

Understanding Global Reaction to the Recent Outbreaks of COVID-19: Insights from Instagram Data Analysis

Abdul Muntakim Rafi¹, Shivang Rana², Rajwinder Kaur³, Q.M. Jonathan Wu⁴, and Pooya Moradian Zadeh⁵

^{1,2,4}Department of Electrical & Computer Engineering,

^{3,5} School of Computer Science

University of Windsor, Ontario, Canada

{rafi11, rana11r, kaur19s, jwu, pooya}@uwindsor.ca

Abstract—The coronavirus disease, also known as the COVID-19, is an ongoing pandemic of a severe acute respiratory syndrome. The pandemic has led to the cancellation of many religious, political, and cultural events around the world. A huge number of people have been stuck within their homes because of unprecedented lockdown measures taken globally. This paper examines the reaction of individuals to the virus outbreak through the analytical lens of specific hashtags on the Instagram platform. The Instagram posts are analyzed in an attempt to surface commonalities in the way that individuals use visual social media when reacting to this crisis. After collecting the data, the posts containing the location data are selected. A portion of these data are chosen randomly and are categorized into five different categories. We perform several manual analyses to get insights into our collected dataset. Afterward, we use the ResNet-50 convolutional neural network for classifying the images associated with the posts, and attention-based LSTM networks for performing the caption classification. This paper discovers a range of emerging norms on social media in global crisis moments. The obtained results indicate that our proposed methodology can be used to automate the sentiment analysis of mass people using Instagram data.

I. INTRODUCTION

Advent of social media presents a promising opportunity for emotion analysis of the mass people. Different approaches have successfully analyzed social media data to detect several conditions [1]. There is a vast wealth of psychological data encoded in visual social media data, such as photographs posted to Instagram. In this paper, we explore the possibility of analyzing visual data from Instagram to understand the global reaction to the recent outbreak of the coronavirus.

It is reasonable to prioritize research into Instagram analysis for this problem. Instagram members currently contribute almost 100 million new posts per day [2], and Instagram's rate of new users joining has recently outpaced Twitter, YouTube, LinkedIn, and even Facebook [3]. Users give more importance to the visual content while sharing posts on Instagram. The posts are accompanied by captions as well. Instagram, among all social media, is more of personal space for the users. People tend to be less formal in their reactions. The data shared here can often have hidden complex meanings, which makes the Instagram posts most suitable for our

task. Instagram data has proved to be a useful source for emotion analysis of the mass people. Instagram tends to be a place where people share personal emotions [4].

According to [5], the origin of the outbreak has been identified as Wuhan. The virus rapidly spread all over the globe. As of May 2020, more than 4 million cases of COVID-19 is confirmed worldwide [6], [7]. World Health Organisation (WHO) has declared a Global Health emergency, as the number of this novel virus keeps ascending globally [8]. Worldwide 212 countries or territories have confirmed cases of this novel virus [6]. A lot of everyday activities around the world have been restricted except for necessity and health circumstances. Acknowledging the rapid spread of the novel virus, there is a considerable impact of this epidemic on social networks and media as well.

Understanding public reactions to a crisis is the first step to solving the situation. The initial assessment includes how the crisis is affecting people's lives. A concern that disrupts the everyday lives of more people is a more severe problem. Therefore, successfully identifying human reactions to a situation can lead to a successful estimation of the seriousness of the problem. Also, the mental health of individuals may deteriorate during crisis moments. As Instagram is a platform where people share their emotions, it is one of the best places to analyze the mental health of mass groups of people. On average, 100 million photos are uploaded daily on Instagram [9]. Manual analysis of all the information shared on a social media platform is almost impossible. Recently, deep learning algorithms have performed magnificently at extracting features from visual data. Therefore, we try to explore the feasibility of deep learning algorithms in analyzing Instagram data in this paper.

The major contributions of our paper are as follows:

- 1) We collect a dataset consisting of 11973 Instagram posts with location.
- 2) We manually annotate 2534 posts for getting insight into our collected Instagram data.
- 3) We explore the performance of deep learning algorithms in performing the analysis without human intervention.

The rest of the paper is organized as follows. Section II describes the related researches conducted on the literature. Section III explains the deep learning terminologies used

*The first 3 authors had equal contribution to the paper. The dataset is available at www.kaggle.com/muntakimrafi/global-reaction-to-covid19-instagram-data-analysis

throughout the paper. Section IV explains our proposed methodology in details, and section V contains all of our experiments and results. Finally, we present our conclusion at Section VI.

