M	8.7M

Table 1: Counts of English tweets containing Emoji, collected in the span of 5/2016–4/2019. ALL counts refer to all platforms, not only the most popular three.

The longevity of the data covers multiple changes in the emoji visualizations across platforms, including some significant changes such as the pistol ( \mapsto ) and the cookie ( \mapsto ), among many other less pronounced changes, e.g., the ROFL emoji ( \mapsto ). These changes provide the unique opportunity to study the effect of the visual realization on the semantics of emoji, as well as drifts unrelated to a forced visual change³.

Data Preprocessing Retweets and tweets lacking text were removed from the corpus. The remaining tweets were lower-cased and punctuation was padded. Additionally, we removed URLs, special characters, and user mentions. Consecutive





³It is important to note that some visual differences more pronounced than others (consider , , , , , ), thus may not trigger a significant change in the way they perceived after the change. Quantifying the graphical change and its cognitive impact is a challenging task. We note, however, that semantic change could serve as a proxy to the impact of the visual change.

emoji were padded with white spaces⁴.

4 Experimental Setting

4.1 Learning Embedding Models

Embedding Models Given n data sets ($n = T \times P$; for T time-spans and P platforms), we train an embedding model⁵ M^i ($i = 1, \dots, n$) for each dataset. In total, we trained 36×3 models – a model for each month and platform in our data.

Sanity Checks Since our models are trained on n relatively small subsets of the data (e.g., an average of less than a million tweets per monthly model for the Android platform, see Table 1), we first verify that the models indeed learn proper embeddings. We do this by applying a set of standard analogy tests⁶. The test contained both word analogies (e.g., $\overrightarrow{king} - \overrightarrow{man} + \overrightarrow{woman} = \overrightarrow{queen}$) and emoji analogies (e.g.,  -  +  = ).




4.2 Quantifying Pair-wise Drift

Given an embedding model M^i and a vocabulary V (shared across datasets and derived models), the matrix D^i holds the pair-wise distances between all pairs of words (and emoji) in dataset i , so $D_{kl}^i = d(v_k, v_l)$ for a given distance function d , e.g., cosine-similarity, and the embeddings of v_k, v_l in M^i .

The baseline shift between two models M^i and M^j is given by the matrix $D^{i \rightarrow j} = D^i - D^j$ (all values are converted to their absolute value). The mean-shift and variance of the baseline shift, $\mu^{i \rightarrow j}$ and $\sigma^{i \rightarrow j}$, are computed over the values of $D^{i \rightarrow j}$.

The binary matrix $\Delta^{i \rightarrow j}$ indicates the *pair-wise* semantic shift for all pairs of words, where the shift for a pair $\{v_k, v_l\}$ is defined as:

$$\Delta_{kl}^{i \rightarrow j} = \begin{cases} 1 & |D_{kl}^{i \rightarrow j} - \mu^{i \rightarrow j}| > \beta \sigma^{i \rightarrow j} \\ 0 & \text{Otherwise} \end{cases}$$

⁴In some special cases a sequence of emoji unicodes with no white spaces but with the zero-width-joiner code U+200D serves as a “modifier”, e.g., the sequence U+1F469U+1F469U+1F466 (codes for woman,woman,boy) is displayed as , while the sequence U+1F469U+200DU+1F469U+200DU+1F466 is displayed as a single emoji – a family unit of a gay/lesbian couple and a boy , a modification of the U+1F46A (family) emoji . We keep these sequences unpadded, in order to maintain the modified emoji as a single emoji token.

⁵All reported results were obtained using the skip-gram model (Mikolov et al., 2013).

⁶To the degree analogy tests do learn inference relations, and see (Levy et al., 2015).

⁷There is a crown, prince and princess emoji but no emoji for king or queen, so we use the crown, woman, girl and princess emoji.

for some $\beta \geq 2$, establishing the statistical significance of the shift⁸.

