EMOJI-DRIVEN VISUAL SENTIMENT INTELLIGENCE FOR LOCAL GOVERNANCE COMMUNICATION AND CITIZEN ENGAGEMENT

NandaGopal G[1]*, V. S. Arulmurgan[2], Mohammadha Hussaini M[3]

[1]Research Scholar, Department of Information and Communication Engineering, Anna University, Tamil Nadu, India
[2]Professor & Head, Department of Electrical and Electronics Engineering, Shree Venkateshwara Hi-Tech Engineering College (Autonomous), Gobichettipalayam, Erode District, Tamil Nadu, India
[3]Associate Professor, Department of Electrical and Electronics Engineering, Government College of Engineering, Erode, Tamil Nadu, India

Corresponding email: gopal.nanda1981@gmail.com[1]
arulpvp@gmail.com[2]
hussaini1008@gmail.com[3]

Abstract - The rapid expansion of social media has significantly transformed communication practices between local governments and citizens, making digital symbols such as emojis an integral part of contemporary public discourse. Emojis serve as concise, language-independent visual cues that convey emotions, opinions, and attitudes, thereby enriching textual and visual communication in online civic spaces. In the context of local self-government, understanding emoji usage in citizen-generated content is increasingly important for interpreting public sentiment, enhancing participatory governance, and improving the responsiveness of local administrations. Despite notable advances in neural network models for text-based emoji prediction, forecasting emojis directly from images shared on digital platforms remains a challenging and underexplored task, particularly within governance-oriented communication analysis. This study proposes an integrated deep learning framework that combines Convolutional Neural Network (CNN)-based image classification with emoji2vec embeddings aligned with word2vec representations to predict emoji labels from images relevant to civic and local governance contexts. In addition, sentiment analysis of associated textual content is employed to improve the accuracy of future emoji prediction and to capture emotional patterns embedded in citizen–government interactions. The proposed approach effectively models semantic and sentiment relationships among emojis, enabling faster and more reliable prediction of image-based emoji usage. By optimizing the search duration and improving interpretability, the framework offers practical value for local authorities seeking data-driven insights into public perception, digital participation trends, and community feedback. Overall, the study contributes to the growing field of computational social analysis by demonstrating how visual sentiment intelligence can support evidence-based decision-making and communication strategies in local self-government.

Keywords: Local Governance, Emoji Prediction, Visual Sentiment Analysis, Citizen Engagement, Convolutional Neural Networks, Social Media Analytics

Introduction

Emojis like have grown in popularity on social networking sites like Twitter in recent years. In web conversations, emojis and text are frequently blended to convey various meanings and emotions. Exploring the connections between emoticons and text-based messages is a crucial endeavor. Emoji-enriched text production, social media sentiment analysis, online information retrieval, and other areas could all benefit from studying the links between emojis and text-based messaging. Studies on emojis often concentrate primarily on examining the semantics, usage, or sentiment of emojis. For instance, [1] investigated how emojis varied in meaning across languages. Emoji text descriptions should be included to improve the quality of emoji embedding, according to [2]. These methods, however, are unable to tap into the natural relationships between texts and emoticons. [3] presented a novel task to predict an emoji that is evoked by a simple tweet in an attempt to close this gap. They completed this job by using a hierarchical LSTM network. So, the performance might not be at its best.

Additionally, there are two issues with only guessing an original emoji for each message. First, people frequently convey their opinions within a message by using various emoticons. For instance, a tweet with the subject line "tired after the exam,[4] but the summer vacation comes!" makes use of several emojis to indicate various feelings. Second, popular emojis frequently evolve.[5-6] For instance, the emoji might be

used a lot around Christmas but very little else. Consequently, it is not ideal to use a fixed multi-class categorization model.
subsequently its origin, social media has expanded enormously and has subsequently developed into a venue for people to express their thoughts, ideas, and sentiments. Emojis were embraced by users to better convey the sentimental context of the text.[7] Emojis are essentially unique characters (or images) that are used to convey context in ways that regular text cannot. In electronic messages and on websites, emojis are ideograms and smileys. The origin of *emoji* derives from Japanese e (絵, picture) + moji (文字, character)[20]. Even though there is a great deal of linguistic difference, emojis, and their definitions are essentially the same in all the major languages. Emojis capture a more universally shared form of communication, particularly amongst cultures that are related.[8] Microblogging is a popular social media activity. These brief blogs have an emoji that expresses the text's emotions together with a small amount of content. As a result, we construct a broad association between the text and the emoji, with the emoji serving as the text's microblog tag.

