

# Semantic Journeys: Quantifying Change in Emoji Meaning from 2012–2018

Alexander Robertson,<sup>1</sup> Farhana Ferdousi Liza,<sup>2</sup>  
Dong Nguyen,<sup>3</sup> Barbara McGillivray,<sup>4,6</sup> Scott A. Hale<sup>5,6</sup>

<sup>1</sup>University of Edinburgh <sup>2</sup>University of Essex <sup>3</sup>Utrecht University  
<sup>4</sup>University of Cambridge <sup>5</sup>University of Oxford <sup>6</sup>Alan Turing Institute  
alexander.robertson@ed.ac.uk, farhana.ferdousi.liza@essex.ac.uk,  
d.p.nguyen@uu.nl, bmcgillivray@turing.ac.uk, scott.hale@oii.ox.ac.uk

## Abstract

The semantics of emoji has, to date, been considered from a static perspective. We offer the first longitudinal study of how emoji semantics changes over time, applying techniques from computational linguistics to six years of Twitter data. We identify five patterns in emoji semantic development and find evidence that the less abstract an emoji is, the more likely it is to undergo semantic change. In addition, we analyse select emoji in more detail, examining the effect of seasonality and world events on emoji semantics. To aid future work on emoji and semantics, we make our data publicly available along with a web-based interface that anyone can use to explore semantic change in emoji.

## 1 Introduction

Semantic change is the development of a word’s meaning over time, e.g. *cute* (*shrewd* to *pleasant*) and *gay* (*bright* to *homosexual*). Such changes are only known to us because instances of both usage can be found in the historical record. It is entirely possible that there are other shifts and meanings of which we have no knowledge, purely because we have no direct evidence. This lack of evidence may be due to the loss of historical documents, or because the speech community it was found in did not record it before themselves dying out.

Emoji offer a unique perspective on this. Unlike words, we know exactly when each individual emoji came into existence and are pre-warned of new ones. The origin of words is far less clear, with new coinings and neologisms difficult to accurately detect. We also have unparalleled access to instances of emoji usage, through social media. We find ourselves in the hitherto unimaginable situation of being able to observe semantic change directly, from the birth of a term right up to the present moment.

Perhaps owing to their newness, prior work on emoji has taken a synchronic approach—studying emoji in general with no regard to time, implicitly assuming that emoji today are the same as ten days or years ago. We depart from this and present a diachronic view of emoji semantics, examining how the meaning of individual emoji evolves in six years of Twitter data. This approach is motivated by the development of techniques for semantic change detection, which so far have been applied only to words.

The main contribution of this work is the application of semantic change detection techniques to emoji and we present the first longitudinal analysis of emoji semantics. We examine semantic development over time and search for patterns in this development. Although our analyses are mainly exploratory, we also analyse in more detail some of the patterns we observe. We conclude with case studies of a few select emoji, presenting their semantic journeys in more detail and account for their specific patterns in more detail by relating them to world events.

Our finding that emoji can undergo semantic change over time will be of interest to linguists, and may motivate the application of linguistic theories of semantic change and development to emoji. They also highlight the need to consider semantic variation and change in emoji when designing NLP systems which include emoji—emoji semantics are not as static and simplistic as might be assumed. To aid future work, we make our data publicly available and release a website<sup>1</sup> where anyone can explore emoji semantic change.

## 2 Previous Work

There is a growing literature on the semantics and pragmatics of emoji. Core themes are the sentiment associated with emoji (Novak et al. 2015); how individual emoji are interpreted by readers, with context (Miller et al. 2016) or without context (Miller et al. 2017); how people use new emoji features like skin-tone modifiers (Robertson, Magdy, and Goldwater 2018, 2020); how particular emoji affect the interpretation of messages (Tigwell and Flatla 2016); and how these findings compare to emoji-like constructs such as Animoji (Herring et al. 2020). There has been less work on the syntax of emoji and though there is evidence of an emergent syntax (Herring and Ge 2020), it is emoji’s semantic and pragmatic properties which are most suggestive of them being, if not a language, at least language-like.

Research into computational methods for the automatic detection of semantic change is currently an active area of research. Researchers have proposed a range of different approaches, from topic-based models (Cook et al. 2014; Lau et al. 2014; Frermann and Lapata 2016), to graph-based models (Mitra et al. 2015; Tahmasebi and Risse 2017), and

Version	Date	Emoji	Monthly embeddings		
			median	min	max
unicode-6.0	Oct 2011	494	57	1	66
unicode-6.1	Jan 2012	13	58	58	58
unicode-7.0	Jun 2014	40	16	1	26
unicode-8.0	Jun 2015	11	22	4	25
emoji-3.0	Jun 2016	3	10	6	10

Table 1: Emoji per release, with summary of number of monthly embeddings per emoji.

word embeddings (Kim et al. 2014; Basile and McGillivray 2018; Kulkarni et al. 2015; Hamilton, Leskovec, and Jurafsky 2016b; Dubossarsky, Weinshall, and Grossman 2017; Tahmasebi 2018; Dubossarsky et al. 2019). Tahmasebi, Borin, and Jatowt (2018) and Kutuzov et al. (2018) provide good overviews of this field and Schlechtweg et al. (2020) report on the results of the first SemEval shared task on semantic change detection. The most successful methods are based on word embeddings, and most work has focused on semantic change over a relatively long time span, typically a few centuries (Hamilton, Leskovec, and Jurafsky 2016a; Perrone et al. 2019; Schlechtweg et al. 2020), with some notable exceptions, such as Shoemark et al. (2019), who perform a systematic evaluation of embedding-based methods for short-term semantic change detection using Twitter data from 2011 to 2017. All methods proposed so far have dealt with semantic change of words, and to our knowledge no work has yet been done on applying this research to emoji specifically.

Emoji are generally considered as time-invariant artefacts, exhibiting variation only between each other or between populations of emoji users: i.e. iOS versus Android (Miller Hillberg et al. 2018), young versus old (Herring and Dainas 2020), male versus female (Chen et al. 2017). However, there is some work on how small close-knit groups apply their own special meaning to specific emoji (Wiseman and Gould 2018).

