

Emoji semantics/pragmatics: investigating commitment and lying

Benjamin Weissman

Rensselaer Polytechnic Institute – Department of Cognitive Science
weissb2@rpi.edu

Abstract

This paper presents the results of two experiments investigating the directness of emoji in constituting speaker meaning. This relationship is examined in two ways, with Experiment 1 testing whether speakers are committed to meanings they communicate via a single emoji and Experiment 2 testing whether that speaker is taken to have lied if that meaning is false and intended to deceive. Results indicate that emoji with high meaning agreement in general (i.e., pictorial representations of concrete objects or foods) reliably commit the speaker to that meaning and can constitute lying. Expressive emoji representing facial expressions and emotional states demonstrate a range of commitment and lie ratings: those with high meaning agreement constitute more commitment and more of a lie than those with less meaning agreement in the first place. Emoji can constitute speaker commitment and they can be lies, but this result does not apply uniformly to all emoji and is instead tied to agreement, conventionality, and lexicalization.

1 Introduction

Despite a multitude of studies focusing on emoji meanings, there has not yet been much research on the nature of these meanings with respect to semantics and pragmatics. The present research steps in this direction by investigating the relationships between emoji meaning and commitment and emoji meaning and lying. This paper presents the results of two studies probing the extent to which emoji constitute speaker commitment to content and the possibility of lying via emoji in order to better understand perceptions of the strength of meaning of emoji.

Studies on emoji meaning have ranged from how emoji supplement text with pragmatic information like irony (Garcia et al., 2022; Weissman & Tanner, 2018), emotional valence (Pfeifer et al., 2022), and indirect meaning (Holtgraves & Robinson, 2020) to more direct investigations of emoji meaning ratings and norms (Rodrigues et al., 2018; Was & Hamrick, 2021). In light of the wide range of communicative functions that emoji can fulfill (e.g., Beißwenger & Pappert, 2019; Dainas & Herring, 2021; Ge & Herring, 2018; Logi & Zappavigna, 2021; Yang & Liu, 2021), the nature of emoji meanings across these varied uses is rich ground for further research.

In research at the semantics/pragmatics interface, commitment offers a way to explore the meaning-making process – what we take a speaker to mean is related to what we take that speaker to have committed to. This has recently been explored with respect to inferences, such as presuppositions, implicatures, and explicatures. The link established from this research thus far ties together the notions of commitment, expression directness, and meaning: a more direct expression yields greater commitment and stronger meaning (e.g., Bonalumi et al., 2020; Boulat & Maillat, 2017; Mazzarella et al., 2018; Moeschler, 2013; Vullioud et al., 2017). An implicature, for example, communicates content less directly than saying that content literally; the speaker is thus less committed to the implicated content than the directly-said content. There is no universal definition for commitment, but it can be explored by testing whether speakers are taken as committed to certain propositions. Another lens with which to view this is deniability (e.g., Boogaart et al., 2021): an indirect (i.e., implicated) expression of content theoretically leaves the speaker room to deny that what the hearer understood is not what they

intended to mean with their utterance, while a direct expression of that content leaves the speaker no such room.

A recent proposal claims that an utterance can only be a lie if the speaker is committed to the relevant content (Reins & Wiegmann, 2021). If a speaker implies something false, but is not committed to that implicated content, the speaker is not taken to have lied. This is consistent with theoretical proposals (e.g., Saul, 2012) and experimental evidence (Weissman & Terkourafi, 2019) claiming that delivering false content through implicature is “merely misleading” rather than outright lying. Recent approaches have also provided support for the idea of lying in different modalities, like Viebahn’s (2019) investigation of lying with pictures – the present study tests extending this claim to emoji, a conventionalized and (to varying extents) lexicalized set of pictures.

The present research weaves together these threads of research to assess the link between emoji and speaker meaning via commitment and lying – are speakers committed to what they “say” if what they “say” is an emoji? Is it possible to lie via emoji? As emoji continue to grow in popularity (and, correspondingly, conventionality), perceptions of emoji constituting commitment and lying may change over time as well, consistent with the finding that emoji meaning changes over time (Robertson et al., 2021).

