Journal Homepage: - www.journalijar.com

INTERNATIONAL JOURNAL OF ADVANCED RESEARCH (IJAR)

Article DOI: 10.21474/IJAR01/14398
DOI URL: http://dx.doi.org/10.21474/IJAR01/14398

RESEARCH ARTICLE

IMAGE SENTIMENTAL ANALYSIS: AN OVERVIEW

V. Sunil Kumar, Vedashree C.R and Sowmyashree S.
Department of Computer Science and Engineering, East West Institute of Technology, Magadi Main Road, Anjana Nagar, Bangalore, Karnataka, 560091, India.

Manuscript Info

Manuscript History
Received: 10 January 2022
Final Accepted: 15 February 2022
Published: March 2022

Key words:-
Visual Content, Emotion, Sentiments, Image Sentiment Analysis, Social Image Sentiment, Image Polarity Classification

Abstract

Visual content, such as photographs and video, contains not only objects, locations, and events, but also emotional and sentimental clues. On social networking sites, images are the simplest way for people to communicate their emotions. Images and videos are increasingly being used by social media users to express their ideas and share their experiences. Sentiment analysis of such large-scale visual content can aid in better extracting user sentiments toward events or themes, such as those in image tweets, so that sentiment prediction from visual content can be used in conjunction with sentiment analysis of written content. Despite the fact that this topic is relatively new, a wide range of strategies for various data sources and challenges have been created, resulting in a substantial body of study. This paper introduces the area of Image Sentiment Analysis and examines the issues that it raises. A description of new obstacles is also included, as well as an assessment of progress toward more sophisticated systems and related practical applications, as well as a summary of the study's findings.

Introduction:-

With the widespread use of social media platforms (such as social networks, forums, and blogs), political parties, businesses, and other organizations are increasingly considering public opinion in their business strategy. Images and videos are increasingly being used by social media users to express their ideas and share their experiences. Sentiment analysis of such large-scale visual content can aid in better extracting user sentiments toward events or themes, such as those in image tweets, so that sentiment prediction from visual content can be used in conjunction with sentiment analysis of written content. A thousand words are considered to be worth a picture. When it comes to expressing human feelings and sentiments, it is unquestionably more valuable. With the introduction of social media, an increasing number of people have begun to utilize images on social media platforms such as Flickr and Instagram to communicate their excitement, rage, and boredom. Many applications in health-care, anthropology, communication studies, marketing, and many sub-areas of computer science, such as computer vision, are becoming increasingly dependent on automatic inference of emotion and from the ever-increasing, vast volumes of user-generated pictures, sentiment data can be gleaned. It introduces self-empathy, for example, making a person more conscious of their sentiments. It also boosts one's self-esteem and resilience, allowing them to recover quickly from poor mental health as well as physical hardship. The goal of sentiment analysis is to extract people's attitudes on a topic or the author's intended emotional effect on the readers. Since opinions impact many human decisions in both economic and social activities, sentiment analysis has various practical uses. The polarity categorization of an input text (e.g., derived from a review, a comment, or a social post) in terms of positive, negative, or neutral polarity is the

Corresponding Author:- V. Sunil Kumar
Address:- Department of Computer Science and Engineering, East West Institute of Technology, Magadi Main Road, Anjana Nagar, Bangalore, Karnataka, 560091, India.

basic task in Sentiment Analysis. This analysis can be done at the level of the document, sentence, or feature. The visual features of users' everyday activities and interests are reflected in the photographs published on social media platforms. Such rapidly rising user-generated graphics constitute a new and significant source of data for analyzing consumers' preferences. Applications such as brand monitoring, market prediction, and political voting predictions are possible thanks to the study of such social data. Companies are particularly interested in tracking public perceptions of their products or services (e.g., brand monitoring), while customers consult other users' feedback to gauge the quality of a service or a product. The polarity classification of an input text in terms of positive, negative, or neutral polarity is the basic task in Sentiment Analysis.

Consider the scenario depicted in Figure 1.1. We can analyze a person's mental wellness based on the mood and sentiment inferred from her images on social media sites, as individuals are increasingly using photos to record their daily lives.