We claim that sentences with similar related emojis in various languages have semantic properties that are similar using this element of emojis. Despite the prevalence of videos on social media platforms like Facebook and Twitter, text still predominates in communication. Emojis can offer richer expressions to help overcome the lack of non-verbal clues that can occur in text-based communication.[9-10] A collection of reserved characters known as emojis are displayed as tiny pictograms that represent various facial expressions. Sarcasm is the subtle style of language that people use to express the opposite of what is conveyed in social media. The job of detecting sarcasm is crucial to raising the standard of online communication. [11]It first aids in our comprehension of the feedback from users' true intentions. User reviews could include statements like "Wow, this product is great," "It is very fast," "Totally worth it," etc. However, these remarks are being made in a satirical manner. Second, satirical posts could affect how individuals feel and respond to the election campaign.
Emojis are now used more frequently than text in social media texts such as blogs, microblogs (e.g.: Twitter), chats (e.g.: WhatsApp and Facebook), and more [12]. An angry face Emoji can convey negative emotions, whilst a happy face Emoji can convey pleasant emotions.[13] Based on the emoji characters chosen, one can clearly understand the emotion expressed in the text.
Social network users are continuously growing in number. Users publish data on a range of subjects. Researchers, businesses, politicians, and public figures can all benefit from the data by better-understanding consumer or market opinion and making the necessary adjustments to increase their appeal. One of the most well-known social networking services is Twitter. With over 500 million tweets posted every day, Twitter usage is growing constantly. The use of tweet sentiment analysis is expanding in various fields, including recommendation and decision support systems. Therefore, it has become crucial and a focus for many researchers to increase the precision of tweet sentiment analysis. By utilizing simply textual data, many studies have attempted to increase the effectiveness of tweet sentiment analysis techniques. [14-15] The Twitter search API was used to gather the H4EAD data over one week. The "saveh4ead" hashtag was used to gather the supportive tweets, and the "noh4ead" hashtag was used to gather the critical tweets. There were 2252 tweets, 246 of which were critical, and 2006 of which were supportive. The outcomes of our trials using the Kaggle dataset, which contains both positive and negative emojis, are pretty encouraging.

METHODS

Emoji Prediction

Emoji usage is rapidly expanding, but there is a need to make them more accessible. There aren't any prediction or classification tools or data available for emoji submission, but there are for non-emoji (text) content. There have been several approaches suggested for similarity modeling and skip-gram models to predict future emoji entries from text, such as tweets and other types of text, but not from photos. The project involves predicting the label for the upcoming emoji addition using photos. The method suggests

W2V for word vector representations and CNN for image classification. Additionally, this technique uses Euclidian distance to compare emoji similarity and Cosine similarity to represent similarity between photos. T-SNE plots show how similar emojis are to one another. Additionally, text sentiment analysis is done to use RNN to forecast future emoji inputs.

Table 1: Emoji usage in different language tweets

Language	😕	😍	😇	😑	😣	😎	😡	😆	😊	🙂	😒	❤️	💞	😓	😴	🤔	💕	😶
Telugu	22.7	0.3	3.9	10.8	5.7	0.5	1.6	0.9	13.4	5.1	3.0	6.7	0.9	2.3	4.2	0.3	16.6	1.4
Hindi	9.7	2.8	4.9	9.5	6.8	2.6	2.9	4.3	9.4	8.6	6.7	7.7	2.3	1.9	7.2	1.5	7.5	3.9
Spanish	9.7	2.7	6.3	6.5	10.8	2.0	3.4	6.1	4.0	4.6	6.0	5.4	3.4	2.8	3.4	6.8	13.1	3.1
English	17.1	4.9	3.8	4.7	15.3	2.6	3.5	3.3	3.2	3.1	3.6	3.0	2.7	2.7	2.8	3.7	10.0	5.7

Emojis have developed into a vital tool for facilitating communication and emotional expression. Attention is drawn to the study of emojis about text comprehension and sentiment analysis. By adding emoticon signals, it suggests an unsupervised paradigm for sentiment categorization. A Multinomial Naive Bayes Classifier is used to test the viability of an emoji training heuristic for multi-class sentiment analysis on Twitter.