An important nuance to acknowledge in an emoji investigation like this is that not all emoji are the same; we should not necessarily expect all emoji to constitute speaker commitment in the same way. At the very least, there appear to be two broad categories of emoji: those that realistically depict real-world objects, animals, foods, etc. and those that more symbolically represent concepts like facial expressions, gestures, and other expressive meanings (see Grosz et al. (2021) for a semantic analysis demonstrating group differences between what they call “face” emoji and “activity” emoji or Maier (2021) for a different analysis between “entity” emoji and “facial expression” emoji). Just as there are different implicature types that do not all contribute to meaning in exactly the same way with exactly the same strength (e.g. Ariel, 2019; Doran et al., 2012; van Tiel et al., 2016), we may, too, expect similar nuance in emoji meaning-making.

Another level of nuance comes from emoji meaning agreement – not all emoji are equally

unambiguous in their links to meaning. 🍓 seems fairly clear and unambiguous in its representation of a strawberry, but an emoji like 😊 may mean different things to different people and in different contexts. To get an appropriately nuanced picture of the link between emoji and meaning, these experiments will test two sets of emoji: one set of high-meaning-agreement non-expressive emoji (objects, foods, and animals) and another set of expressive emoji (facial expressions and bodily gestures) that demonstrate a wide range of meaning agreement.

This paper presents the results of two experiments aimed at assessing the relationship between emoji and meaning commitment. The first experiment asks directly about commitment and deniability (partially following the approach used by Reins & Wiegmann (2021)); the second experiment probes lie ratings.

If emoji are found to yield uniformly less commitment and lower lie ratings than words, that would suggest that emoji are less direct in their meaning than words and as such contribute less strongly to speaker meaning. If emoji are found to yield uniformly as much commitment as words (with equal lie ratings), that would suggest that emoji are as direct in their meaning as words. A third possibility is that different emoji yield a range of attributions of commitment and lie ratings, which would suggest that emoji are capable of delivering speaker meaning but not all emoji do so in exactly the same way – we would thus end up with a more nuanced account of how emoji contribute to meaning, potentially related to emoji meaning agreement and consistency.

2 Pretest

Two meaning agreement surveys were carried out to hone in on the set of emoji to be used in the experiments, one for the non-expressive group and one for the expressive group. In both cases, participants were presented with a list of emoji and instructed to type in the meaning for each. Results were manually categorized by two raters into bins for each emoji and these bins were then ordered according to frequency. These results thus yield an agreement proportion for every emoji tested; if 92% of participants write that 🍅 means “tomato,” that emoji has an agreement of 0.92. Synonyms (e.g., “happy” and “glad”) were binned together but

similar non-synonyms (e.g., “happy” and “smiling”) were binned separately.

Data for the non-expressive group was collected on 80 emoji from 49 participants (average age = 30.29 (sd = 13.5), 41 female, 8 male) as part of another experiment; participants were compensated with course credit. was chosen manually. The 80 emoji chosen included foods, objects, and animals that were expected to have reasonably high agreement across participants. A set of 20 emoji all with agreement over 0.75 (average agreement = 0.88 (sd = 0.05)) were selected for use in the experiment.

Data for the expressive group was collected from a new set of 28 participants (average age = 31 (sd = 11.6), 22 female, 6 male). These participants were compensated \$2.55 USD for their participation in the survey, for an average prorated compensation of \$12.90/hr. 75 emoji, primarily gesture and facial expression emoji, were tested; from this group, a set of 20 emoji were selected for use in the experiment (average agreement = 0.62 (sd = 0.23)). Importantly, this set covered a range of agreement from 0.25 to 0.96.

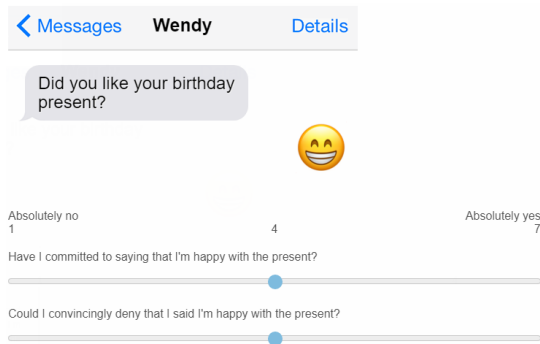


Figure 1 - Example stimulus from Experiment 1.