II. LITERATURE REVIEW

Researchers have found that mental health is strongly associated with social activities [10]. Photographs posted to Instagram contain different dimensions that might be analyzed for psychological insight. The photographs can be analyzed to determine different characteristics, such as the presence of people in the photographs or videos, the setting in nature or indoors, the time of the day, etc. Some research have been conducted in extracting statistical properties from images for performing various types of analysis [11]. Instagram metadata offers additional information like how many comments or likes the post received. Finally, platform activity measures of different areas can yield clues to Instagram user's mental state. These dynamics were briefly discussed in [12]. In particular, Lup et al. [13] and Andalibi et al. [14] proposed to use Instagram data for mental health analysis. Recently, some works in the literature started using multi-modal analysis. In [15], the authors showed that using both textual and visual analysis allowed them to distinguish between different types of users and their interests.

Instagram posts have been used to find out the emerging norms in social network dynamics. In [4], researchers have used specific hashtags to pull off data from Instagram and discovered the nature of the reaction of general people to birth and death. They have raised a crucial privacy concern by showing that people create the first digital footprint of their children even before they are born. They also showed that people share their emotions on Instagram. In [16], the authors showed that Instagram plays a vital role in body dissatisfaction and women's drive for thinness. In their research, it has been shown how people's lives are affected by the posts they see on social media.

Previously, the manual or statistical feature-based analysis was used to infer decisions about social media data. The advancement of deep learning has made it possible to learn from images and associated text in a self-supervised way. The idea originated from the works in [17], where the researchers proposed that instead of predicting the ImageNet classes of the image, deep learning models can be used to infer the Word2Vec representations of their corresponding labels. Researchers used the idea to analyze social media data. In [18], the authors proposed a method to mine data from different social media platforms and modalities related to New York City and examined the behavior of the citizens. In [19], the authors used geo-located social media photographs to identify the most visited places by tourists in major European cities. In [20], the authors used popular hashtags on Instagram to find out cultural differences between different cities and neighborhoods. In [21], the authors proposed a joint embedding of pictures of food and recipes. They were able to achieve interesting results on an image-recipe retrieval task. In [22], the authors analyzed

Instagram posts to study how the different neighborhoods, events, and cultures of the city are presented on Instagram. Researchers detected race, age, and gender of people from New York City Instagram images in [23]. They have also analyzed the social diversity of different neighborhoods and compare them to census-based metrics.

III. PRELIMINARIES

A. Convolutional Neural Network (CNN)

CNNs are a class of deep learning networks that can infer intricate visual patterns and symmetries from a set of images. It learns filters of specific kernel size and understands these patterns after passing the images through the filters [24]. The CNN network that is used in this project is ResNet-50 [25].

Mathematically, an image is convolved at every layer in a CNN and features are made more accessible for dense layer at the final classification task. Suppose x_m^l is the m^{th} input feature in the l^{th} layer then the n^{th} output feature in that y_n^l is calculated as:

$$y_n^l = \sum_m^{M^{l-1}} w_{n,m}^l * x_m^l + b_n^l$$

Here, where M^{l-1} is the number of input maps, $*$ denotes convolution operation, and b_n^l is the bias of the n^{th} output map in the l^{th} layer.

B. Recurrent Neural Network (RNN)

RNNs are a class of deep learning networks that consists of a series of memory cells, which allows input from the previous memory cells and can infer the context from time-series data, then subsequently generate output [26].

For defining how RNN computes mathematically its output, following equation is described:

$$o^t = f(h^t; \theta)$$

$$h^t = g(h^{t-1}, x^t; \theta)$$

where o^t is the output of the RNN at time t , x^t is the input to the RNN at time t , and h^t is the state of the hidden layer(s) at time t .

C. Attention based LSTM networks

Aspect based sentiment analysis requires a sort of mechanism which can have a weighted focus on variegated words in a particular set of sentences. This mechanism is called attention, which gives different weights to words in a sentence while considering various types of aspects of the sentence [27].