### 3 Data

We start with the publicly available dataset of monthly Twitter embeddings created by Shoemark et al. (2019), who used them to evaluate different embedding and semantic change detection methods. Summary statistics are shown in Table 1.

Due to sporadic data collection issues, we do not have complete data for some months. These are therefore excluded. Due to tokenisation issues<sup>2</sup>, we can retain only those emoji which are only found as a single Unicode sequence. For example, 🍌 is actually composed of 🍌 and 🍌. However, 🍌 can appear on its own, so its embedding will have been trained on 🍌 and 🍌. This risks clouding the semantics encoded by the embedding. This issue affected 45 single codepoint emoji in total, plus all sequences they appear in. We have no data for Emoji 1.0 or 2.0 releases and very

<sup>2</sup>Which are understandable, given Shoemark et al. (2019) focused on words, not emoji.

little for 3.0—these contain mainly multi-codepoint emoji sequences.

The first month for which we have data is January 2012. As emoji were only widely available outside Japan from October 2011, when Apple released iOS 5, the trained embeddings very nearly capture the “birth” of many Unicode 6.0 and 6.1 emoji and we have an almost complete timeline of embeddings for most of these. In our analyses, we therefore focus on Unicode 6.x emoji but do include all emoji in our first explorations of the data.

## 4 Methods

The aim of this paper is to characterise the development of emoji semantics from 2012 to 2018. We apply methods from prior work on words’ semantic change. Due to the large number of emoji we examine, we use unsupervised clustering techniques to automatically identify patterns within the developmental trajectories we extract. In this section, we provide the technical details of our techniques.

### 4.1 Measuring semantic change

We measure semantic change at each month of an emoji’s existence, following the local neighbourhood measure of semantic change described by Hamilton, Leskovec, and Jurafsky (2016a). This technique has been widely used, for example to track semantic change over time in gender/ethnic stereotypes (Garg et al. 2018) or trace the dynamics of global armed conflicts (Kutuzov, Velldal, and Øvrelid 2017). Crucially, it has also been used to track semantic change on social media and found to perform well in that context (Shoemark et al. 2019).

The local neighbourhood measure of semantic change works as follows. For a target word (or emoji, in our case), the first available monthly embedding is set as an “anchor” as per Shoemark et al. (2019). All subsequent monthly embeddings are compared against this. The comparison involves constructing two “second-order” vectors for each embedding, whose components are the cosine similarities between the target word and its  $k$  nearest neighbours. The cosine similarity between these two vectors represents the local neighbourhood change for that timestep. By always calculating semantic change scores relative to the second-order vector generated from a word’s first appearance (its anchor point), we can measure that item’s semantic change score over time. From here on, we refer to this as the semantic change (SC) score.

We use 25 neighbours per second-order vector (as per Hamilton, Leskovec, and Jurafsky (2016a)). We exclude hashtags from the list of nearest neighbours as they are generally ephemeral, linked to specific events or trends and serve a variety of platform-specific purposes (Chang 2010; Kwak et al. 2010; Kowald, Pujari, and Lex 2017). We also exclude emoji, to better understand emoji relative to words. The 25 nearest neighbours are taken from a filtered selection of the 500 nearest neighbours<sup>3</sup>.

<sup>3</sup>Because we require each neighbour to appear in both time steps, a second-order vector may have fewer than 25 components because vocabularies change at each time step. In practice this af-




Figure 1: Mean semantic change score of emoji over time. The standard deviation of the data at each point is shown by the shaded areas. Upper-left plot: all emoji. Others: specific emoji releases. Dashed vertical lines denote addition of new emoji to the Unicode Standard by the Unicode Consortium. The solid vertical line denotes the date of release for that Unicode/Emoji version. First occurrences of emoji lag behind their Unicode addition as they are not immediately implemented by vendors such as Apple and Google.

## 4.2 Clustering

We can construct a time series using SC at each time point, showing an emoji’s change over time. Clustering them, to identify common patterns, proved difficult using standard techniques due to missing data points (some emoji have no embeddings for some months) and time series of differing lengths (emoji can have different anchor points and different end points). Imputation of missing values through linear interpolation can help with the former, but not the latter. One approach would be to focus only on the 125 emoji for which we have complete data, or to separately cluster groups of emoji based on their having the same number of data points. We feel this would limit the scope of our analysis—that one emoji has 65 months of data and another 64 should not warrant separate analyses.

Instead we use a simple technique to generate fixed-length vectors for any emoji. First, we impute missing values using linear interpolation, but only for months between an emoji’s first appearance (i.e. the anchor point) and the last month that emoji has data. This prevents us from accidentally “inventing” a starting/ending pattern for any emoji.

Next, because the data is highly specific to each emoji, we apply exponential smoothing. This technique slides a window across each time series, fitting that local subset of points to a polynomial by least squares. The new time series has a smoother appearance while retaining the overall characteristics of the original. This makes it more suitable for applying similarity-based algorithms, where we are not necessarily concerned about the finer details of the data. This is a common step in signal processing. We use the Savitzky Golay method (Savitzky and Golay 1964) for smoothing. We use a small sliding window of 5 and a polynomial of degree 3 to avoid over-fitting as this would exaggerate some of the local

facts  $\sim 2,000$  of the  $\sim 25,000$  vector pairs we construct and 90% of these have at least 15 components.

characteristics of the time series we wish to filter out: see Hassanpour, Zehtabian, and Sadati (2012) for an overview.

The smoothed data amplifies the general characteristics of the semantic trajectories, but they still have very different scales and start/begin at different points in time. We therefore follow the recommendation of Keogh and Kasetty (2003)—each smoothed time series was individually z-normalised, which scales and shifts a time series such that it has zero mean and zero standard deviation.

Finally, we calculate the similarity of all pairs of emoji timelines using Dynamic Time Warping (DTW) (Salvador and Chan 2007). This method allows comparison of time series of different lengths using any distance metric—we use Euclidean distance. For any given emoji, we take the 10 other emoji with the most similar time series, as determined by DTW. After one-hot encoding the names of the most similar emoji to create a feature vector, we use hierarchical clustering to group them. This algorithm is suited to finding unevenly-sized clusters: we do not expect all semantic patterns in emoji to describe uniform numbers of emoji.