3 Experiment 1

Experiment 1 investigated the extent to which emoji messages yield commitment and deniability.

3.1 Methods

The emoji selected from the pretests were turned into experimental stimuli by creating question-answer pairs where the answer is a single emoji. These are presented as text messages that I (the experimenter) have received and sent. For each item, participants were asked to provide two ratings, corresponding to direct commitment (“Have I committed to saying X?”) and deniability

(“Could I convincingly deny that I said X?”). Ratings were provided on a sliding scale. An example is provided below in Figure 1.

100 participants (average age = 32.4; 72 female, 25 male, 2 non-binary, 1 not reported) recruited from Prolific participated in Experiment 1. Participants were compensated \$2.00 USD for their participation in the short survey, for an average prorated compensation of around \$18/hr.

Each participant saw 20 items. For 10 of these, the question matched the most-commonly-provided meaning for the emoji; for 10 of these, the question asked about some other meaning that was clearly not present in the text message. These mismatch trials should, in theory, always receive a score at the low end of the scale – it would be rare for a participant to interpret 😄 as “excited” or 🐘 as “elephant.” The mismatching set was included as fillers to counterbalance the experimental items, and the specific ratings given for any item in that condition are dependent upon the exact mismatching alternative provided; as such, these are not discussed further in analysis here. So that every emoji could be seen equally in both conditions, 4 lists were created. Each participant saw 10 expressive emoji items and 10 non-expressive emoji items, counterbalanced across the match/mismatch conditions.

3.2 Results

Results given on the sliding scale were converted to a traditional 1-7 scale. The direct commitment and (inversed) deniability ratings were merged into a single commitment rating for each participant for item. Overall, as shown in Figure 2, the non-expressive set (average = 6.09, sd = 1.11) yielded significantly higher commitment ratings than the expressive set (5.17, sd = 1.59) as modeled by a linear mixed effects model with random slopes for emoji type and random intercepts for participant ($F(1,99) = 114.2, p < .001$).

The expressive set was investigated further to clarify the relationship between agreement and commitment. There was a significant effect of pre-test agreement ratings on commitment ratings according to the linear mixed effects model with random slopes for emoji type and random intercepts for participant ($F(1,998) = 20.5, p < .001$). Emoji that demonstrated lower population-wide agreement yielded lower commitment ratings in this task; emoji with higher agreement yielded

higher commitment ratings. This correlative relationship is shown in Figure 3.




Figure 3 - Violin plot of commitment/deniability ratings for expressive and non-expressive emoji in Experiment 1. Horizontal bar in each column represents median; boxes extend to first and third quartiles.

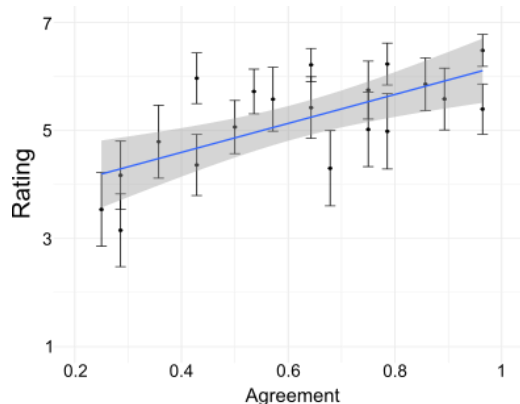


Figure 4 - Correlative relationship between pretest emoji meaning agreement and Experiment 1 commitment/deniability ratings.

4 Experiment 2

Experiment 2 investigates the extent to which it is possible to lie via emoji.