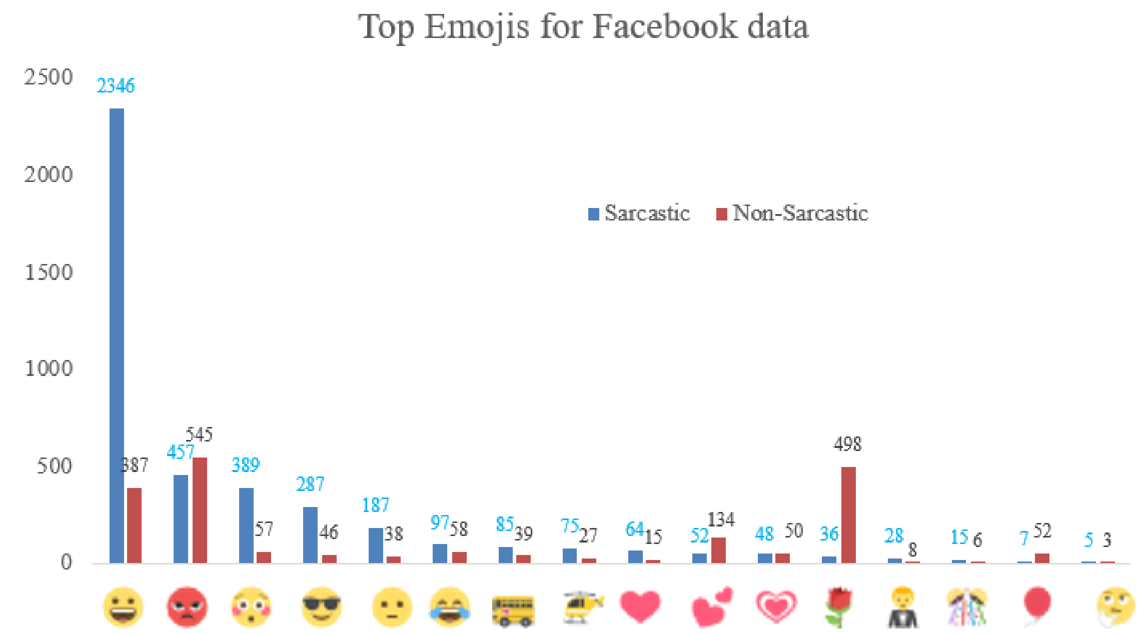

Figure 1. Comparison of top Emojis for Facebook data.

By applying skip grammar to the Unicode standard's definitions of the emojis, it learns how to represent them. Emoji representation can enhance sentiment analysis, according to the results of a sentiment analysis that uses the generated emoji representation with word embedding from Google News. This demonstrates how emojis can be appropriated to facilitate communication while using interview data. The entire architecture of this proposal is seen in Figure 2.

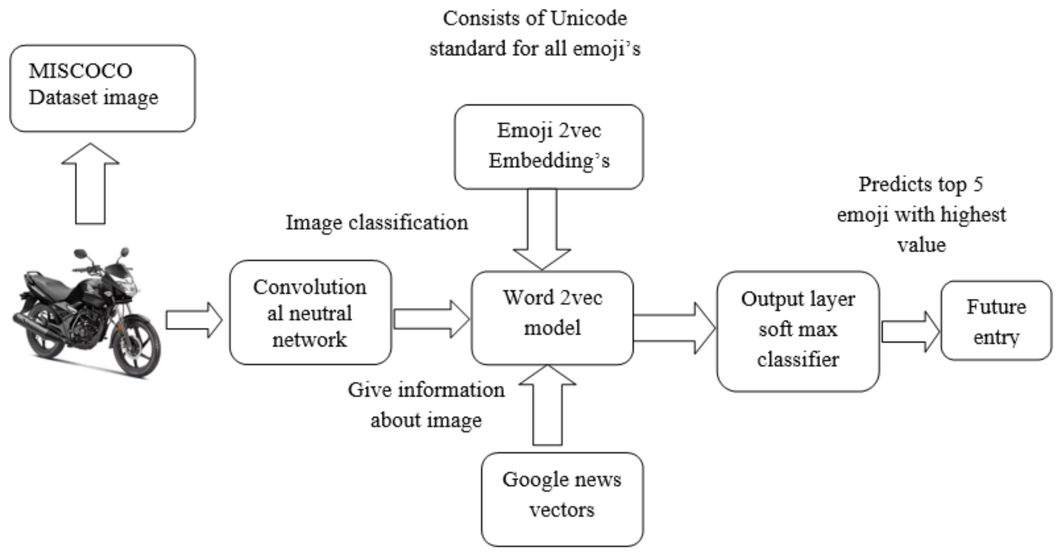

Figure 2: a general framework for predicting future emoji entries from photos

ConvNets, a type of convolutional NN architecture, uses COCO images as its input layer. Neurons that makeup convents have biases and weights that can be learned. In contrast to a standard NN, ConvNet layers include neurons arranged in three dimensions: width, height, and depth. The convolutional Layer, Normalisation Layer, Pooling Layer, and ReLU Layer are the four crucial levels of the ConvNet architecture, as depicted in Figure 2. The Max Pooling layer is the output layer in CNN that provides descriptions of images about image IDs from the COCO dataset. The last instance of a sequential layer in this model is an average pooling layer rather than a maximum pooling layer, and each instance of a sequential layer has four primary CNN layers. To create a feature map from the inputs, the convolutional layer has a convolutional filter. All values in the input volume are subject to a maximum function by the ReLU layer. The largest value from each neuronal cluster at the preceding layer is used by the max pooling layer. The W2V model receives the CNN-obtained picture descriptions as an input. A three-layer neural network model called the Word2Vec has one hidden layer. Word embeddings are created for the data by W2V, which teaches a linguistic context of words.