This approach is in a similar spirit to Mylonas, Wallace, and Kollias (2004) who used pairwise distance metrics between individual features of high dimensionality datasets as a form of dimensionality reduction, stating “the primary aim of clustering algorithms is not to correctly classify data, but rather to identify the patterns that underlie in it and produce clusters of similar data samples. Therefore, ‘wrong elements’ in clusters may be acceptable, as long as the overall cluster correctly describes an existing and meaningful pattern”. This concept is more formally presented as exploiting ultrametric embeddings by Murtagh, Downs, and Contreras (2008) who, again, use the concept of feature-by-feature similarity to find informative subsets of features to perform “data condensation” on very high-dimensionality data. The main difference in our method is that we induce a compact form of representation based on similarity of each entire

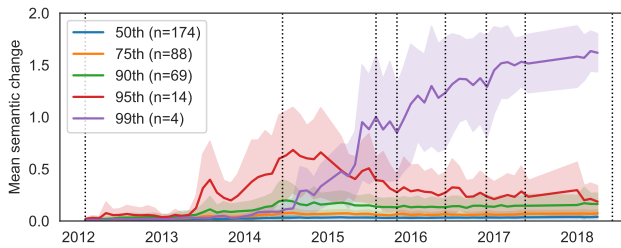


Figure 2: Mean semantic score change over time for Unicode 6.x emoji, grouped by percentile based on standard deviation of their month-to-month semantic change scores. The standard deviation of the data at each point is shown by the shaded areas.

item in our dataset, rather than on the individual features of the items.

## 5 Analyses

We first explore the overall picture of semantic change in emoji, using the techniques described in Section 4.1 on page 2. We focus on 350 emoji with anchor points in 2012, since we have the most data for these. This will allow us to observe the long-term semantic development of a large group of emoji from very close to their release in late 2011 until 2018. We then extend this by looking more closely at individual emoji in this group, to get an understanding of any differences within the group. For these analyses we use SC scores over time, at the aggregate and individual level.

Next, we apply the clustering approach described in Section 4.2 on the previous page to identify more specific patterns of semantic change. For example, do some emoji change very drastically compared to their initial selves as observed in 2012? Perhaps some have essentially fixed semantics. We are interested in quantifying how many patterns there are and how the emoji are distributed across these patterns.

### 5.1 An initial inventory of semantic development

Using the method described in Section 4.1 on page 2, we plot the mean month-to-month SC for all emoji, and then group by the Unicode version in which they first appeared (Figure 1).

Considering all emoji together, a general pattern emerges: an initial period where semantics changes very little, followed by a gradual increase in SC scores. Later additions to the standard, though we have less data for them, appear to be following the same general trend as the more established Unicode 6.0 emoji. The standard deviations for SC indicate that this change is not evenly distributed. It is likely to be driven by some subset of emoji. This is especially the case for Unicode 6.0 emoji. Emoji in other releases appear to undergo less semantic change, but again there is a gradual increase in the mean SC scores accompanied by an increase in the standard deviation of those scores.

To tease apart the aggregate picture and determine which emoji contribute to the variance in SC, we calculate per-




Figure 3: Semantic change over time for 5 random emoji in the 50th (top), 75th, 90th, 95th and 99th (bottom) percentile, based on standard deviation of month-to-month semantic change for Unicode 6.x emoji with 2012 anchor points.

emoji standard deviation of the SC scores and group them into five percentiles. We focus on emoji which have anchor points in 2012—these are all Unicode 6.0 or 6.1 and comprise 348 emoji<sup>4</sup> out of a possible total of 507 emoji. Figure 2 shows the mean SC for these subsets. The variation observed in Figure 1 on the preceding page is due to a very small subset of emoji with a wider range of SC scores over time, rather than all emoji having highly variable SC. This is illustrated very clearly in Figure 3, where we show the SC scores for a random sample of each percentile.

The vast majority of emoji, those not in the 95th percentile or above, exhibit very little semantic change over the 6 years, though we do observe some small peaks here and there. The more variable emoji of the 95th percentile appear to exhibit larger, more prolonged, peaks of change. The 99th percentile, just 4 emoji, show significant and sustained change in the second half of their lives.

From these more initial patterns we can begin to build an inventory of emoji semantic development. For most emoji, their semantics is stable from the beginning. Though SC scores may trend upwards over time, this is very slight. For a few emoji, their semantics has highly changed since their first introduction. And in between these two extremes are emoji which, at some point in their life so far, have experienced a rise and fall (to different extents) in their SC scores, suggesting they diverged from their initial meaning but only temporarily.

### 5.2 A detailed inventory of semantic development

Our analysis so far has looked at very broad patterns of semantic development. The finding that only a small number of emoji appear to undergo substantive semantic change, and that these changes can happen at different points in time and

<sup>4</sup>Not 350, as 🏠 and 🚗 only have one embedding each due to low frequency.