4.1 Methods

The same question/answer text messages from Experiment 1 were used in Experiment 2 with context added. The context was in the form of a few sentences presented above the text message picture and worked to establish whether the answer that I provide in the text message is true or false. In the true condition, the context reveals that my answer

is true (i.e., matches what really happened or what I really believe to be true); in the false condition, the context revealed that my answer is false (i.e., does not match what happened or what I believe to be true) and includes motivation for me to deceive the interlocutor. An example of the “lie” condition

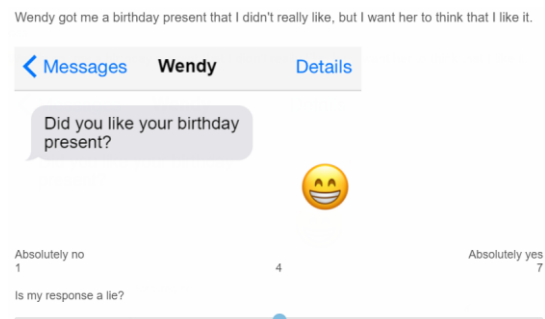


Figure 2 - Example stimulus from Experiment 2.

is shown in Figure 4.

202 new participants (average age = 33.7, sd = 10.3; 151 female, 50 male, 1 genderqueer), none of whom participated in Experiment 1, were recruited from Prolific. Data from 7 participants was discarded due to consistently unreliable answers to filler items, yielding a final dataset from 195 participants.

As in Experiment 1, 4 lists were created, each with 10 expressive and 10 non-expressive emoji, counterbalanced across 10 true and 10 false responses. In this experiment, however, those 4 lists were repeated with word responses instead of emoji responses to allow for a word vs. emoji comparison.

4.2 Results

A linear mixed effects model with random intercepts for participant estimated a significant interaction between type (expressive/non-expressive) and presentation (word/emoji) ($F(1,1753) = 3.90, p = 0.048$). Lie ratings for emoji were significantly lower than lie ratings for words, but this difference was significantly greater for expressive items than non-expressive items. Figure 5 portrays this relationship graphically.




Figure 5 - Boxplot depicting lie ratings for expressive and non-expressive emoji and words. Horizontal line in each column represents median; boxes extend to first and third quartiles.

The expressive set was again investigated further to explore the relationship between commitment and lie ratings for emoji messages. A linear mixed effects model with random intercepts for participant yielded a significant effect of commitment on lie ratings ($F(1,432) = 15.13, p < .001$) in the expected direction – emoji to which participants in Experiment 1 attributed a higher degree of commitment yielded higher lie ratings in Experiment 2 when the message meaning is revealed to be false.

5 Discussion

These two experiments have provided evidence that emoji can constitute speaker commitment and it is possible to lie via emoji, but the extent to which they contribute to speaker meaning is not the same for all emoji. Emoji that have high meaning agreement in the first place (from another perspective, emoji that are farther along in the lexicalization or conventionalization process) contribute more directly to speaker meaning; the correlative relationships between agreement, commitment, and lie ratings highlight this finding. A fully lexicalized emoji with high meaning agreement and consistency does not leave much room for varying interpretation and accordingly yield the speaker less deniability. Though the focus of this paper is on emoji themselves, further study in this direction can work towards establishing a taxonomy of semantic/pragmatic commitment across modalities.

In these results, there is a distinction between expressive emoji (e.g., facial expressions) and non-expressive emoji (e.g., objects) with respect to commitment and lying, but this is mediated by their meaning agreement in the first place. In other words, non-expressive emoji do not inherently contribute more to commitment than expressive emoji, but this difference surfaces because non-expressive emoji are more likely to be direct and unambiguous representations of their assigned meanings. On the other hand, expressive emoji are likely to contain more ambiguity and potential polysemy, yielding less agreement over their meanings in the first place. This finding is highlighted by the significant correlation between agreement and commitment among the expressive emoji tested.

Further research can explore the context-sensitivity of these relationships. Since context significantly affects both lie ratings (Weissman, 2019a) and emoji interpretations (e.g., Miller et al., 2017; Weissman, 2019b), a more nuanced look at this complex relationship is likely warranted. As recent work has begun exploring multi-emoji sequences and the extent to which those are (un)natural forms of expressing content (e.g., Cohn et al., 2019; Herring & Ge, 2020; McCulloch & Gawne, 2018), this endeavor could benefit from a commitment-based analysis as well. Discussions of emoji interpretation challenges in the courtroom has already begun (e.g., Foltz & Fray, 2020). In today’s age, it does not seem far-fetched to imagine a public figure mired in an emoji-related scandal – such an occasion would certainly provide a fascinating case study for the deniability and commitment of emoji.