(a) “Girlfriend crying a lot when I proposed to her”, (b) “Crying baby after her toy was taken”.
Fig 1.1 An example shows affective gap

State of The Art:-

Using a Transfer Learning Approach, Visual Sentiment Analysis for Social Images 2016 IEEE International Conferences on Big Data and Cloud Computing (BDCloud), Social Computing and Networking (SocialCom), and Sustainable Computing and Communications (SustainableCom) Jyoti Islam, Yanqing Zhang (SustainCom). They present a unique visual sentiment analysis system that employs a transfer learning technique to predict sentiment in this research. To avoid over fitting, they use hyper-parameters obtained from a very deep convolution neural network to initialize our network model. They undertake extensive experiments on a Twitter image dataset, demonstrating that our model outperforms the current state-of-the-art.

Yuhai Yu1, Hongfei Lin, Jiana Meng, and Zhehuan Zhao present Visual and Textual Sentiment Analysis of a Micro blog Using Deep Convolution Neural Networks. They use deep learning models in a convolution neural network (CNN) to analyze sentiment in Chinese micro blogs using both textual and visual content in this paper. For textual sentiment analysis, they use a CNN built on top of pre-trained word vectors, and for visual sentiment analysis, they use a deep convolution neural network (DNN) with generalized dropout. They next test our sentiment prediction framework on a dataset of text and related images taken from a popular Chinese social media network (Sina Weibo), demonstrating state-of-the-art results on this Chinese sentiment analysis benchmark.

Quanzeng You and Jiebo Luo, Hailin Jin, and Jianchao Yang, Robust Image Sentiment Analysis Using Progressively Trained and Domain Transferred Deep Networks. They begin by designing an appropriate CNN architecture for picture sentiment analysis in this work. They get half a million training data by labeling Flickr photographs with a baseline sentiment algorithm. We adopt a progressive technique to fine-tune the deep network to make use of such noisy machine labeled data. We also increase Twitter image performance by inducing domain transfer using a small number of manually annotated Twitter images. They've done a lot of testing with manually labeled Twitter photos. The results reveal that the suggested CNN outperforms competing algorithms when it comes to image sentiment analysis.

Social Media Image Sentiment Analysis Baoxin Li and Yilin Wang They investigate the topic of deducing human sentiments from a huge collection of Internet photographs using both image attributes and contextual social network information in this proposal (such as friend comments and user description). Despite significant progress in detecting user sentiment based on word data, sentiment analysis of image content has mostly been overlooked. They apply the major improvements in text-based sentiment prediction tasks to the more difficult problem of predicting the underlying sentiments behind the photos. They show that neither visual nor textual features are sufficient for reliable sentiment labeling on their own. As a result, they avoid employing both and define the sentiment prediction problem in two scenarios: supervised and unsupervised.

The majority of efforts in Image Sentiment Analysis are based on prior studies on emotional-aware image retrieval (Colombo et al., 1999; Schmidt and Stock, 2009), which attempt to establish links between emotions and picture visual attributes in order to improve image retrieval and classification. The authors of (Machajdik and Hanbury, 2010) investigated the role of visual elements in emotive image classification. The picture characteristics were chosen based on experimental observation and insights into emotional responses to different hues. Anger, Awe, Amusement, Excitement, Contentment, Disgust, Fear, and Sad were the eight emotional classifications examined in the image classification (Mikels et al., 2005). Sun et al. (2016) suggested an algorithm that integrates data from the entire image (global features) and prominent items in salient regions.

The research in (Ortis et al., 2018) looked at the problem of noise in the language connected with social photos by users, which is caused by their subjectivity. The authors developed an image sentiment classifier that uses a representation based on both visual and textual features extracted from an image, and used four deep models trained on different tasks to evaluate the results obtained by using the text provided by users (subjective) and the text extracted from the visual content. Visual Sentiment Analysis is a relatively new field of study. The majority of the research in this new subject is based on prior studies on emotional semantic image retrieval [17–20], which tie low-level image features to emotions in order to perform automatic image retrieval and categorization.

The first work on Visual Sentiment Analysis was published in 2010 [22] and seeks to identify images as "good" or "negative." The authors investigated the relationships between image sentiment and visual content in this study. Based on the accompanying text, they assigned numerical sentiment scores to each image (i.e. meta-data). The authors achieved this by extracting sentiment score values from the text linked with photographs using the SentiWordNet [23] lexicon.

In [24] presents a study on the elements that are beneficial in the task of emotive classification of photographs. To scientifically choose the picture qualities, researchers used ideas from the experimental observation of emotional responses to colors and art. The authors used the eight emotional output categories outlined in [25] to do the emotional image classification.