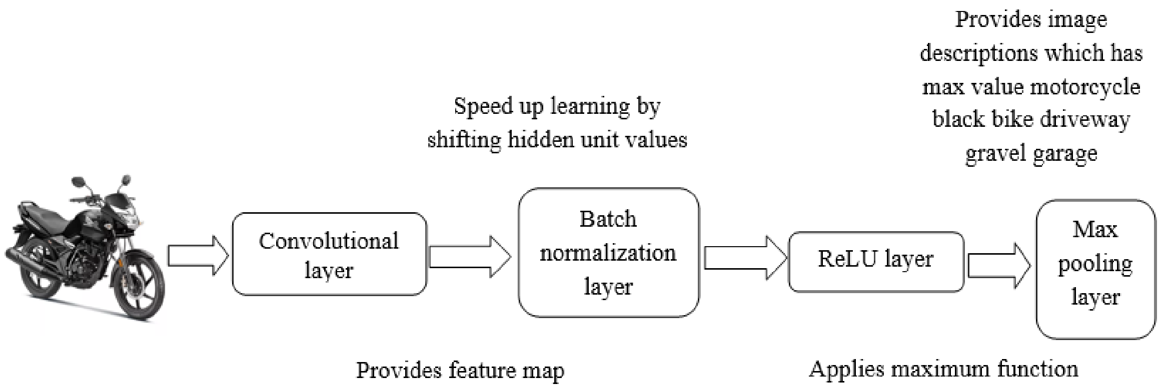

Figure 3: Convolutional neural network

Smiley Embedding Network

Deep neural network architectures can be utilized to efficiently learn the emoji embedding with lower risks of overfitting the data given the quantity of the obtained dataset. Learning an embedding function, f(), that converts an image, x Xdx, to an embedding, e Ede, or f: Xdx Ede, is our formal goal. so that dx and de are the dimensions of the image and emoji spaces, respectively. It is practical to implement f() by using the proxy task of explicit emoji prediction (Figure 3a). This offers two major advantages over previous methods, such as metric learning in the emoji domain. As a result, it is now able to assess the embedding and develop a brand-new zero-shot challenge for visual sentiment learning.

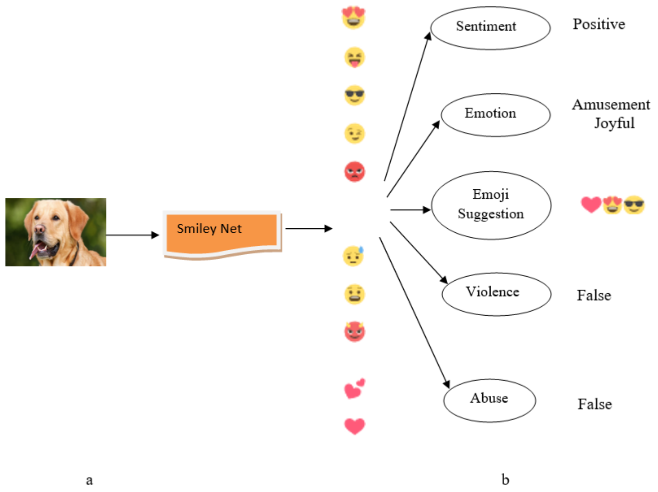

Figure 4: From large-scale and chaotic social media data, our model (SmileyNet) (a) learns how to include photos in emoji's low-dimensional space. Transfer learning (b) can then be used to take advantage of this embedding for a variety of target tasks where it is necessary to derive emotions from visual data, including sentiment and emotion analysis.

The information in RNN loops back on itself. Making decisions is based on both the present input and the memory material that has been acquired from previous inputs. GloVe vector embeddings are employed in our model to train the text sentiment analysis on the textual data. GloVe is a distributed word representation approach that is derived from Global Vectors. The input data are used to train the network model together with glove embeddings. RNN is made up of a simple RNN layer, which is followed by layers for density, dropout, and activation. The output dimensionality unit of the space is created by the simple RNN layer using the embedding matrix of the input data. The network's representational capacity is increased by the dense layer. Overfitting of the data is prevented by the dropout layer. Using softmax activation, the activation layer generates probabilities for each word in the vocabulary. The words with the highest likelihood are output in Python and given an emoji name by mapping the output word to the relevant Unicode character. Figure 5 depicts the RNN architecture for the text's sentiment analysis.