Acknowledgments

Thank you to Camdyn Vanek for their assistance in developing materials to be used in this research.

References

- Ariel, M. (2019). Different prominences for different inferences. *Journal of Pragmatics*, 154, 103–116.
- Beißwenger, M., & Pappert, S. (2019). How to be polite with emoji: A pragmatic analysis of face work strategies in an online learning environment. *European Journal of Applied Linguistics*, 7(2), 225–253.
- Bonalumi, F., Scott-Phillips, T., Tacha, J., & Heintz, C. (2020). Commitment and communication: Are we

- committed to what we mean, or what we say? *Language and Cognition*, 12(2), 360–384.
- Boogaart, R., Jansen, H., & van Leeuwen, M. (2021). “Those are your words, not mine!” Defence strategies for denying speaker commitment. *Argumentation*, 35(2), 209–235.
- Boulat, K., & Maillat, D. (2017). She said you said I saw it with my own eyes: a pragmatic account of commitment. In *Formal models in the study of language* (pp. 261–279). Springer.
- Cohn, N., Engelen, J., & Schilperoord, J. (2019). The grammar of emoji? Constraints on communicative pictorial sequencing. *Cognitive Research: Principles and Implications*, 4(1).
- Dainas, A. R., & Herring, S. C. (2021). Interpreting emoji pragmatics. *Approaches to Internet Pragmatics*, 107-144.
- Doran, R., Ward, G., Larson, M., McNabb, Y., & Baker, R. E. (2012). A novel experimental paradigm for distinguishing between what is said and what is implicated. *Language*, 88(1), 124–154.
- Foltz, K. A., & Fray, J. (2020). Emoji goes to court: An analysis of emoji in court proceedings and implications for legal practice. *Communication Law Review*, 20, 109-125.
- Garcia, C., Turcan, A., Howman, H., & Filik, R. (2022). Emoji as a tool to aid the comprehension of written sarcasm: Evidence from younger and older adults. *Computers in Human Behavior*, 126(August 2021), 106971.
- Ge, J., & Herring, S. C. (2018). Communicative functions of emoji sequences on Sina Weibo. *First Monday*, 23(11).
- Grosz, P., Kaiser, E., & Pierini, F. (2021). Discourse anaphoricity and first-person indexicality in emoji resolution. *Proceedings of Sinn Und Bedeutung*, 25.
- Herring, S., & Ge, J. (2020). Do Emoji Sequences Have a Preferred Word Order? 14th International AAAI Conference on Web and Social Media, (June), 1–5.
- Holtgraves, T., & Robinson, C. (2020). Emoji can facilitate recognition of conveyed indirect meaning. *PLoS ONE*, 15(4), 1–13.
- Logi, L., & Zappavigna, M. (2021). A social semiotic perspective on emoji: How emoji and language interact to make meaning in digital messages. *New Media and Society*, (September).
- Maier, E. (2021). Emoji as pictures. Under review. <https://ling.auf.net/lingbuzz/006025>.
- Mazzarella, D., Reinecke, R., Noveck, I., & Mercier, H. (2018). Saying, presupposing and implicating: How pragmatics modulates commitment. *Journal of Pragmatics*, 133(June), 15–27.
- McCulloch, G., & Gawne, L. (2018). Emoji grammar as beat gestures. *CEUR Workshop Proceedings*, 2130, 3–6.
- Miller, H., Kluver, D., Thebault-Spieker, J., Terveen, L., & Hecht, B. (2017). Understanding Emoji Ambiguity in Context: The Role of Text in Emoji-Related Miscommunication. *Proceedings of the Eleventh International AAAI Conference on Web and Social Media*, 152–161.
- Moeschler, J. (2013). Is a speaker-based pragmatics possible? Or how can a hearer infer a speaker’s commitment? *Journal of Pragmatics*, 48(1), 84–97.
- Pfeifer, V. A., Armstrong, E. L., & Lai, V. T. (2022). Do all facial emoji communicate emotion? The impact of facial emoji on perceived sender emotion and text processing. *Computers in Human Behavior*, 126, 107016.
- Reins, L. M., & Wiegmann, A. (2021). Is lying bound to commitment? Empirically investigating deceptive presuppositions, implicatures, and actions. *Cognitive Science*, 45(2).
- Robertson, A., Liza, F. F., Nguyen, D., McGillivray, B., & Hale, S. A. (2021). Semantic journeys: quantifying change in emoji meaning from 2012-2018. *Emoji2021*. http://workshop-proceedings.icwsm.org/pdf/2021_07.pdf
- Rodrigues, D., Prada, M., Gaspar, R., Garrido, M. V., & Lopes, D. (2018). Lisbon Emoji and Emoticon Database (LEED): Norms for emoji and emoticons in seven evaluative dimensions. *Behavior Research Methods*, 50(1), 392–405.
- Saul, J. M. (2012). *Lying, Misleading, and What is Said: An Exploration in Philosophy of Language and in Ethics*. Oxford: Oxford University Press.
- van Tiel, B., van Miltenburg, E., Zevakhina, N., & Geurts, B. (2016). Scalar diversity. *Journal of Semantics*, 33(1), 137–175.
- Viebahn, E. (2019). Lying with pictures. *British Journal of Aesthetics*, 59(3), 243–257.
- Vullioud, C., Clément, F., Scott-Phillips, T., & Mercier, H. (2017). Confidence as an expression of commitment: why misplaced expressions of