Taking H be a matrix comprising of hidden vectors $[h_1, \dots, h_N]$ that the LSTM has yielded, where d is the dimensionality of hidden layers & N is the length of the particular input sentence. Let v_a represents the embedding of aspect and e_N be a identity vector. The attention mechanism will produce an attention weight vector α and a weighted hidden representation r . Thereby, it allows the attention

mechanism to focus on a particular segment of the sentence while considering one specific type of aspect[27].

$$M = \tanh \left(\begin{bmatrix} W_h H \\ W_v v_a \otimes e_N \end{bmatrix} \right)$$

$$\alpha = \text{softmax}(w^T M)$$

$$r = H\alpha^T$$

where, M, α, r, W_h, W_v and w are projection parameters.

The final sentence representation is given by:

$$h^* = \tanh(W_p r + W_x h_N)$$

where, h^*, W_p and W_x are projection parameters to be learned during training. h^* is defined as the feature representation of a sentence with given an input aspect. Then linear layer to convert sentence vector to e , followed by softmax layer to transform e into a conditional probability is implemented.

$$y = \text{softmax}(W_s h^* + b_s)$$

where W_s and b_s are the parameters for softmax layer. For aspect-based sentiment analysis, the input aspect embedding is appended with each word input vector.

D. Transfer Learning

Learning from a pre-trained neural network can be applied to a different application. Usually, the initial layers of networks are kept frozen, and the final layers are trained again on the new data. This method enables transferring the learning from one application to another similar type of application without requiring much data.



Fig. 1: Different categories of data.

IV. PROPOSED METHODOLOGY

Our proposed methodology can be defined by dividing the whole procedure into four tasks. We scrape data from Instagram, clean the data, annotate the data, and apply deep learning models on the collected data. We obtain results for understanding the sentiment from images and text captions both.

A. Data collection

We use igramscraper, a python API [28] to collect data from Instagram using hashtags #wuhan, #corona, and #coronavirus. We scrap the data from February 16 to March 20. We save the metadata related to each post in excel files and the corresponding media data (for videos we save the first frame). The metadata we store for each post are the post identifier, Instagram short code (i.e., B8o8MQHJbMc), created time, data type (image or video), source link to the post, low resolution image source link, high resolution image source link, post caption, owner id, number of likes received, number of comments received, and location name. All the posts we scrap from Instagram are public posts and visible to everyone.

B. Data Cleaning

After gathering the data, we perform data cleaning. The data cleaning process is divided into several sub-tasks as following:

- Data deduplication: In this step, we remove the data which are repeated.
- Data completeness: The part of data that is captured incompletely during the scraping process, can either be deleted or imputed to fill in meaning. In our case, missing data imputation may provide us wrong impressions about sentiments. For this reason, we delete the incomplete data points.
- Corrupted Data: The corrupted data points are also removed.




Fig. 2: Class distribution of image annotations.




Fig. 3: Class distribution of caption annotations.

C. Data Annotation

We categorize our annotated data into five categories, as shown in Fig. 1. Humor, funny, and memes posts are classified as the first class. News articles, awareness posts, and information about cloud-19 and lockdown are marked as the second class. Happy, supportive, appreciable, positive, and nostalgic posts are put together in class 3. Negative posts regarding China, xenophobic, racist behavior, blaming, and protest against isolation are regarded as category 4. The class distribution of the image and caption annotations are illustrated in Fig. 2 and Fig. 3. As depicted in the graphs, it can be inferred that the occurrence of "Positive" type and "Neutral/News" type annotations are the maximum, while "Negative" type annotation is observed the least.

During the annotation process, we use Google Lens for translating the media content to help us with the annotation process. Google translate Chrome extension is used for translating the caption. Further discussion about which type of content is segregated in which kind of annotation is given below (see Fig. 4).

1) *Memes/Humor Type*: This category comprises of any sort of meme or humor content, which mostly included people sharing memes surrounding toilet papers, corona beer, along with hilarious alternatives to sanitizers, and social distancing.

2) *News/Neutral Type*: This class covers posts, which were shared just to spread awareness about healthy habits like washing hands, wearing masks, and social distancing to be practiced during the pandemic. It also covers news information shared by a news channel on Instagram, along with advertisement about sanitizers, closure announcement of stores, info about toilet papers, masks, sanitizers available in a particular store. It also includes news reports about lower pollution and emptier streets.