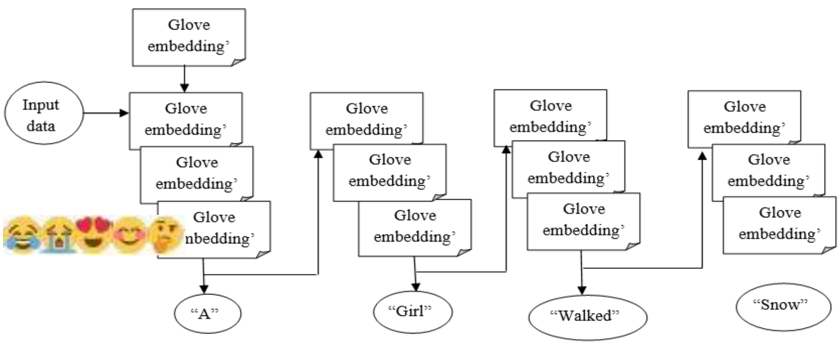

Figure 5: Recurrent neural network

Sample Distribution

Several social media platforms, like Instagram, Flickr, and Twitter, provide a significant source of extensive emoji data. Emojis are thought to be sent more than 700 million times every day on Facebook, and they are used in about half of Instagram posts. Here, we choose our Twitter samples in a way that only includes tweets with emojis and tweets that are linked to at least one image. The motivation for this came from the observation that these elements frequently function as important context cues that go beyond the visual information associated with the chosen emoji. Using the aforementioned criteria, the 2.8 million Tweets from the first half of 2018 were located. The label distribution of the data is displayed in Figure 6a (). We can observe that this data is significantly biased towards a small number of categories and has a long tail distribution, with the top 5 most popular emojis accounting for about 40% of the samples that were collected. This is a significant challenge for the majority of conventional machine learning techniques since an unbalanced training dataset may cause the learning process to simply predict the labels that occur most frequently rather than discovering a more insightful representation. Additionally, we see that the content of samples tends to be highly biased towards a small number of significant temporal events when data are collected over a short period. As a result, there is less variation in the photos, which improves the model's capacity to generalize effectively across domains.

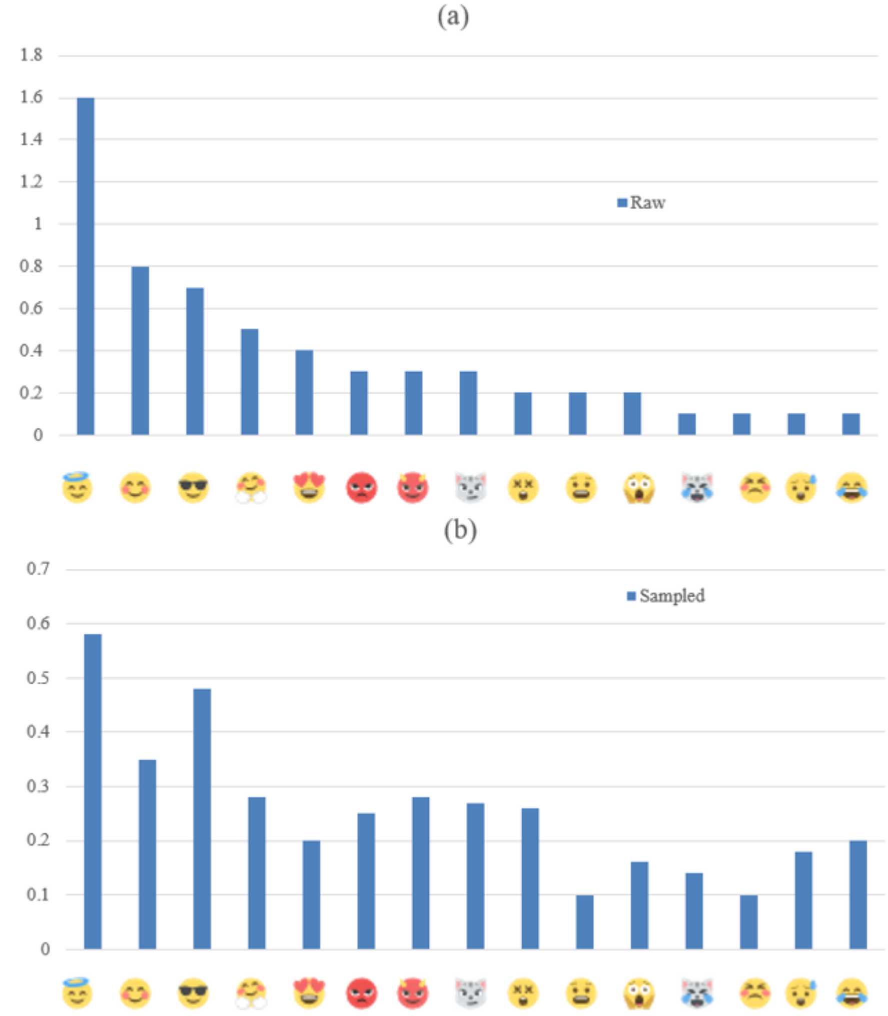

Figure 6: Emoji frequency in the (a) temporally balanced sampled dataset and (b) a raw sample of data. The study makes use of dataset (b).