confidence backfire. *Evolution and Human Behavior*, 38(1), 9–17.

Was, C. A., & Hamrick, P. (2021). What Did They Mean by That? Young Adults’ Interpretations of 105 Common Emoji. *Frontiers in Psychology*, 12(May), 1–7.

Weissman, B. (2019a). The role of linguistic meaning in the concept of lying. PhD Dissertation. University of Illinois at Urbana-Champaign.

Weissman, B. (2019b). Peaches and eggplants or. . . something else? The role of context in emoji interpretations. *Proceedings of the Linguistic Society of America*, 4(1), 29.

Weissman, B., & Tanner, D. (2018). A strong wink between verbal and emoji-based irony: How the brain processes ironic emoji during language comprehension. *PLoS ONE*, 13(8).

Weissman, B., & Terkourafi, M. (2019). Are false implicatures lies? An experimental investigation. *Mind & Language*, 34(2), 221–246.

Yang, X., & Liu, M. (2021). The pragmatics of text-emoji co-occurrences on Chinese social media. *Pragmatics*, 31(1), 144–172.

A Supplementary Material

Materials and data are available at: https://osf.io/wtnzy/?view_only=48586c4d47f946fd9efc1044c995e32e

B Appendix

This appendix lists the emoji used in the experiment, the most-commonly-provided meaning for each emoji from the pre-test, and the agreement rating for each emoji from the pre-test. Expressive emoji appear in Table 1; non-expressive in Table 2.





















Emoji	Meaning	Agreement
	angry	.964
	silly	.964
	cool	.893
	shocked	.857
	celebrate	.786
	lips are sealed	.786
	kiss	.750
	angel	.750
	neutral	.679
	happy	.643
	crying	.643
	upset	.571
	tired	.536
	afraid	.500
	stressed	.429
	no words	.429
	cringe	.358
	annoyed	.286
	sassy	.286
	goofy	.250

Table 1: expressive emoji used in both experiments





















Emoji	Meaning	Agreement
	church	0.959
	hospital	0.959
	tomato	0.939
	spider	0.918
	key	0.918
	castle	0.918
	donut	0.918
	salad	0.898
	bike	0.898
	basketball	0.898
	astronaut	0.878
	pancakes	0.878
	taco	0.878
	grapes	0.878
	snake	0.857
	dragon	0.857
	roller coaster	0.837
	violin	0.816
	football	0.776
	paper clip	0.776

Table 2: non-expressive emoji used in both experiments