


Understanding Twitch Emotes in Sentiment Analysis

Originally published at <https://www.unite.ai/understanding-twitch-emotes-in-sentiment-analysis/>




The public's [growing use](#) of emojis, emoticons, emotes, memes, GIFs and other non-verbal ways to communicate on social media platforms has, in recent years, increasingly confounded the efforts of data scientists to understand the global sociological landscape; at least, to the extent that worldwide sociological trends can be discerned from public discourse.

Though Natural Language Processing (NLP) has become a powerful tool in sentiment analysis over the last decade, the sector has difficulty not only in keeping up with an [ever-evolving lexicon](#) of slang and linguistic shortcuts across multiple languages, but also in attempting to decode the meaning of *image-based* posts on social media platforms such as Facebook and Twitter.

Since the [limited number](#) of highly populous social media platforms are the only truly hyperscale resource for this kind of research, it's essential for the AI sector to at least attempt to maintain pace with it.

In July, a paper from Taiwan offered a [new method](#) to categorize user sentiment based on 'reaction GIFs' posted to social media threads (see image below), using a database of 30,000 tweets to develop a way to predict reactions to a post. The paper found that image-based responses are in many ways easier to gauge, since they are less likely to contain sarcasm, a [notable challenge](#) in sentiment analysis.



CLICK TO VIEW ANIMATED GIF

Researchers from Taiwan studied the use of animated reaction GIFs as 'reductive indicators' of sentiment in a 2021 paper.

Earlier this year, a research effort led by Boston University [trained machine learning models](#) to predict image memes that are likely to go viral on Twitter; and in August, British researchers examined the growth of emojis in comparison to emoticons (there's a [difference](#)) on social media, compiling a large-scale 7-language dataset of pictographic Twitter sentiment.

Twitch Emotes

Now, US researchers have developed a machine learning methodology to better understand, categorize and measure the ever-evolving pseudo-lexicon of *emotes* on the hugely popular Twitch network.

Emotes are neologisms used on Twitch to express emotion, mood, or in-jokes. Since they are by definition new expressions, the challenge for a machine learning system is not necessarily to endlessly catalogue new emotes (which may only be used once, or else fall out of usage rapidly), but to gain a better understanding of the framework that endlessly generates them; and to develop systems capable of recognizing an emote as a 'temporarily valid' word or compound phrase whose emotional/political temperature may need to be gauged entirely from context.

					
FeelsGoodMan		FeelsBadMan			
Token	Similarity	Token	Similarity	Token	Similarity
EZY	0.65	Sadge	0.896	Sadge	0.721
FeelsAmazingMan	0.647	PepeHands	0.896	PepeHands	0.681
Clap	0.634	Smoge	0.806	FeelsBadMan	0.68
EZ	0.624	:{(	0.806	:{(	0.643
FeelsOkayMan	0.616	:{(	0.78	sadge	0.638
widepeepoHappy	0.589	sadge	0.758	widepeepoSad	0.638
pepeW	0.58	peepoSad	0.758	peepoSad	0.631
=D	0.568	={(	0.736	t_t	0.603
Okayge	0.567	:{(	0.73	:{(	0.582
OkayChamp	0.561	:/	0.724	😞	0.579

Neighbors of the 'FeelsGoodMan' emote, whose meaning can be altered by obscure suffixes. Source: <https://arxiv.org/pdf/2108.08411.pdf>

The [paper](#) is titled *FeelsGoodMan: Inferring Semantics of Twitch Neologisms*, and comes from three researchers at Spiketrapp, a social media analysis company in San Francisco.

Bait and Switch

Despite their novelty and often-brief lives, Twitch emotes frequently recycle cultural material (including older emotes) in a way that can steer sentiment analysis frameworks in the wrong direction. Tracing the shift in the meaning of an emote as it evolves can even reveal a complete inversion or negation of its original sentiment or intent.

For instance, the researchers note that the original alt-right [misuse](#) of the eponymous *FeelsGoodMan* Pepe-the-frog meme has almost completely lost its original political flavor in the context of its usage on Twitch.