Image Preprocessing

Because irrelevant data, noisy data, and unreliable data can produce false findings during the training phase of data mining, data preparation is a crucial component of the process. Therefore, preprocessing the data is required before training. Every image in the data collection had to be resized to be the same size and dimensions due to the inconsistent dimensionality of the photos. The Python Imaging Library (PIL) uses ANTIALIAS to automatically resize every picture to 256256 pixels.

Dataset Used

For picture inputs, we used the MSCOCO 2017 dataset, which includes 15k training images and 5k testing images (cocodataset.org). Word vector representations are made using Google-News-vectors (GoogleNews-vectors-negative300.bin.gz).

Temporal Sampling

This suggests retrieving the data from a comparatively long period while uniformly sampling the tweets from shorter temporal windows to combat content uniformity. In particular, we gather tweets made between January 1st, 2016, and July 31, 2018. We divided the period into 30-day time frames that were sequential. For additional relief from label imbalance, we choose up to 4,000 tweets at random from each window's emoji categories. Additionally, we permit valid samples to have up to 5 emojis, so some samples will contain multiple labels. This process produced 5.2 million emoji labels and around 4 million photos in total. The label distribution of the sampled dataset is depicted in Figure 6b. As we can see, our dataset is more evenly distributed throughout the different categories than the raw data distribution. We create a normalized correlation matrix of all emojis in the gathered data2 to have a better understanding of the association between labels. The two most popular emojis, and, co-occur with the majority of the categories, as can be seen by looking at the correlation matrix. The correlation matrix also identifies other semantically similar categories, including [] and [].

W2V Model

To forecast future emoji entries, Word2Vec, another neural network model, receives information about the image from CNN. Emoji2vec, an emoji vector space model, is employed for this assignment in place of words. Emoji embedding E2V uses the uniform Unicode representation of every emoji. W2V, which is based on E2V, offered a semantically effective paradigm for connecting text from CNN to emoji characters. W2V was created with the help of Gensim and NLTK. The word vectors utilized for training are the three million 300-dimension word vectors from Google News. For each piece of information (e.g., noun, adjective, verb, etc.) that NLTK gleaned from CNN about the provided image, W2V uses a Part-of-speech tag. It also eliminates information that is unnecessary and unreliable. Because these words—rather than all other words in sentences—constitute the necessary information about a picture, NLTK chooses the proper nouns, adjectives, and verbs. Gensim uses word2vec to vectorize the remaining words in the textual data. Five emojis that are comparable to or closest to each word in the textual data from CNN are obtained by W2V for each word. W2V is a three-layer NN model with a word from CNN's textual data as the input layer. The word vectors in the hidden layer are organized in a matrix with hidden weights, with each distinct word in the Google News dataset as a column. A multi-class classifier called a Soft max regression classifier serves as the output layer.

CNN Model

After being preprocessed into 256256 pixels, resized images are submitted to the CNN model. A particular kind of neural network called a CNN is composed of neurons with biases and weights that can be taught. CNN is an image classifier that has many different labels and classifications. In terms of implementation, Pytorch was used to generate CNN. Each sequential layer that makes up the CNN network design contains the max-pooling layer, the ReLU activation layer, the convolutional layer, and the batch normalization layer. The convolution filter is used by the CNN Convolutional layer to turn the input images into a feature map. Sliding the filter across the input at each pixel point computes the element-wise matrix multiplication and total. The total is shown on the feature map. The inputs from the feature map are normalized in the following layer, called Batch Normalisation, by shifting the hidden unit values, taking the batch mean into account, and dividing by the standard deviation. The training and learning process is accelerated by this. The ReLU layer's maximum function is applied to each value in the input volume. Every negative activation is changed to a value of 0. This boosts the network's overall nonlinear properties as well as the model's nonlinear properties. A series of layers follow each succeeding layer.

EXPERIMENTAL RESULTS

Every image was preprocessed to get rid of any non-uniform dimensions. CNN finds it challenging to analyze photographs of various sizes. The default dimension for every image used for prediction is therefore 256256. The scaled image is then submitted to CNN for image classification following input image preprocessing. CNN uses training data to categorize the image and provides detailed textual information. Euclidian distance is used to determine how similar two emojis are to one another and to locate emojis that are similar to one another. This aids in accurately predicting and analyzing the upcoming emoji labels. Table 2 provides an example of the results of the emoji Euclidian distance measurement.

Table 2: Euclidian distance measure of emojis

Emoji 1	Emoji 2	Euclidean measure
	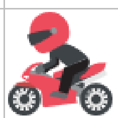	5.9348
		6.2304

	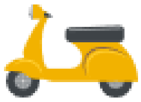	3.0394
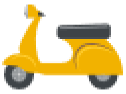		8.0934
	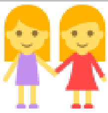	5.9283
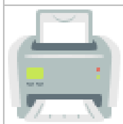	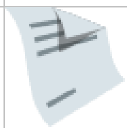	9.2323
		12.4958
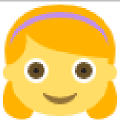		7.2352

The TSNE plot is used to visualize emojis to examine the relationships between them and identify emoji clusters based on a similarity metric. The local structure is very well captured by TSNE, and the clustered data is likewise very distinct. The Cosine Similarity measure is used to determine how similar two photographs are. Because the cosine similarity measure uses categories rather than just two variables, it is employed. Table 2 provides the cosine similarity between the photos. This model forecasts future emoji input labels by examining the text's sentiment or emotions. The entire written information is divided into five categories, including joyful, in love, furious, sports, and eating.