3) *Positive Type*: This category includes any Instagram post, which is targeted to spread not only physical but also mental positivity among the users. It covers posts regarding how people are spending their leisure time while quarantining themselves, which includes reading novels, dancing, exercising, playing video games, having movie watch parties,

and singing. It also includes Instagrammers posting about the appreciation of the front-line workers, medical teams working during the pandemic, and wearing unique masks. This type of annotation also includes people posting their nostalgic past travel pictures and sharing pictures of how they are spending more time with their adorable pet dogs and cats. Medical doctors and nurses posting their selfies are also included in this category.

4) *Negative Type*: The negatively annotated data includes posts criticizing the political systems, xenophobic posts against China, anger posts against imposed lockdown, fear-mongering users by sharing content with people dying, ranting about loss of freedom, conducting public protests against the government. People blaming government and media for hiding factual data and spreading conspiracy theories also fall into this class.

5) *Random Type*: This class is used to separate Instagram influencers' posts and advertisements, which targets the promotion of unessential products by using the hashtags #wuhan and #coronavirus. It shows that the hashtags were misused to promote irrelevant agendas.

D. Algorithm

We process the video data by taking the first frame of the video and save it as a single image which is stored along with the scraped image dataset. We extract audio from the video and run a subtitle generator on it. Thus, subtitles that are up to 2 minutes are generated for the video and merged on top of stored caption data for the related post. Pytesseract API is used to extract text written on the image and saved along with the caption data stored for that post. Before deploying the deep learning models, first pre-processing steps are applied to caption text data and is translated into English using python API and trimmed up to word length of 300 words. Our proposed algorithms look at a single frame (image or the first frame of the video) for each post and the corresponding caption data.

The text data is used for training attention-based LSTM network for textual sentiment analysis. The image data is used for image classification using CNNs. We choose ResNet-50 for the image classification task. The proposed algorithm is illustrated in Fig. 5.

V. EXPERIMENTS AND RESULTS

We perform several analyses to get insights into our collected dataset. We also explore the performance of deep learning algorithms for checking the feasibility of automatic analysis. We use the Tensorflow library for implementing deep learning algorithms. All of the experiments regarding training and implementation of the models are performed in a hardware environment, which includes Intel Xeon, 2.60 GHz CPU and Nvidia Quadro M4000 (8 GB Memory) GPU.

A. Location based analysis:

We collect longitude and latitude for the location of the posts to perform several location-based analyses on our collected data. It enables us to see which regions have been most



Fig. 4: Samples from different classes.

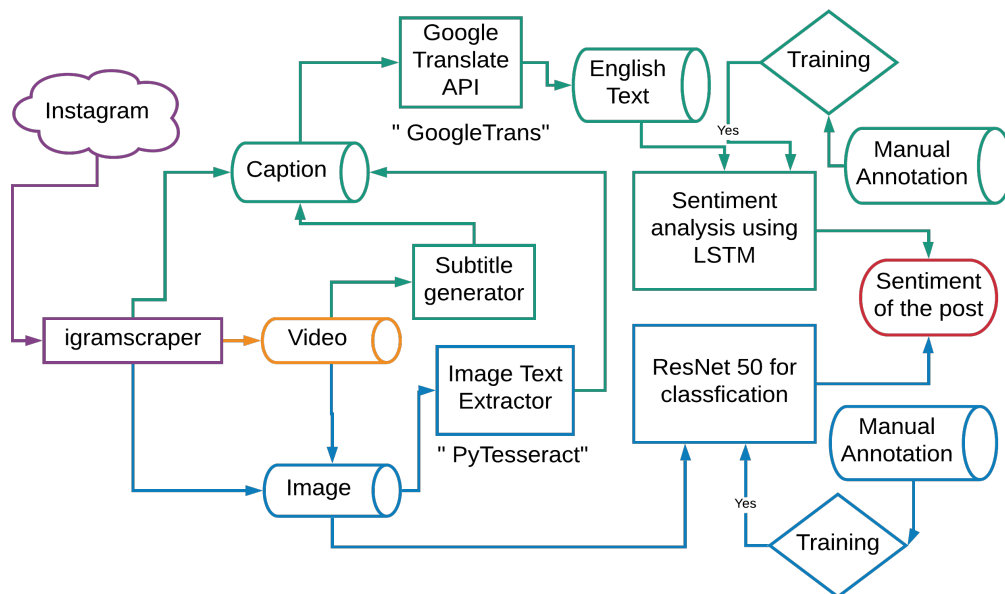


Fig. 5: Proposed algorithm flowchart.