Given that a person's emotional response to an image may include multiple emotions, the authors of [27] compared three methods for predicting such emotion distributions: a Support Vector Regressor (based on hand-crafted features related to edge, color, texture, shape, and saliency), and a CNN for both classification and regression. They also advocated changing an image's texture and colors to vary the generated emotion distribution. Given a source image and a target image, the proposed approach converts the source image's color tone and textures to those of the target image. As a result, the altered image generates feelings that are more similar to the target image than the original. Using four similarity metrics between distributions, this strategy has been quantitatively tested.

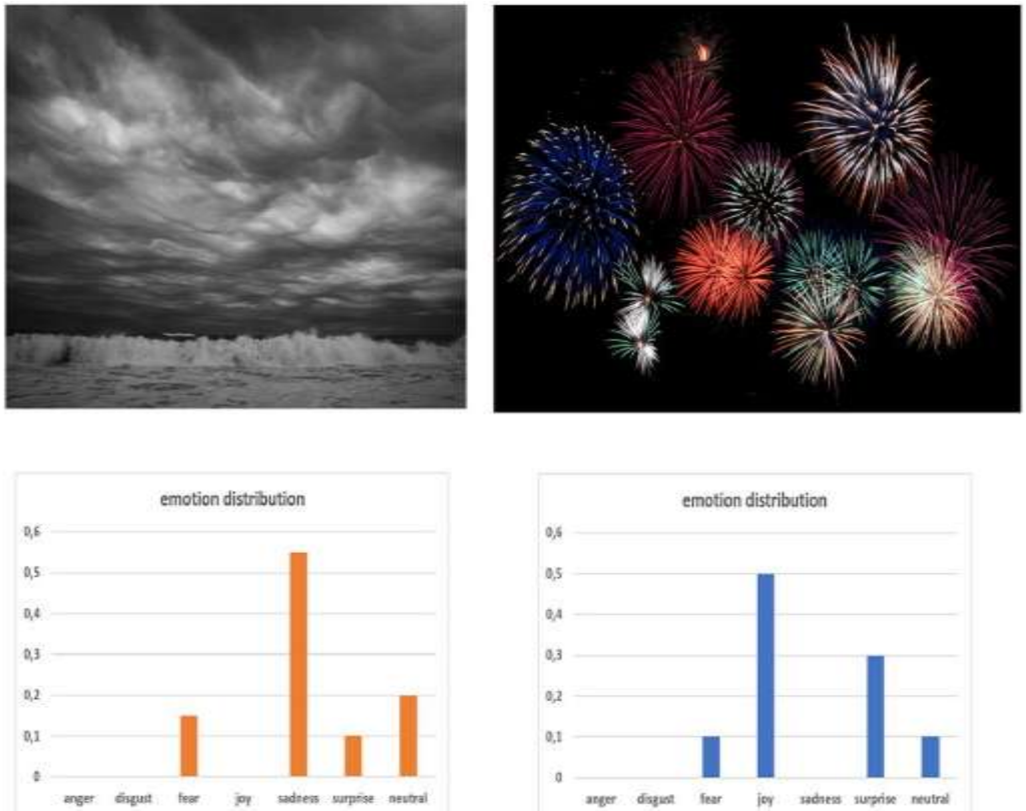

Fig2:- Examples of image emotion distributions.

Katsurai and Satoh [29] used visual, linguistic, and sentiment characteristics to create a latent embedding space with the highest correlation between projected features from different views. This work uses the CCA (Canonical Correlation Analysis) technique to create a 3-view embedding that serves as a tool for encoding inputs from various sources (i.e., a text and an image with similar meaning/sentiment are projected near each other in the embedding space) as well as a method for obtaining a sentiment representation of images (by simply projecting an input feature to the latent embedding space). This format is used to train a linear SVM classifier to determine if something is positive or negative.

The works in [30, 31] classify images using several emotional categories to conduct emotional image classification. Instead of training a model to predict only one sentiment label, as described by Peng et al. [27] in 2015, the authors analyzed a distribution over a series of pre-defined emotional labels. They suggested a multi-task system that optimizes categorization and distribution prediction at the same time to achieve this goal. The authors introduced two Conditional Probability Neural Networks (CPNN): Binary CPNN (BCPNN) and Augmented CPNN (ACPNN) in [30]. (ACPNN). A CPNN is a hidden layer neural network that takes either features or labels as input and outputs the label distribution. Indeed, a CPNN's goal is to forecast the probability distribution over a set of labels.