Table 3: Cosine similarity between images

Image 1	Image 2	Score 1	Score 2	Cosine measure
		21.984489 14.239480 11.349892 23.234121 3.029348	6.3509234 5.3059832 7.3458923 11.345903 5.2341155	0.94
		2.934857 3.938567 8.029348 9.938572 21.230948	6.233094 6.123311 5.231522 9.039863 2.213454	0.98

		3.325896 5.123532 6.345311 8.453634 14.645454	8.675456 8.654542 8.543436 8.543557 8.324568	0.95
		8.765769 2.653543 1.4354687 2.877476 6.543565	4.214235 3.2315342 1.243534 9.23523 8.231534	0.93
		3.675675 6.675798 8.754539 3.351231 9.435234	11.235457 7.235457 1.346543 2.347567 4.325798	0.92

Figure 7 displays a graphic representation of the degree of resemblance between the two photos. photos that are similar to one another display a low distance value, whereas photos that are dissimilar display a large difference in the similarity values. Recurrent Neural Network is used to forecast future emoji entries via sentiment analysis. This model keeps both the previous and current outputs in its memory. The network can learn long-term relationships in a sequence thanks to its memory, enabling it to anticipate the next word in a sentence or classify sentiments while taking into account the complete context.

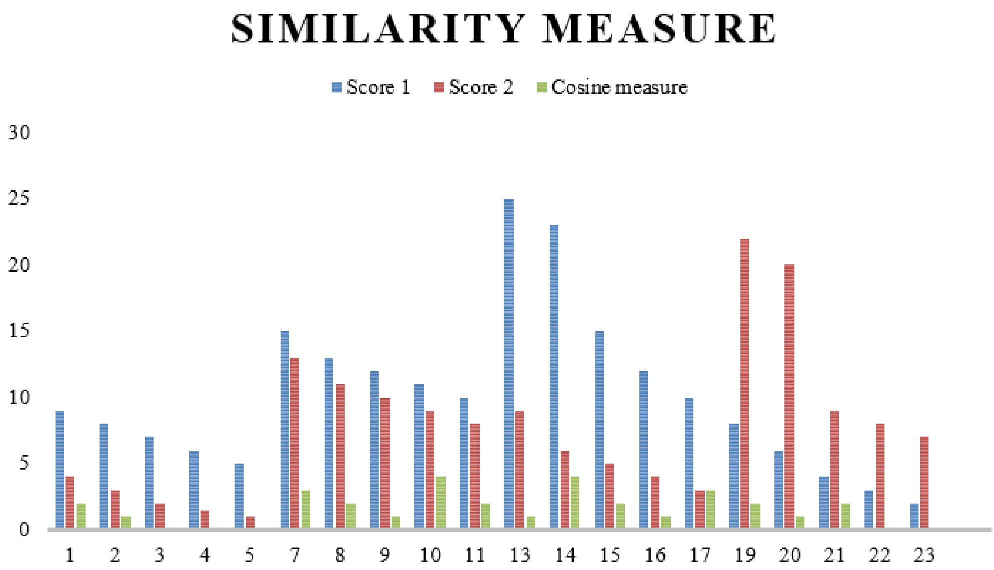

Figure 7: Comparison of similarity measure

Emojis and feeling

The "5 agrees" split images are all automatically integrated into the emoji area using our SmileyNet. Top SmileyNet predictions for certain test photographs from Twitter are displayed in Figure 8. Our method generates accurate predictions that reflect the tone of the image as a whole. When the input image consists of sub-images or when the main sentiment region is out of focus, our algorithm suggests emojis of sentiment with various polarities. This might be connected to the all-encompassing SmileyNet strategy. The idea is that processing based on region or attention may help prioritize the most important visual area for outcome predictions. Last but not least, predictions may be beneficial for both conventional applications like sentiment analysis and cutting-edge ones like identifying violence or abuse in images.

Figure 8: Top 5 emojis predicted per image with SmileyNet

Table 4: Top emojis that correspond to each emotion category.