active and learn the degree of attention posts from different regions received from Instagram users. We use the number of

likes and comments as indications for measuring attention. The data is visualized based on its longitude and latitude

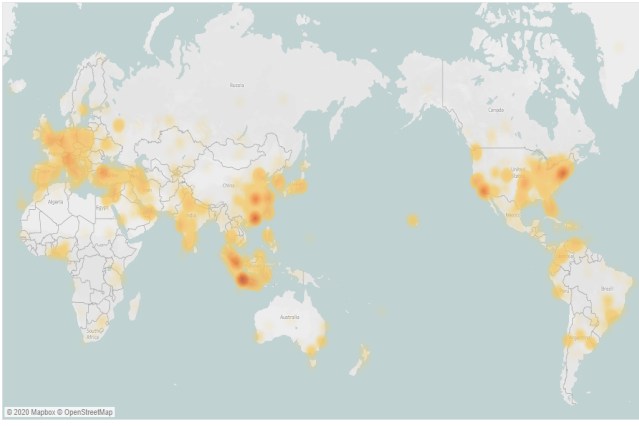


Fig. 6: Geographic representation of Instagram activity for collected data based on number of posts.

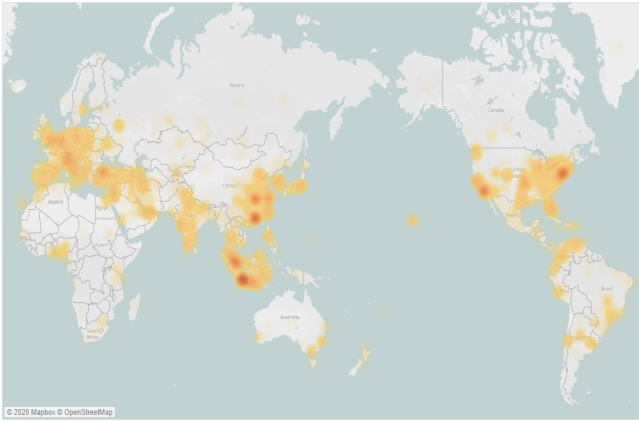


Fig. 7: Geographic representation of Instagram activity for collected data based on the number of likes received.

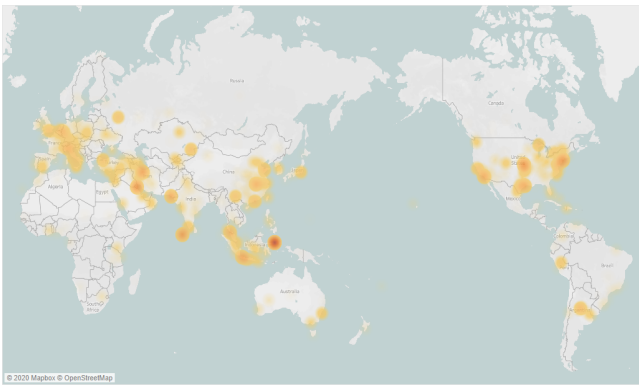


Fig. 8: Geographic representation of Instagram activity for collected data based on the number of comments received.

using Tableau. The geographic representation of data shows the intensity of global discussion worldwide on the Instagram platform. We illustrate the Instagram activity based on the number of posts, the number of likes received, and the number of comments received in Fig. 6-8, respectively. In

Fig. 6, we can see that Indonesia, China, United States, Turkey, United Kingdom, Germany, Malaysia, and Italy have been most active concerning the number of posts posted on Instagram. It is also visible that Russia was quite inactive, given its 46M users [9]. We see somewhat similar trends in Fig. 7 and Fig. 8.

B. Country based analysis

To understand the demographic viewpoint further, we extract the country name based on the location of the Instagram posts using the pycountry library. The top 15 countries with the maximum number of posts are illustrated with respect to the count of the different class of annotations in Fig. 9. Mostly, News/neutral, memes, and positive content are observed among the countries. The United States deviates from this trend as more negative posts are observed than positive posts.

C. Further insights into the data

For getting further insights into our collected dataset, we plot several graphs by taking combinations of the following parameters:

- class of images in the Instagram posts (denoted by `Img_ANN`),
- class of captions of the posts (denoted by `Cap_ANN`),
- number of likes (denoted by `likes_count`),
- number of comments (denoted by `comments_count`).