While $\Delta_{kl}^{i \rightarrow j}$ indicates the pair-wise shift, we are interested in identifying which of the words in the pair underwent the shift. In order to do that we simply choose the word that changed with respect to more words. That is, for each k, l for which $\Delta_{kl}^{i \rightarrow j} = 1$, we decide whether it is k or l that shifted by computing:

$$\phi(k, l, \Delta^{i \rightarrow j}) = \begin{cases} k & \sum_{p \neq k} \Delta_{kp}^{i \rightarrow j} > \sum_{p \neq l} \Delta_{pl}^{i \rightarrow j} \\ l & \sum_{p \neq k} \Delta_{kp}^{i \rightarrow j} < \sum_{p \neq l} \Delta_{pl}^{i \rightarrow j} \end{cases}$$

Naturally, it is possible that both k and l have shifted, although $\phi(k, l, \Delta^{i \rightarrow j}) = k$. The shift of l will be acknowledged by $\phi(p, l, \Delta^{i \rightarrow j})$ returning l for some $p \neq k$, that is, the shift of a word is checked for every occurrence of it in $\Delta^{i \rightarrow j}$.

5 Results and Discussion

Using the method described in Section 4.2, we identify dozens of drifting emoji. Table 2 presents some illuminating examples, along with the semantic context before and after the change. In the remainder of this section we discuss some of these in more detail.










Token	Before Change	After Change
rally	politics, geo-locations	college, sports
Florence	tourism	hurricanes, disaster
boomer	terror, gaming	condescension, millennials
	anger, death, failure, gun rights	games, gaming, online trends
	drug abuse, weapons	medicine
	sweets, desserts	baking, food




Table 2: Examples of words and emoji that went through a semantic shift between 2016 and 2019.

Expected Drift Following (Kulkarni et al., 2015) example’s of Twitter drifts (*rally*: politics \mapsto gaming; *sandy*: beaches \mapsto hurricanes), we explicitly verified the existence of three assumed *word* drifts in our data: *rally*, *Florence*⁹, and the recent [*ok*] *boomer* trend. Indeed, we identified the drifts, as expected (see Table 2, Top).




The Right to Bear Arms:     The pistol emoji presents a unique case, as its visual changed significantly, from a realistic pistol to a toy ( \mapsto ). Indeed, this drift was identified across

⁸We use the Shapiro-Wilk test to verify the normality of the distribution of the values of $\Delta^{i \rightarrow j}$.

⁹Florence is the code of a devastating 2018 hurricane, similar to Sandy in Oct. 2012.

platforms, in the *respective time of the graphical change* – Apple changed the emoji in July 2016, while other platforms changed it during 2018. A careful examination of the semantic relations between the pistol and other words and emoji is illuminating. Before the visual change, the tokens most similar to the pistol emoji were *fml*¹⁰, *#shots-fired*,  (coffin),  (rage, pouting face),  (skull and cross-bones), and *2a* (short for the 2nd Amendment, used by gun rights advocates¹¹). Interestingly, we observe a growing semantic gap between the pistol and the former five emoji. We also observe a sharp drop in similarity between the pistol emoji and *2a* in the two months right after the visual change, followed by a slow but stable recovery of the similarity, although it does not reach the semantic similarity of the original emoji. This pattern of recovery deserves a more substantive cultural analysis.

Semantic Drift and Ambiguity Some of the emoji identified as drifting presented a non stable behaviour – drifting back and forth through the months without a clear semantic field. Regressing the time series of the change resulted in coefficient that is close to zero. We find that the semantic instability of these emoji correlates with the lack of cohesiveness as measured by the average distance of the emoji from the k ¹² terms (emoji and words) closest to it in the first month, and the average distance between each pair within that set of k terms. We leave further investigation to future work.

Seasonality and Global Events We have identified similar textual drifts as those reported by (Kulkarni et al., 2015), although our datasets span different time frames (2011–2013 vs. 2016–2019). These similarities suggest a pattern of cyclic changes, seasonal or related to global events. We observe similar cyclic changes for seasonal emoji like  (falling leaves),  (skier), or emoji related to world singular events like  (Portuguese flag) winners of the UEFA EURO 2016. These changes are reversed soon after the season or the event, would it be an election season (*rally*), a weather catastrophe (*Sandy*, *Florence*), or global sport event (Portugal, UEFA). While Kulkarni et al. acknowl-

¹⁰Acronym for ‘fuck my life’, commonly used to indicate an embarrassment, disappointment or anger.

¹¹The 2nd Amendment to the Constitution of the United States stipulating that “the right of the people to keep and bear Arms, shall not be infringed” is one of the most divisive issues in American politics.