Emotion	Most correlated Emojis						
Excitement	0.23	0.34	0.12	0.18	0.10	0.34	0.45
Contentment	0.34	0.18	0.18	0.19	0.13	0.71	0.23
sadness	0.93	0.92	0.87	0.78	0.98	0.23	0.56
Anger	0.44	0.67	0.78	0.78	0.89	0.56	0.67
Fear	0.12	0.11	0.22	0.67	0.16	0.67	0.78
Disgust	0.56	0.78	0.78	0.67	0.22	0.87	0.89
Amusement	0.34	0.67	0.33	0.98	0.17	0.11	0.14
Awe	0.45	0.49	0.23	0.67	0.34	0.56	0.21

CONCLUSION

To forecast upcoming emoji inputs from photos, two models are built. Emoji and textual information are both employed in media to represent information about images. Therefore, when compared to text, photos can be utilized in place of text to effectively forecast future emoji entries. Thus, by predicting the next entry from photos in the shortest amount of time possible and optimizing speed and performance, this model aids us. CNN, SmileyNet, and Word2Vec are the three models used to create the Emoji Model. Additionally, we show that our embedding may be employed for sentiment analysis in a zero-shot learning environment without any additional training due to its interpretability. Last but not least, preliminary findings demonstrate that our embedding can support emerging applications where deriving emotion from visual data is crucial, such as visual abuse and violence detection. We anticipate that the outcomes of this study will be relevant to the fields of social media research and emoji modality comprehension in addition to the computer vision and visual sentiment analysis communities.

Reference

1. Ramaswamy, S., Mathews, R., Rao, K., & Beaufays, F. *Federated learning for emoji prediction in a mobile keyboard.* arXiv preprint arXiv:1906.04329, 2019.
2. Tavan, E., Rahmati, A., & Keyvanrad, M. A. *Persian emoji prediction using deep learning and emoji embedding.* In *2020 10th International Conference on Computer and Knowledge Engineering (ICCKE)*, pp. 350–355. IEEE, 2020.
3. Lee, S., Jeong, D., & Park, E. *MultiEmo: Multi-task framework for emoji prediction. Knowledge-Based Systems*, 242: 108437, 2022.
4. Wu, C., Wu, F., Wu, S., Huang, Y., & Xie, X. *Tweet emoji prediction using hierarchical model with attention.* In *Proceedings of the 2018 ACM International Joint Conference and 2018 International Symposium on Pervasive and Ubiquitous Computing and Wearable Computers*, pp. 1337–1344, 2018.
5. Tomihira, T., Otsuka, A., Yamashita, A., & Satoh, T. *What does your tweet emotion mean? Neural emoji prediction for sentiment analysis.* In *Proceedings of the 20th International Conference on Information Integration and Web-Based Applications & Services*, pp. 289–296, 2018.
6. Barbieri, F., Ballesteros, M., Ronzano, F., & Saggion, H. *Multimodal emoji prediction.* arXiv preprint arXiv:1803.02392, 2018.
7. Peng, D., & Zhao, H. *Seq2Emoji: A hybrid sequence generation model for short text emoji prediction. Knowledge-Based Systems*, 214: 106727, 2021.
8. Raj, H. W., & Balachandran, S. *Future emoji entry prediction using neural networks. Journal of Computer Science*, 16(2): 150–157, 2020.
9. Coman, A. C., Zara, G., Nechaev, Y., Barlacchi, G., & Moschitti, A. *Exploiting deep neural networks for tweet-based emoji prediction.* In *International Workshop on Semantic Evaluation*, vol. 4, p. 1, 2018.
10. Barbieri, F., Espinosa Anke, L., Camacho-Collados, J., Schockaert, S., & Saggion, H. *Interpretable emoji prediction via label-wise attention LSTMs.* In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pp. 4766–4771, 2018.
11. Akter, M., Hossain, M. S., & Andersson, K. *Hand-drawn emoji recognition using convolutional neural network.* In *2020 IEEE International Women in Engineering (WIE) Conference on Electrical and Computer Engineering (WIECON-ECE)*, pp. 147–152. IEEE, 2020.
12. Choudhary, N., Singh, R., Bindlish, I., & Shrivastava, M. *Contrastive learning of emoji-based representations for resource-poor languages.* In *International Conference on Computational Linguistics and Intelligent Text Processing*, pp. 129–141. Cham: Springer Nature Switzerland, 2018.
13. Matsumoto, K., Yoshida, M., & Kita, K. *Classification of emoji categories from tweet based on deep neural networks.* In *Proceedings of the 2nd International Conference on Natural Language Processing and Information Retrieval*, pp. 17–25, 2018.

14. Golazizian, P., Sabeti, B., Ashrafi Asli, S. A., Majdabadi, Z., Momenzadeh, O., & Fahmi, R. *Irony detection in Persian language: A transfer learning approach using emoji prediction.* In *Proceedings of the Twelfth Language Resources and Evaluation Conference*, pp. 2839–2845, 2020.
15. Duarte, L., Macedo, L., & Oliveira, H. *Emoji prediction for Portuguese.* In *International Conference on Computational Processing of the Portuguese Language*, pp. 174–183. Cham: Springer International Publishing, 2020.