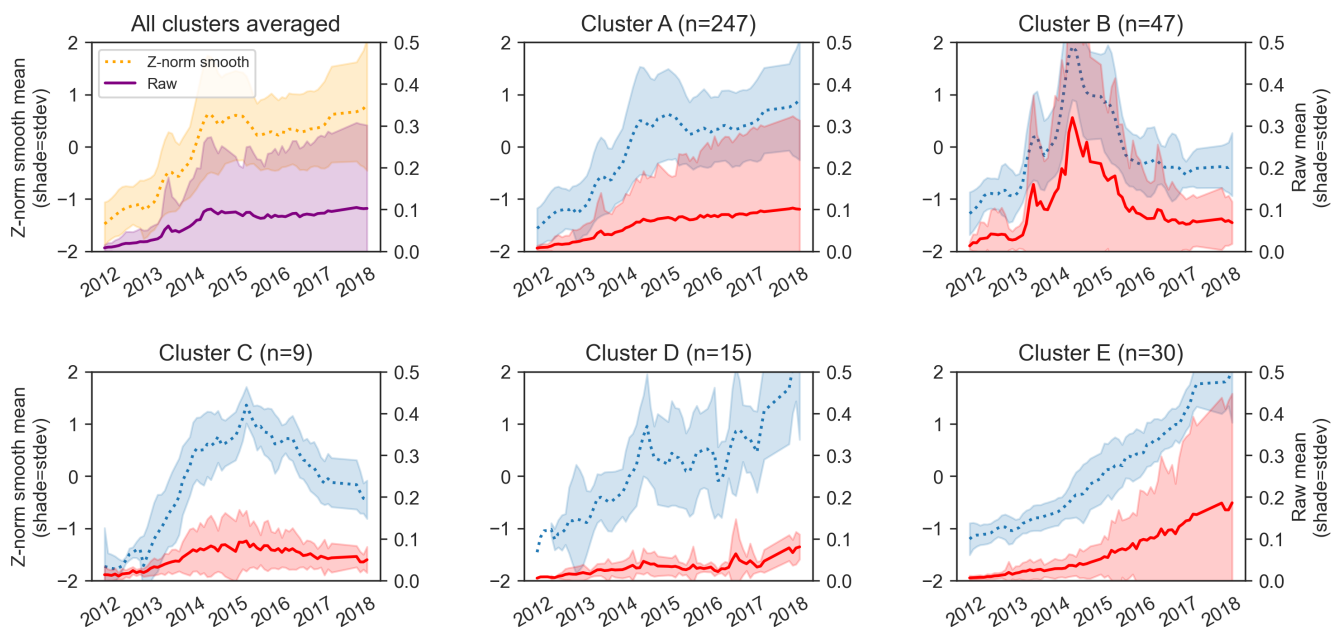


Figure 4: Progression of semantic change in 348 Unicode 6.x emoji, with 2012 anchor points, from 2012 to 2018. Each plot shows a cluster’s characteristic shape (dashed lines) and the mean of the actual observed semantic change scores (solid lines). Top left shows the averages over all clusters. The standard deviation of the data at each point is shown by the shaded areas.

to different extents (see Figure 3 on the preceding page), motivates our effort to classify emoji according to their semantic change patterns.

As we are dealing with many emoji, we employ the clustering technique detailed in Section 4.2 on page 3 to automate the process. For each emoji we generate a characteristic shape of the timeline of its monthly SC scores. We compare all pairs of characteristic shapes to find the ten emoji with the most similar timeline shape. These ten emoji are one-hot encoded to create a fixed-length feature vector for use in hierarchical agglomerative clustering. We use the same 348 emoji as in the previous analysis. We identify 5 clusters—this number was chosen manually, such that the mean characteristic shapes of each cluster were visually dissimilar but without there being very few members in each cluster. Results are shown in Figure 4.

The trend for all emoji matches that of Figure 1 on page 3—emoji semantics changes gradually before settling down. But again there is a large degree of variation. Within this variation, we see five distinct patterns.

Cluster A (gradually establishing, rank 1/5 with  $n=247$ ) matches the general trend of the aggregate data—most emoji change relatively slowly, though there is some variation in the extent of this change. Though their semantics has diverged from their initial form to different degrees (as shown by the standard deviation), they are relatively stable.

Cluster B (extreme sudden peak, rank 2/5 with  $n=47$ ) contains emoji which, overall, remain similar to their initial versions. However, at some point they experiences an extreme but temporary change in their semantics.

Emoji in cluster C (slight gradual peak, rank 5/5 with  $n=9$ ) also experience a temporary change in semantics but to a far lesser degree than Cluster B emoji. The extent of the changes is also more similar within this cluster, with less variation, as shown by the far lower standard deviation.

Cluster D (static, rank 3/5 with  $n=15$ ) changed very little for the larger part of their life but may be undergoing change now. These emoji have been otherwise very stable.

Cluster E (not yet established, rank 3/4 with  $n=30$ ) have been changing from their initial semantics ever since the start. This change appears to be accelerating but to different extents for the emoji in this cluster.

This analysis gives a more nuanced picture of emoji semantic development. Emoji generally have fixed, stable semantics within the time period we observed. The main factor contributing to variation in Cluster A could be due to seasonality—this cluster contains 🍌 and 🍌 and their SC scores swing dramatically once per year, but on average, are very similar to their 2012 SC scores. See Figure 6 on page 7 for the 🍌-specific data.

## 6 The role of concreteness in emoji semantic change

Our analyses showed that subsets of the 350 emoji introduced around 2012 underwent different semantic journeys over the next six years. This could be due to the lack of any prescriptive authority (or intent) for emoji, as discussed in Section 1 on page 1, but we conjecture that some emoji have more potential for having fixed semantics.

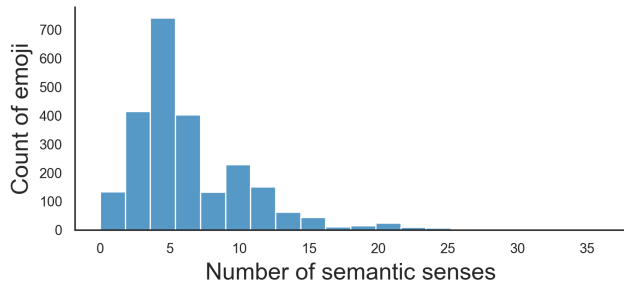


Figure 5: Distribution of the number of senses per emoji, based on EmojiNet data linked to BabelNet senses.

Previous computational studies on the distributional semantic properties of concrete and abstract words (Naumann, Frassinelli, and Schulte im Walde 2018) have found evidence for consistent differences between these two groups. More specifically, Naumann, Frassinelli, and Schulte im Walde (2018) found partial support for the hypothesis that concrete words have significantly less diverse contexts compared to abstract words. Given their specific characteristics, emoji represent an important linguistic category to investigate alongside words with respect to properties along the concrete/abstract scale (Wicke and Bolognesi 2020). To investigate the relation between concrete versus abstract emoji and their degree of semantic change, we devised a pipeline drawing on EmojiNet (Wijeratne et al. 2017), BabelNet (Navigli and Ponzetto 2012) and the concreteness dataset by Brysbaert, Warriner, and Kuperman (2014). As a first step we used EmojiNet to determine how many semantic senses emoji have. The distribution is shown in Figure 5.