The following Fig. 10 shows the class-overlap between the images and captions for the posts. It is visible that there are not many mismatches between the image annotation and the caption annotation for the Instagram posts.

To understand the pattern between the number of likes received and the number of comments received, a scatter-point graph between these two parameters is plotted (see Fig. 11). From the chart, we can infer two things. Firstly, negative posts tend to receive more comments than likes, as people provide their own opinion on that post without liking it. Secondly, positive posts receive more likes than comments.

A scatter-plot is shown with the following variables: the number of likes received, class of the image and the class of the and caption for each post (see 12). We can see that posts with negative images but positive captions receive a high number of likes. It indicates that positive caption can change the class of the overall post and make it positive, despite containing a negative image.

D. Performance of deep learning algorithms:

The results we obtained from our experiments on the collected data are shown in Table I. We perform sentiment analysis on images and captions separately. For textual sentiment analysis; we use attention-LSTM [27] network and train it on Stanford's Movie Review Dataset [29]. We achieved 87.7% accuracy. Later, we applied transfer learning

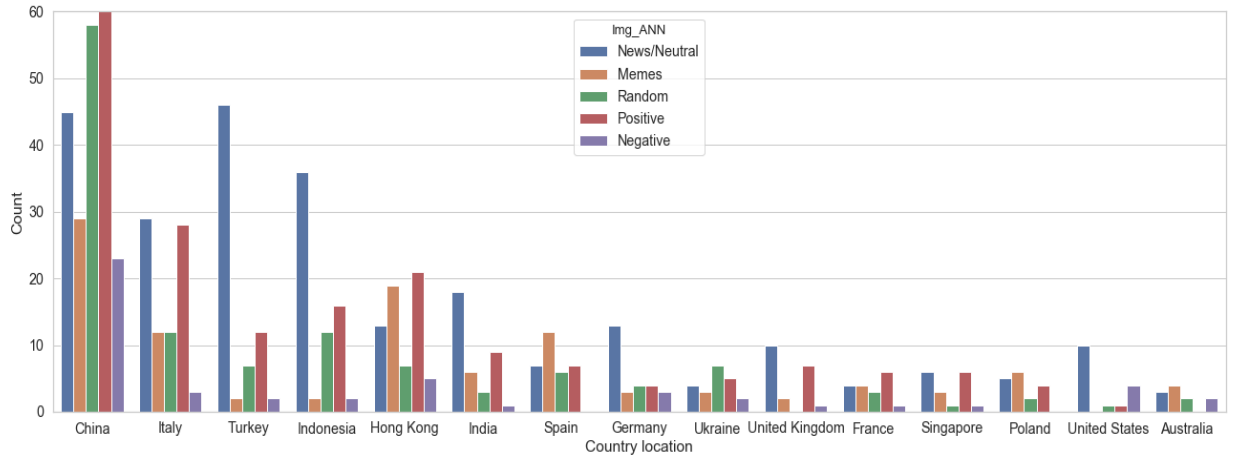


Fig. 9: Bar Graph depicting sentiment in different countries (y-axis is capped at 60)

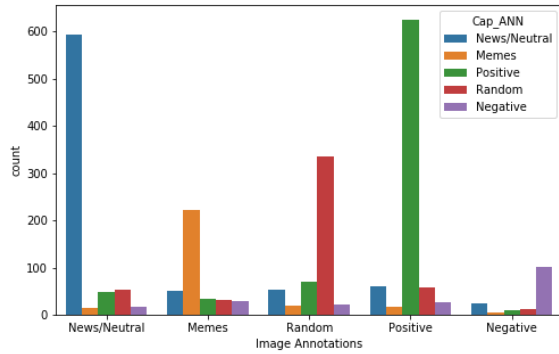


Fig. 10: Class-overlap between the images and the captions for each post.




Fig. 11: Number of comments vs. number of likes.

and train the model on our annotated dataset. We achieved an accuracy of 73.6%.

TABLE I: Performance of proposed deep learning algorithms.