¹²We set $k = 50$.

edge that their Twitter dataset spans a much shorter time, compared to the Google Book Ngram corpus, they do not account for the seasonal or temporal nature of some changes. We believe that these seasonal shifts, of words or emojis, cannot be coupled with the substantive cultural shifts.

6 Limitations and Future Work

There is a number of limitations to this work, some are inherent to short papers.

First, reporting on drift in a quantitative way (e.g., “we find n emoji drifting in our data” is hard and misleading due to the exploratory nature of the task and the fact that drift is a complex phenomenon – some drifts are cyclic and some are permanent; some are more pronounced and some are nuanced as the drift occurs on a specific sense of the emoji. We were trying to illustrate the various effects through qualitative examples. This approach is consistent with prior work, for example (Garg et al., 2018).

Second, the literature proposes a number of approaches to measure semantic drift (see Section 2), each has its pros and cons. In this short paper we propose another measure, inspired by (Gonen et al., 2020), and apply it on emoji. Future work should address changes in words and study the differences between the different measures.

Another limitation of the proposed method is a result of the fixed embeddings used. An emoji may have multiple senses, some may drift while others may be stable. Future work should account for the contextual differences in the use of an emoji within the basic time unit.

Finally, we study emoji since they provide a glimpse to the way visual changes affect semantics. However, some emoji went through major changes while other were changed slightly (compare the three visualizations 🍪, 🍪, 🍪 for the same code U+1F36A). Quantifying the *visual* distance is not a trivial task, although it may be required to get more accurate results (and see Footnote 3).

7 Conclusion

This work presents the first analysis of semantic drift of emoji, accounting for the *graphical variability* of emoji over time and platform. We proposed a novel statistical method for catching drifts and offered a substantive discussion of some of the factors that are unique to emoji use and to the modeling of drift in general. We note that the method

we proposed is not specific to emoji and can be applied to any token, as demonstrated in 2 (top).

References

- Wei Ai, Xuan Lu, Xuanzhe Liu, Ning Wang, Gang Huang, and Qiaozhu Mei. 2017. Untangling emoji popularity through semantic embeddings. In *Eleventh International AAAI Conference on Web and Social Media*.
- John Algeo. 1977. Blends, a structural and systemic view. *American Speech*, 52(1/2):47–64.
- John Algeo. 1980. Where do all the new words come from? *American Speech*, 55(4):264–277.
- John Langshaw Austin. 1975. *How to do things with words*. Oxford university press.
- Francesco Barbieri, Luís Marujo, Pradeep Karuturi, William Brendel, and Horacio Saggion. 2018. Exploring emoji usage and prediction through a temporal variation lens. In *Proceedings of the 1st International Workshop on Emoji Understanding and Applications in Social Media (Emoji 2018); 2018 Jun 25; Stanford, CA, USA.[Aachen]: CEUR; 2018*. CEUR Workshop Proceedings.
- 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)*.
- Laurel J Brinton and Elizabeth Closs Traugott. 2005. *Lexicalization and language change*. Cambridge University Press.
- Zhenpeng Chen, Xuan Lu, Wei Ai, Huoran Li, Qiaozhu Mei, and Xuanzhe Liu. 2018. Through a gender lens: learning usage patterns of emojis from large-scale android users. In *Proceedings of the 2018 World Wide Web Conference*, pages 763–772. International World Wide Web Conferences Steering Committee.
- Östen Dahl. 1979. Typology of sentence negation. *Linguistics*, 17(1-2):79–106.
- Thomas Dimson. 2015. [Machine learning for emoji trends](#).
- Haim Dubossarsky, Yulia Tsvetkov, Chris Dyer, and Eitan Grossman. 2015. A bottom up approach to category mapping and meaning change. In *NetWordS*, pages 66–70.
- Haim Dubossarsky, Daphna Weinshall, and Eitan Grossman. 2017. Outta control: Laws of semantic change and inherent biases in word representation models. In *Proceedings of the 2017 conference on empirical methods in natural language processing*, pages 1136–1145.