The authors of [32] built on their prior work [28], in which they trained a CNN for sentiment analysis and then empirically examined the contributions of each layer. They employed the activations in each layer to train distinct linear classifiers in particular. The authors also investigated the influence of weight initialization for fine-tuning by altering the task (i.e. the output domain) for which the fine-tuned CNN was originally trained in this study. After that, based on the experimental findings and observations, the authors propose an enhanced CNN design. After that, based on the empirical findings, the authors propose an upgraded CNN design.

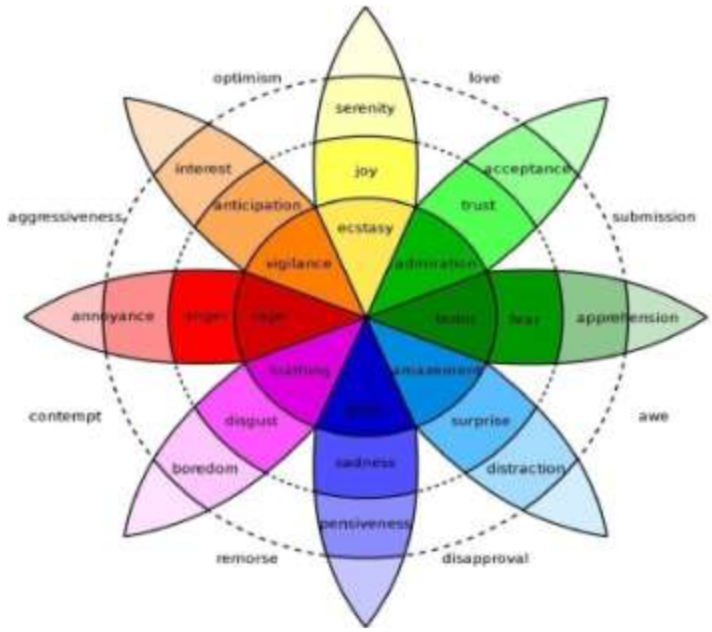

Fig3:- Plutchik's wheel of emotions.

Previous surveys [33-35] provided a quick overview of the field before focusing on specific topics. The proposed study, on the other hand, starts with the most basic task in Visual Sentiment Analysis (classifying an image as positive or negative) and then expands on it by analyzing each aspect in depth, covering a larger number of topics, derived tasks, and new challenges in Visual Sentiment Analysis in great depth (Fig.3).

Problem Presentation:-

The key factors that influence sentiment are discussed in this section. This enables for a more focused examination of the linked sub-issues that make up the Image Sentiment Analysis problem, as well as the development of more robust techniques. A target (or topic) and a sentiment are the two main components of an opinion. Opinions can come from a variety of people, and they can change over time. This means that the system must consider the person who has the view as well as the time when the opinion is expressed. There are various distinctions in the case of sentiment analysis on visual content. Sentiment Analysis can use NLP (Natural Language Processing) techniques to extract distinct parts of a sentence and link each word with a specific meaning when the input is text. Associating visual cues to emotion categories or polarity scores is difficult when the input is an image. Each of the previously described sentiment components (entity, aspect, holder, and time) is discussed in the context of Image Sentiment Analysis in the following paragraphs.

Sentiment Entity and Aspects:-

The entity is the analysis's subject (or target). The entity in the case of Image Sentiment Analysis is the input image. An entity can be thought of as a collection of "components" and "attributes" in general. The collection of "aspects" is made up of the entity's parts, attributes, and the special aspect "GENERAL." With multiple levels of visual elements, this structure can be transported to the image domain. A group of sub-images can be used to define the parts of an image. Several computer vision techniques, such as background/foreground extraction, picture segmentation, multi-object detection, and dense captioning, can be used to create this set (Karpathy and Fei-Fei, 2015). An image's attributes refer to its visual quality features, which are frequently derived via extracting low-level information. A emotion score for each aspect can be calculated using this structured image hierarchy. Finally, the partial scores can be integrated properly to achieve the sentiment classification (e.g., data can be used as input features of a regression model). The concept connected with the image context can be considered. Several works on personal contexts (Ortis et al., 2017; Furnari et al., 2018) and scene recognition can be employed for this, and the inferred concepts can be used to extract the corresponding sentiment from the visual perspective. Image parts and characteristics can also be used to extract sentiment scores. An alternate way to representing picture components as a group of sub-images is to rely on a textual description of the represented scene. A photo's description can be tailored to a specific goal of visual comprehension. We can get several descriptions of the same image from different points of view by modifying the job. The structure described above can then be created by combining these complementing notions. By performing textual Sentiment Analysis, the majority of present efforts in evaluating social media make use of textual information manually paired with photographs. Although the text linked with social photos is commonly used in the state-of-the-art to improve the semantics inferred from photographs, because it is contributed by the users, it can be a very noisy source. The authors of (Ortis et al., 2018) presented a work on Image Polarity Prediction utilising Objective Text derived directly from photos, and compared such text to the Subjective (i.e., user-provided) text in- formation usually employed in state-of-the-art techniques. This approach provides a user-independent source of text that specifies the semantics of images, which is important for addressing challenges linked to the inherent subjectivity of language connected with images.