#	Model	Dataset	Acc.
1	Textual Sentiment (Attention LSTM) [27]	Movie Review Dataset [29]	87.7%
2	Textual Sentiment (Attention LSTM) [27]	Our dataset	73.6%
3	Image Sentiment Analysis (ResNet-50) [30]	Reddit Memes Dataset [31]	82.1%
4	Image Sentiment Analysis (ResNet-50) [30]	Our Dataset	79.2%




Fig. 12: Scatterplot based on number of likes between image annotations and the corresponding caption annotations [y-scale is limited to 5000 likes]

For conducting sentiment analysis on images, we used ResNet-50 [25] as suggested in the following paper [30]. We used Reddit Memes Dataset [31] first to train the network and obtained 82.1 % accuracy. Afterwards, we fine-tuned

the pre-trained Resnet-50 on our dataset and achieved 79.2 % accuracy. It should be noted that we exclude the ‘random’ class while implementing our algorithms, and the networks are trained for classifying four classes.

VI. CONCLUSION

In this paper, we have studied the reaction of Instagram users, in light of the recent outbreak of the coronavirus. It is essential to evaluate the public emotion in times of crisis and social media data, particularly Instagram, can be a valuable source for conducting such analysis in real-time. Consequently, we have collected in total, around 100k Instagram posts. Among them, around 14% of the posts containing location data were used for our location-based analysis. We manually annotated around 20% of the posts with location to perform our reaction analysis to the virus outbreak. The collected images and videos on Instagram reveal a range of emotions, highlighting social awareness, and mass mayhem. We found that users from most countries share positive and awareness posts. We also used deep learning algorithms to conduct the classification. Results show that the large-scale sentiment analysis can be automated with the help of deep learning algorithms.

In our future work, we want to improve the accuracy of the deep learning models and perform the emotion analysis at a larger scale.