- Ben Eisner, Tim Rocktäschel, Isabelle Augenstein, Matko Bosnjak, and Sebastian Riedel. 2016. emoji2vec: Learning emoji representations from their description. In *Proceedings of The Fourth International Workshop on Natural Language Processing for Social Media*, pages 48–54.
- Vyvyan Evans. 2017. *The emoji code: The linguistics behind smiley faces and scaredy cats*. Picador USA.
- Bjarke Felbo, Alan Mislove, Anders Søgaard, Iyad Rahwan, and Sune Lehmann. 2017. Using millions of emoji occurrences to learn any-domain representations for detecting sentiment, emotion and sarcasm. In *2017 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics.
- Nikhil Garg, Londa Schiebinger, Dan Jurafsky, and James Zou. 2018. Word embeddings quantify 100 years of gender and ethnic stereotypes. *Proceedings of the National Academy of Sciences*, 115(16):E3635–E3644.
- Yoav Goldberg and Jon Orwant. 2013. A dataset of syntactic-ngrams over time from a very large corpus of english books.
- Eric Goldman. 2018. Emojis and the law. *Wash. L. Rev.*, 93:1227.
- Hila Gonen, Ganesh Jawahar, Djamel Seddah, and Yoav Goldberg. 2020. Simple, interpretable and stable method for detecting words with usage change across corpora. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 538–555.
- Sharath Chandra Guntuku, Mingyang Li, Louis Tay, and Lyle H Ungar. 2019. Studying cultural differences in emoji usage across the east and the west. In *Proceedings of the International AAAI Conference on Web and Social Media*, volume 13, pages 226–235.
- William L Hamilton, Jure Leskovec, and Dan Jurafsky. 2016a. Cultural shift or linguistic drift? comparing two computational measures of semantic change. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing. Conference on Empirical Methods in Natural Language Processing*, volume 2016, page 2116. NIH Public Access.
- William L Hamilton, Jure Leskovec, and Dan Jurafsky. 2016b. Diachronic word embeddings reveal statistical laws of semantic change. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1489–1501.
- Hans Henrich Hock and Brian D Joseph. 2009. *Language history, language change, and language relationship: An introduction to historical and comparative linguistics*, volume 218. Walter de Gruyter.
- Otto Jespersen. 1917. *Negation in English and other languages*. Kobenhavn: Host.
- Vivek Kulkarni, Rami Al-Rfou, Bryan Perozzi, and Steven Skiena. 2015. Statistically significant detection of linguistic change. In *Proceedings of the 24th International Conference on World Wide Web*, pages 625–635. International World Wide Web Conferences Steering Committee.
- Omer Levy, Steffen Remus, Chris Biemann, and Ido Dagan. 2015. Do supervised distributional methods really learn lexical inference relations? In *Proceedings of the 2015 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 970–976.
- Mingyang Li, Sharath Guntuku, Vinit Jakhetiya, and Lyle Ungar. 2019. Exploring (dis-) similarities in emoji-emotion association on twitter and weibo. In *Companion Proceedings of The 2019 World Wide Web Conference*, pages 461–467. ACM.
- Gretchen McCulloch. 2019. *Because Internet: Understanding the New Rules of Language*. Penguin Random House.
- Jean-Baptiste Michel, Yuan Kui Shen, Aviva Presser Aiden, Adrian Veres, Matthew K Gray, Joseph P Pickett, Dale Hoiberg, Dan Clancy, Peter Norvig, Jon Orwant, et al. 2011. Quantitative analysis of culture using millions of digitized books. *science*, 331(6014):176–182.
- Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. Distributed representations of words and phrases and their compositionality. *arXiv preprint arXiv:1310.4546*.
- Hannah Jean Miller, Jacob Thebault-Spieker, Shuo Chang, Isaac Johnson, Loren Terveen, and Brent Hecht. 2016. “blissfully happy” or “ready tofight”: Varying interpretations of emoji. In *Tenth International AAAI Conference on Web and Social Media*.
- Alexander Robertson, Farhana Ferdousi Liza, Dong Nguyen, Barbara McGillivray, and Scott A Hale. 2021. Semantic journeys: Quantifying change in emoji meaning from 2012–2018. In *Emoji Workshop, Fifteenth International AAAI Conference on Web and Social Media*.
- Antonio Blanco Salgueiro. 2010. Promises, threats, and the foundations of speech act theory. *Pragmatics. Quarterly Publication of the International Pragmatics Association (IPrA)*, 20(2):213–228.
- Yang Xu and Charles Kemp. 2015. A computational evaluation of two laws of semantic change. In *CogSci*.