We consider an emoji to have undergone a high degree of semantic change if it is in the top 10% in terms of standard deviation of its month-to-month SC scores. This criterion selects 53 emoji. As EmojiNet provides senses but not lemmas which we can easily look up in the concreteness score dataset, we cross-reference the EmojiNet senses with BabelNet which does provide lemmas. Not all emoji senses appear in the concreteness dataset, however, so we generate scores for 16 emoji in total, with results shown in Table 2. The mean concreteness score of these emoji is 4.59 ( $\sigma=0.54$ ), while the full concreteness dataset is 3.04 ( $\sigma=1.04$ ).

We evaluated the null hypothesis that the expected value of the sample of the semantically changed emoji is equal to the mean of the full dataset results, using a two-sided  $t$ -test. This resulted in a  $t$ -statistic of 13.52 and a  $p$ -value  $\ll 0.01$ , meaning that that we can reject the null hypothesis. In other words, semantically changed emoji seem to have a statistically significantly different (higher) average concreteness score compared to all emoji. As we discuss in section 8, a comparison between this result and the degree of semantic change of abstract vs. concrete words is an area of further research.

## 7 Case studies

We now examine in more detail how specific emoji have changed from 2012 to 2018. We select these to explore dif-

Emoji	Score	Emoji	Score	Emoji	Score
	5		4.83		4.48
	5		4.8		4.48
	5		4.77		4.26
	5		4.73		3.61
	5		4.57		3.52
	4.86				

Table 2: Concreteness scores of most semantically changed emoji, based on based on the concreteness ratings of Brysbaert, Warriner, and Kuperman (2014).

ferent patterns of seasonality ( ) and fads () in emoji semantics, as well as one case () of a dramatic shift in semantics.

### 7.1

The frog emoji was assigned to Cluster B in Section 5.2 on page 4. These emoji exhibit faddish behaviour, with a sudden intense peak of semantic change. In Figure 6 on the next page we show the semantic change for . According to Emojipedia<sup>5</sup>, it has been associated with several memes: “but that’s none of my business”<sup>6</sup> and “Pepe the frog”<sup>7</sup>, so we also show Google Trends data on some related terms: “trump” and “pepe the frog”. Both show increased interest around 2016/2017, but it does not appear that this was a direct cause of the overall pattern of change seen in —by then it was already very different from its initial version.

The most similar words in 2012 and 2013 were *lizard*, *bunny*, *frog*, but this changed in 2014. Most similar then were *kermit*, *none/nun*, *snitch*, *nvm*, *lowkey*. These are related to the “but that’s none of my business” meme, but this changed again from 2015 onward when the most similar words become more abusive: the *n*-word, *bitch*, *hate* all appear along with *pepe*. As of 2017, it was returning not to its 2012 animal-focused meaning, but to the 2014 meme-based meaning as well as taking on a new meaning related to sharing gossip and “spilling the tea”.

### 7.2

The skull emoji was assigned to Cluster A, the largest group and containing emoji that are generally undergoing gradual change. We selected this one for closer study because it is well known to have a figurative rather than literal meaning and is associated with dying of laughter or embarrassment<sup>8</sup>.

This change happened very quickly in 2013, as can be seen in Figure 7 on the following page. The most similar words in 2012 were *zombie*, *corpse*, *bury*, *undead*, *murder*.

<sup>5</sup><https://emojipedia.org/frog/>

<sup>6</sup><https://knowyourmeme.com/memes/but-thats-none-of-my-business>

<sup>7</sup><https://knowyourmeme.com/memes/pepe-the-frog>

<sup>8</sup><https://emojipedia.org/skull/>

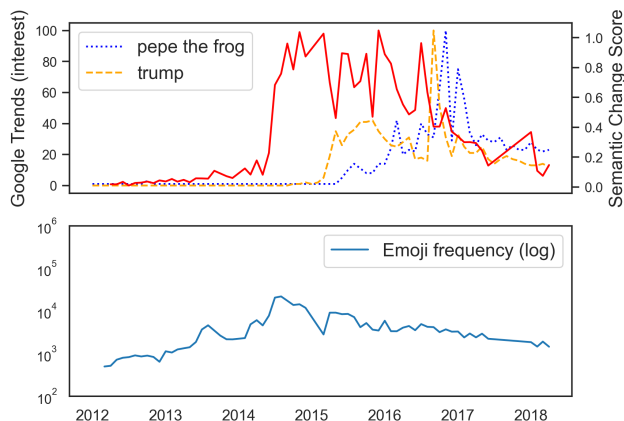


Figure 6: Top: Semantic change (raw, non-interpolated, non-znormed scores) of 🍄 (solid red line). Dashed lines show Google Trends data for “pepe the frog” and “trump”. Bottom: log token frequency per month.

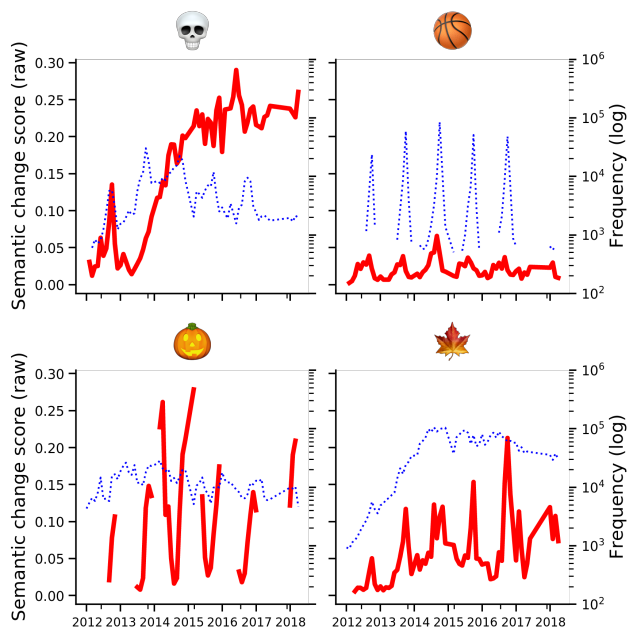


Figure 7: Semantic change patterns (raw, non-interpolated, non-znormed scores - red solid line) and log token frequency (dashed blue line) for four select emoji.