The use of the phrase, together with an image of a cartoon frog from a 2005 comic by artist Matt Furie, [became a far-right meme](#) in the 2010s. Though Vox [wrote](#) in 2017 that the right's appropriation of the meme had survived Furie's self-avowed [disassociation](#) with such use, the San Francisco researchers behind the new paper have found otherwise*:

'Furie's cartoon frog was adopted by rightwing posters on various online forums like 4chan in the early 2010s. Since then, Furie has campaigned to reclaim the meaning of his character, and the emote has seen an upsurge in more mainstream [non hate usage](#) and positive usage on Twitch. Our results on Twitch agree, showing that "FeelsGoodMan" and its counterpart "FeelsBadMan" are mainly being used literally.'

Trouble Downstream

This kind of 'bait and switch' regarding the generalized 'features' of a meme can impede NLP research projects that have already categorized it as 'hateful', 'right wing' or 'nationalist [US]', and which have dumped that information into long-term open source repositories. Later NLP projects may not choose to audit the older data's currency; may not have any practical mechanism to do

so; and may not even be aware of the need.

The upshot of this is that using 2017 Twitch-based datasets to formulate a 'political categorization' algorithm would attribute notable alt-right activity on Twitch, based on the frequency of the *FeelsGoodMan* emote. Twitch may or may not be [full of alt-right influencers](#), but, according to the researchers of the new paper, you can't prove it by the frog.

The 'Pepe' meme's political significance appears to have been casually discarded by Twitch's 140 million users (41% of whom [are under 24](#)), who have effectively re-stolen the work from the original thieves and painted it in their own colors, without any particular agenda.


Method and Data

The researchers found that labeled Twitch emote data was 'virtually non-existent', despite the conclusion of an [earlier study](#) that there are [eight million total emotes](#), and 400,000 were present in the single week of Twitch output in the week chosen by those earlier researchers.

A [2017 study](#) addressing emote prediction on Twitch limited itself to predicting only the top 30 Twitch emotes, scoring just 0.39 for emote prediction.

Addressing the shortfall, the San Francisco researchers took a new approach to the older data, splitting it 80/20 between training and testing, and applying 'traditional' machine learning methods, which had not been used before to study Twitch data. These methods included [Naive Bayes](#) (NB), [Random Forest](#) (RF), [Support Vector Machine](#) (SVM, with linear kernels), and [Logistic Regression](#).

This approach outperformed previous Twitch sentiment baselines by 63.8%, and enabled the researchers to subsequently develop the LOOVE (Learning Out Of Vocabulary Emotions) framework, which is able to identify neologisms and 'enrich' existing models with these new definitions.



Architecture of the LOOVE (Learning Out Of Vocabulary Emotions) framework developed by the researchers.

LOOVE facilitates the unsupervised training of word embeddings, and also accommodates periodic retraining and fine-tuning, obviating the need for labeled datasets, which would be logistically impractical, considering the scale of the task and the rapid evolution of emotes.

In the service of the project, the researchers [trained](#) an emote 'Pseudo-Dictionary' on an unlabeled Twitch dataset, in the process generating 444,714 embeddings of words, emotes, emojis and emoticons.

Further, they augmented a [VADER lexicon](#) with an [emoji/emoticon lexicon](#), and in addition to the aforementioned EC dataset, also exploited three other publicly available datasets for [ternary](#) sentiment classification, from Twitter, Rotten Tomatoes and a sampled YELP dataset.

Given the great variety of methodologies and datasets used in the study, the results are variegated, but the researchers assert that their best-case baseline outperformed the nearest prior metric by 7.36 percentage points.

The researchers consider that the ongoing value of the project is the development of LOOVE, based on word-to-vector (W2V) embeddings trained on over 313 million Twitch chat messages with the help of [K-Nearest Neighbor](#) (KNN).

The authors conclude:

'A driving feature behind the framework is a emote pseudo-dictionary which can be used to derive sentiment for unknown emotes. Using this emote pseudo-dictionary, we created a sentiment table for 22,507 emotes. This is the first case of emote understanding on this scale.'

** My conversion of inline citations to hyperlinks.*

Writer on machine learning, domain specialist in human image synthesis. Former head of research content at Metaphysic.ai, until its dissolution into DNEG's Brahma.ai.

Portfolio site: martinanderson.ai

Contact: martin@martinanderson.ai



Discover More