REFERENCES

- [1] M. Moreno, D. Christakis, K. Egan, L. Brockman, and T. Becker, “Associations between displayed alcohol references on facebook and problem drinking among college students: 125,” *Alcoholism: Clinical & Experimental Research*, vol. 35, no. 1, p. 2056305117691544, 2017.
- [2] S. Alhabash and M. Ma, “A tale of four platforms: Motivations and uses of facebook, twitter, instagram, and snapchat among college students?” *Social Media+ Society*, vol. 3, no. 1, p. 2056305117691544, 2017.
- [3] D. Chaffey, “Global social media research summary 2016,” *Smart Insights: Social Media Marketing*, 2016.
- [4] T. Leaver and T. Highfield, “Visualising the ends of identity: pre-birth and post-death on instagram,” *Information, Communication & Society*, vol. 21, no. 1, pp. 30–45, 2018.
- [5] P. Zhou, X.-L. Yang, X.-G. Wang, B. Hu, L. Zhang, W. Zhang, H.-R. Si, Y. Zhu, B. Li, C.-L. Huang *et al.*, “A pneumonia outbreak associated with a new coronavirus of probable bat origin,” *nature*, vol. 579, no. 7798, pp. 270–273, 2020.
- [6] “2019-ncov global cases study (by johns hopkins csse),” <https://gisanddata.maps.arcgis.com/apps/opsdashboard/index.html>, accessed: 2020-01-31.
- [7] “Coronavirus live updates: Third case of coronavirus confirmed in london, ont., bringing canadian cases to four,” <https://nationalpost.com/news/world/coronavirus-live-updates-wuhan-virus-china-2019ncov>, accessed: 2020-01-23.
- [8] “Statement on the second meeting of the international health regulations (2005) emergency committee regarding the outbreak of novel coronavirus (2019-ncov),” [https://www.who.int/news-room/detail/30-01-2020-statement-on-the-second-meeting-of-the-international-health-regulations-\(2005\)-emergency-committee-regarding-the-outbreak-of-novel-coronavirus-\(2019-ncov\)](https://www.who.int/news-room/detail/30-01-2020-statement-on-the-second-meeting-of-the-international-health-regulations-(2005)-emergency-committee-regarding-the-outbreak-of-novel-coronavirus-(2019-ncov)), accessed: 2020-01-31.
- [9] “statista,” <https://www.statista.com/statistics/578364/countries-with-most-instagram-users/>.
- [10] R. Castillo, D. Carlat, T. Millon, C. Millon, S. Meagher, S. Grossman, A. P. Association *et al.*, “Diagnostic and statistical manual of mental disorders,” *American Psychiatric Association Press, Washington, DC*, 2007.
- [11] B. Wild, M. Eichler, H.-C. Friederich, M. Hartmann, S. Zipfel, and W. Herzog, “A graphical vector autoregressive modelling approach to the analysis of electronic diary data,” *BMC medical research methodology*, vol. 10, no. 1, p. 28, 2010.
- [12] A. G. Reece and C. M. Danforth, “Instagram photos reveal predictive markers of depression,” *EPJ Data Science*, vol. 6, no. 1, pp. 1–12, 2017.
- [13] K. Lup, L. Trub, and L. Rosenthal, “Instagram# instasad?: exploring associations among instagram use, depressive symptoms, negative social comparison, and strangers followed,” *Cyberpsychology, Behavior, and Social Networking*, vol. 18, no. 5, pp. 247–252, 2015.
- [14] N. Andalibi, P. Ozturk, and A. Forte, “Depression-related imagery on instagram. 2015 feb 28 presented at: The 18th acm conference companion on computer supported cooperative work & social computing (pp.),” 2015.
- [15] R. Gomez, L. Gomez, J. Gibert, and D. Karatzas, “Learning from barcelona instagram data what locals and tourists post about its neighbourhoods,” in *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018, pp. 0–0.
- [16] Sensis, “Sensis social media report 2017,” 2017.
- [17] M. Norouzi, T. Mikolov, S. Bengio, Y. Singer, J. Shlens, A. Frome, G. S. Corrado, and J. Dean, “Zero-shot learning by convex combination of semantic embeddings,” *arXiv preprint arXiv:1312.5650*, 2013.
- [18] Y.-H. Kuo, Y.-Y. Chen, B.-C. Chen, W.-Y. Lee, C.-C. Wu, C.-H. Lin, Y.-L. Hou, W.-F. Cheng, Y.-C. Tsai, C.-Y. Hung *et al.*, “Discovering the city by mining diverse and multimodal data streams,” in *Proceedings of the 22nd ACM international conference on Multimedia*, 2014, pp. 201–204.
- [19] J. C. García-Palomares, J. Gutiérrez, and C. Mínguez, “Identification of tourist hot spots based on social networks: A comparative analysis of european metropolises using photo-sharing services and gis,” *Applied Geography*, vol. 63, pp. 408–417, 2015.
- [20] S. Chang, “Instagram post data analysis,” *arXiv preprint arXiv:1610.02445*, 2016.
- [21] A. Salvador, N. Hynes, Y. Aytar, J. Marin, F. Ofli, I. Weber, and A. Torralba, “Learning cross-modal embeddings for cooking recipes and food images,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017, pp. 3020–3028.
- [22] J. D. Boy and J. Uitermark, “Reassembling the city through instagram,” *Transactions of the Institute of British Geographers*, vol. 42, no. 4, pp. 612–624, 2017.
- [23] V. K. Singh, S. Hegde, and A. Atrey, “Towards measuring fine-grained diversity using social media photographs,” in *Eleventh International AAAI Conference on Web and Social Media*, 2017.
- [24] J. Gu, Z. Wang, J. Kuen, L. Ma, A. Shahrourdy, B. Shuai, T. Liu, X. Wang, G. Wang, J. Cai *et al.*, “Recent advances in convolutional neural networks,” *Pattern Recognition*, vol. 77, pp. 354–377, 2018.
- [25] K. He, X. Zhang, S. Ren, and J. Sun, “Deep residual learning for image recognition,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 770–778.
- [26] A. Sherstinsky, “Fundamentals of recurrent neural network (rnn) and long short-term memory (lstm) network,” *arXiv preprint arXiv:1808.03314*, 2018.
- [27] Y. Wang, M. Huang, X. Zhu, and L. Zhao, “Attention-based lstm for aspect-level sentiment classification,” in *Proceedings of the 2016 conference on empirical methods in natural language processing*, 2016, pp. 606–615.
- [28] “instagramscraper,” <https://pypi.org/project/igramscraper/>.
- [29] A. L. Maas, R. E. Daly, P. T. Pham, D. Huang, A. Y. Ng, and C. Potts, “Learning word vectors for sentiment analysis,” in *Proceedings of the 49th annual meeting of the association for computational linguistics: Human language technologies-volume 1*. Association for Computational Linguistics, 2011, pp. 142–150.
- [30] V. Gajaria and A. Gupta, “Emotion detection and sentiment analysis of images.”
- [31] “Reddit Memes Dataset,” <https://www.kaggle.com/sayangoswami/reddit-memes-dataset/version/1>.