In 2013, these were joined by versions of *lmao* and entirely replaced by similar terms in 2014.

### 7.3 🏀

Another Cluster A emoji, this has undergone very little change on average but has seen a small temporary change each year so far (Figure 7). By examining the monthly most similar words over all years, rather than by the entire year, we observe that the change aligns with the NBA basketball season in the USA, which runs from October to April. Each year from August to October, NBA becomes the most similar term, along with NBA teams and players. At other times of the year, this emoji is used more broadly with terms such as *volleyball*, *soccer* and *softball*.

### 7.4 🎃

The jack-o-lantern emoji (Figure 7) is also in Cluster A. It is even more seasonal than basketball—so seasonal, in fact, that we have no data for it for half of each year. The nearest neighbours each month are almost all related to Halloween or other holidays, such as Christmas and Thanksgiving. Less similar are more general fantasy and horror terms. It is possible that these are the “background” semantics that we would observe if we had set a lower frequency threshold for this emoji.

### 7.5 🍂

The maple leaf emoji (Figure 7) is in Cluster D, a small group of emoji which have generally been stable when analysed in terms of their general semantic development. This is true of the maple leaf emoji, on average—it remains similar to its initial version, subject to some seasonal changes.

In this sense the maple leaf emoji is both very similar and very different to the basketball emoji. They both have a seasonal change in meaning but while basketball’s two modes are quite similar, the maple leaf’s modes are not. From August to November, its most similar words are *autumn*, *winter*, *chilly* and *pumpkin*. From December to July, it is used in connection with cannabis and its most similar terms are *weed*, *smoke*, *blunt* and *stoner*.

## 8 Discussion

Our analyses establish a general picture of semantic development in emoji, while case studies showed that semantic change can be linked to different kinds of seasonality, or to world events, or a shift from literal to figurative usage. The SC patterns identified here could be further developed into more specific sub-types and give an even more detailed understanding of emoji semantics.

One limitation of this work is that we only considered Unicode 6.x emoji, but it could be extended to others. Another is that our data collection ended in 2018, so we cannot determine the effect of the global pandemic on emoji semantics. For example, the 2020–2021 NBA season was delayed until December 2020—did this have an effect on the seasonality of the basketball emoji?

We expected abstract emoji to undergo more change than concrete emoji, but found the opposite. The visual affordances of emoji may account for this: 🌿 can denote marijuana because of shared visual similarities, in this case the shape of the leaves. The Lisbon Emoji and Emoticon Database (LEED) (Rodrigues et al. 2018) contains subjective norms (including visual complexity and concreteness) for some emoji which could be used to explore the connection between emoji appearance and semantic change. Because our embeddings were trained on English data and LEED norms are based on native Portuguese speakers only, we did not perform this analysis ourselves. However, Rodrigues et al. (2018) make their study materials available through the Open Science Framework, which should make it easy to extend LEED to additional languages.

As with words, only a minority of emoji exhibit semantic change. Therefore, a major extension of the work here would be to compare emoji semantic development to that of words. Semantic change in words is often described as an S-shaped curve (Feltgen, Fagard, and Nadal 2017). The semantic development of 💀 as shown in Figure 7 on the preceding page, from literal to figurative death, does seem S-shaped. Future work could explore the universality of semantic change.

Finally, although our analyses cover 6 years of data, it may be too soon to draw any firm conclusions on semantic development in emoji. It may require decades or centuries of additional data to determine the true nature of semantic change in emoji. However, as stated in the introduction, we are uniquely positioned to study this phenomenon from the very beginning.

## 9 Conclusion

In this paper, we presented the first longitudinal analysis of emoji semantics. Building on advances in NLP on representation learning and methods for measuring semantic change, we studied how the meaning of individual emoji evolved month by month in six years of Twitter data. Most emoji remained stable, and only a small number of emoji appeared to undergo substantive semantic change in our period of analysis. Our analyses also illustrated variation due to seasonal trends, and how emoji semantics is prone to trends and fads. We also found that semantically changed emoji seem to have higher average concreteness scores; future work should investigate the reason for this further.

Our study is a first step towards a richer understanding of emoji by taking into account their diachronic development. Our analyses open up many future avenues of exploration that could be further explored using our data, which we make publicly available through an interactive website<sup>9</sup> where the raw data can also be downloaded.

## Acknowledgements

Alexander Robertson was supported by the EPSRC Centre for Doctoral Training in Data Science, funded by the UK Engineering and Physical Sciences Research Council (grant

<sup>9</sup><https://emoji-semantic-change.herokuapp.com> - see Appendix A for more details.

EP/L016427/1) and the University of Edinburgh. Farhana Ferdousi Liza acknowledges the support of the Business and Local Government Data Research Centre (ES/S007156/1) funded by the Economic and Social Research Council (ESRC) for undertaking this work. Scott Hale was supported by the Volkswagen Foundation. This work was supported by The Alan Turing Institute under the EPSRC grant EP/N510129/1.

## References

- Basile, P.; and McGillivray, B. 2018. *Discovery Science*, volume 11198 of *Lecture Notes in Computer Science*, chapter Exploiting the Web for Semantic Change Detection. Springer-Verlag.
- Brybaert, M.; Warriner, A. B.; and Kuperman, V. 2014. Concreteness ratings for 40 thousand generally known English word lemmas. *Behavior Research Methods* 3(46): 904–911. doi:10.3758/s13428-013-0403-5.
- Chang, H.-C. 2010. A new perspective on Twitter hashtag use: Diffusion of innovation theory. *Proceedings of the American Society for Information Science and Technology* 47(1): 1–4. doi:<https://doi.org/10.1002/meet.14504701295>. URL <https://asistdl.onlinelibrary.wiley.com/doi/abs/10.1002/meet.14504701295>.
- Chen, Z.; Lu, X.; Shen, S.; Ai, W.; Liu, X.; and Mei, Q. 2017. Through a gender lens: An empirical study of emoji usage over large-scale android users. *arXiv preprint arXiv:1705.05546*.
- Cook, P.; Lau, J. H.; McCarthy, D.; and Baldwin, T. 2014. Novel word-sense identification. In *Proceedings of COLING 2014, the 25th International Conference on Computational Linguistics: Technical Papers*, 1624–1635.
- Dubossarsky, H.; Hengchen, S.; Tahmasebi, N.; and Schlechtweg, D. 2019. Time-Out: Temporal Referencing for Robust Modeling of Lexical Semantic Change. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, 457–470. Florence, Italy: Association for Computational Linguistics.
- Dubossarsky, H.; Weinshall, D.; and Grossman, E. 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*, 1136–1145.
- Feltgen, Q.; Fagard, B.; and Nadal, J.-P. 2017. Frequency patterns of semantic change: corpus-based evidence of a near-critical dynamics in language change. *Royal Society open science* 4(11): 170830.
- Frermann, L.; and Lapata, M. 2016. A Bayesian model of diachronic meaning change. *Transactions of the Association for Computational Linguistics* 4: 31–45.
- Garg, N.; Schiebinger, L.; Jurafsky, D.; and Zou, J. 2018. Word embeddings quantify 100 years of gender and ethnic stereotypes. *Proceedings of the National Academy of Sciences* 115(16): E3635–E3644.



- Hamilton, W. L.; Leskovec, J.; and Jurafsky, D. 2016a. Cultural Shift or Linguistic Drift? Comparing Two Computational Measures of Semantic Change. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, 2116–2121. Austin, Texas: Association for Computational Linguistics. doi:10.18653/v1/D16-1229. URL <https://www.aclweb.org/anthology/D16-1229>.
- Hamilton, W. L.; Leskovec, J.; and Jurafsky, D. 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, 1489–1501.
- Hassanpour, H.; Zehtabian, A.; and Sadati, S. 2012. Time domain signal enhancement based on an optimized singular vector denoising algorithm. *Digital Signal Processing* 22(5): 786 – 794. ISSN 1051-2004. doi:<https://doi.org/10.1016/j.dsp.2012.03.009>. URL <http://www.sciencedirect.com/science/article/pii/S1051200412000711>.
- Herring, S.; and Ge, J. 2020. Do Emoji Sequences Have a Preferred Word Order? In *Proceedings of the 3rd International Workshop on Emoji Understanding and Applications in Social Media*.
- Herring, S. C.; and Dainas, A. R. 2020. Gender and Age Influences on Interpretation of Emoji Functions. *ACM Transactions on Social Computing* 3(2): 1–26.
- Herring, S. C.; Dainas, A. R.; Long, H. L.; and Tang, Y. 2020. Anemoji Adoption and Use: Gender Associations with an Emergent Technology. In *Proceedings of the 3rd International Workshop on Emoji Understanding and Applications in Social Media*.
- Keogh, E.; and Kasetty, S. 2003. On the need for time series data mining benchmarks: a survey and empirical demonstration. *Data Mining and knowledge discovery* 7(4): 349–371.
- Kim, Y.; Chiu, Y.; Hanaki, K.; Hegde, D.; and Petrov, S. 2014. Temporal Analysis of Language through Neural Language Models. In *LTCSS@ACL*, 61–65. Association for Computational Linguistics.
- Kowald, D.; Pujari, S. C.; and Lex, E. 2017. Temporal Effects on Hashtag Reuse in Twitter: A Cognitive-Inspired Hashtag Recommendation Approach. In *Proceedings of the 26th International Conference on World Wide Web, WWW '17*, 1401–1410. Republic and Canton of Geneva, CHE: International World Wide Web Conferences Steering Committee. ISBN 9781450349130. doi:10.1145/3038912.3052605. URL <https://doi.org/10.1145/3038912.3052605>.
- Kulkarni, V.; Al-Rfou, R.; Perozzi, B.; and Skiena, S. 2015. Statistically significant detection of linguistic change. In *Proceedings of the 24th International Conference on World Wide Web*, 625–635. International World Wide Web Conferences Steering Committee.
- Kutuzov, A.; Velldal, E.; and Øvrelid, L. 2017. Tracing armed conflicts with diachronic word embedding models. In *Proceedings of the Events and Stories in the News Workshop*, 31–36. Vancouver, Canada: Association for Computational Linguistics. doi:10.18653/v1/W17-2705. URL <https://www.aclweb.org/anthology/W17-2705>.
- Kutuzov, A.; Øvrelid, L.; Szymanski, T.; and Velldal, E. 2018. Diachronic word embeddings and semantic shifts: A survey. In *Proceedings of the 27th International Conference on Computational Linguistics*, 1384–1397. Santa Fe, New Mexico, USA: Association for Computational Linguistics.
- Kwak, H.; Lee, C.; Park, H.; and Moon, S. 2010. What is Twitter, a Social Network or a News Media? In *Proceedings of the 19th International Conference on World Wide Web, WWW '10*, 591–600. New York, NY, USA: Association for Computing Machinery. ISBN 9781605587998. doi:10.1145/1772690.1772751. URL <https://doi.org/10.1145/1772690.1772751>.
- Lau, J. H.; Cook, P.; McCarthy, D.; Gella, S.; and Baldwin, T. 2014. Learning word sense distributions, detecting untested senses and identifying novel senses using topic models. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, volume 1, 259–270.
- Miller, H.; Kluver, D.; Thebault-Spieker, J.; Terveen, L.; and Hecht, B. 2017. Understanding emoji ambiguity in context: The role of text in emoji-related miscommunication. In *Eleventh international AAAI conference on web and social media*.
- Miller, H.; Thebault-Spieker, J.; Chang, S.; Johnson, I.; Terveen, L.; and Hecht, B. 2016. “Blissfully happy” or “ready to fight”: Varying Interpretations of Emoji. *Proceedings of ICWSM 2016*.
- Miller Hillberg, H.; Levonian, Z.; Kluver, D.; Terveen, L.; and Hecht, B. 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). doi:10.1145/3274393. URL <https://doi.org/10.1145/3274393>.
- Mitra, S.; Mitra, R.; Maity, S. K.; Riedl, M.; Biemann, C.; Goyal, P.; and Mukherjee, A. 2015. An automatic approach to identify word sense changes in text media across timescales. *Natural Language Engineering* 21(5): 773–798.
- Murtagh, F.; Downs, G.; and Contreras, P. 2008. Hierarchical clustering of massive, high dimensional data sets by exploiting ultrametric embedding. *SIAM Journal on Scientific Computing* 30(2): 707–730.
- Mylonas, P.; Wallace, M.; and Kollias, S. 2004. Using k-nearest neighbor and feature selection as an improvement to hierarchical clustering. In *Hellenic Conference on Artificial Intelligence*, 191–200. Springer.
- Naumann, D.; Frassinelli, D.; and Schulte im Walde, S. 2018. Quantitative Semantic Variation in the Contexts of Concrete and Abstract Words. In *Proceedings of the Seventh Joint Conference on Lexical and Computational Semantics*, 76–85. New Orleans, Louisiana: Association for Computational Linguistics. doi:10.18653/v1/S18-2008. URL <https://www.aclweb.org/anthology/S18-2008>.
- Navigli, R.; and Ponzetto, S. P. 2012. BabelNet: The Automatic Construction, Evaluation and Application of a Wide-Coverage Multilingual Semantic Network. *Artificial Intelligence* 193: 217–250.

Novak, P. K.; Smailović, J.; Sluban, B.; and Mozetič, I. 2015. Sentiment of emojis. *PLoS one* 10(12): e0144296.

Perrone, V.; Palma, M.; Hengchen, S.; Vatri, A.; Smith, J. Q.; and McGillivray, B. 2019. GASC: Genre-Aware Semantic Change for Ancient Greek. In *Proceedings of the 1st International Workshop on Computational Approaches to Historical Language Change*, 56–66. Florence, Italy: Association for Computational Linguistics. URL <https://www.aclweb.org/anthology/W19-4707>.

Robertson, A.; Magdy, W.; and Goldwater, S. 2018. Self-Representation on Twitter Using Emoji Skin Color Modifiers. In *Proceedings of the 12th International Conference on Web and Social Media, ICWSM 2018*. AAAI Press.

Robertson, A.; Magdy, W.; and Goldwater, S. 2020. Emoji Skin Tone Modifiers: Analyzing Variation in Usage on Social Media. *ACM Transactions on Social Computing* 3(2): 1–25.

Rodrigues, D.; Prada, M.; Gaspar, R.; Garrido, M. V.; and 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.

Salvador, S.; and Chan, P. 2007. Toward accurate dynamic time warping in linear time and space. *Intelligent Data Analysis* 11(5): 561–580.

Savitzky, A.; and Golay, M. J. E. 1964. Smoothing and Differentiation of Data by Simplified Least Squares Procedures. *Analytical Chemistry* 36(8): 1627–1639. doi:10.1021/ac60214a047. URL <https://doi.org/10.1021/ac60214a047>.

Schlechtweg, D.; McGillivray, B.; Hengchen, S.; Dubossarsky, H.; and Tahmasebi, N. 2020. SemEval-2020 Task 1: Unsupervised Lexical Semantic Change Detection. In *Proceedings of the 14th International Workshop on Semantic Evaluation*. Barcelona, Spain: Association for Computational Linguistics.

Shoemark, P.; Ferdousi Liza, F.; Nguyen, D.; Hale, S.; and McGillivray, B. 2019. Room to Glo: A Systematic Comparison of Semantic Change Detection Approaches with Word Embeddings. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, 66–76. Association for Computational Linguistics.

Tahmasebi, N. 2018. A Study on Word2Vec on a Historical Swedish Newspaper Corpus. In *CEUR Workshop Proceedings. Vol. 2084. Proceedings of the Digital Humanities in the Nordic Countries 3rd Conference, Helsinki Finland, March 7-9, 2018*. Helsinki: University of Helsinki, Faculty of Arts.

Tahmasebi, N.; Borin, L.; and Jatowt, A. 2018. Survey of Computational Approaches to Lexical Semantic Change. In *Preprint at ArXiv 2018*. URL <https://arxiv.org/abs/1811.06278>.

Tahmasebi, N.; and Risse, T. 2017. Finding Individual Word Sense Changes and their Delay in Appearance. In *Proceedings of the International Conference Recent Advances in Natural Language Processing, RANLP 2017*, 741–749.

Tigwell, G. W.; and Flatla, D. R. 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*, 859–866.

Wicke, P.; and Bolognesi, M. 2020. Emoji-based semantic representations for abstract and concrete concepts. *Cognitive Processing* 21: 615–635. doi:<https://doi.org/10.1007/s10339-020-00971-x>.

Wijeratne, S.; Balasuriya, L.; Sheth, A.; and Doran, D. 2017. EmojiNet: An Open Service and API for Emoji Sense Discovery. In *11th International AAAI Conference on Web and Social Media (ICWSM)*, 437–446. Montreal, Canada.

Wiseman, S.; and Gould, S. J. J. 2018. Repurposing Emoji for Personalised Communication: Why 🍷 Means “I Love You”. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems, CHI ’18*, 152:1–152:10. New York, NY, USA: ACM. ISBN 978-1-4503-5620-6. doi:10.1145/3173574.3173726. URL <http://doi.acm.org/10.1145/3173574.3173726>.

## Appendix A

The website <https://emoji-semantic-change.herokuapp.com> presents an interactive interface for exploring the data used in this paper. The monthly semantic change scores from 2012 to 2018 are available for all emoji and multiple emoji can be displayed and compared as timelines. In addition, for each emoji the most similar words are shown, aggregated by month or by year, to give a sense of what an emoji’s semantics are at different times. A Github link is provided, where all raw data can be downloaded for users who wish to perform their own analysis.

## Appendix B




Figure 8: Clustering of 348 Unicode 6.x emoji, based on semantic development from 2012 to 2018.