Sentiment Holder:-

Almost all works on Image Sentiment Analysis ignore the sentiment holder or only evaluate the image publisher's sentiment implicitly. At least two holders can be identified in this context: the image owner and the image viewer. Understanding the linkages between the owner's planned sentiment and the actual sentiment created in the viewers is critical when considering the example of an advertising campaign. Social media networks now provide a very effective means of retrieving real-time and large-scale data about people's reactions to subjects, events, and advertising campaigns. This field of study can help researcher’s better grasp the relationship between the picture owner's affect concepts and the evoked viewer's, allowing for new user-centric applications. Personalization is aided by user profile, which is critical in the field of recommendation systems.

Time:-

Despite the fact that practically all of the aforementioned works overlook it, the feeling created by an image can alter with time. This sentiment component can be overlooked most of the time, yet it is critical in certain situations. An image of the World Trade Center, for example, is likely to generate distinct emotions depending on whether it is exhibited before or after 9/11. Despite the fact that there are no works on Picture Sentiment Analysis that evaluate changes in image sentiments over time due to the task's specialization and a lack of image datasets, there are some works that utilize image analysis over time for specialized cognitive and psychology applications. For example, in (Reece and Danforth, 2017), a statistical framework was used to predict depression by examining the sequence of Instagram photographs. The outcomes of this study imply that individual psychological differences are reflected in how people use social media, and that these differences may be computationally recognized by looking at a user's posting history. The authors of (Khosla et al., 2012) investigated which items and parts of an image are favorably or negatively connected with memorability, allowing memorability maps to be created for each image.

Challenges:-

So far, we've examined the present state-of-the-art in Visual Sentient Analysis, explaining the difficulties at hand as well as the many methodologies and features that have been used. This section seeks to provide some new difficulties and strategies that can be explored.

Popularity:-

Social marketing initiatives are one of the most prominent application fields for Visual Sentiment Analysis. Several companies are interested in analyzing the level of people involvement with respect to social posts connected to their products in the context of social media communication. This can be determined by counting the number of times a post has been viewed, liked, or shared, or by analyzing the comments. These data can be linked with statistics from web search engines and firms' websites to uncover correlations between social advertising efforts and their intended goals (e.g. brand reputation, website/store visits, product dissemination and sale, and so on). [36–38].

An image's popularity is a difficult quantity to define, let alone measure or infer. Humans, on the other hand, can estimate what visual content other people will enjoy in given situations (e.g. marketing campaigns, professional photography). This implies that there are some universally appealing elements in visuals. Researchers have been trying to figure out what characteristics make an image popular thus far. The study in [21] looks at the impact of 16 features on image popularity prediction. Image context, image content (i.e. image content provided by scene, face, and dominant color detectors), user context, and text features were all taken into account (i.e. image tags). By dividing the dataset into photos with high and low popularity values, the authors turned the task into a binary classification problem. The authors used the number of views and comments as a measure of popularity. Their research shows that comments are more predictable than views, and hence comments are more closely linked to the researched characteristics. The experimental results reveal that for both comments and views counts classifications, the accuracy values achieved by just considering textual characteristics (i.e. tags) outperform the performances of the classification based on other features and combinations of them.

The ACM Multimedia 2017 SMP Challenge collected the Social Media Prediction (SMP) dataset, which is a large-scale collection of social posts. This dataset contains about 850 thousand posts and 80 thousand users, as well as photographs from VSO [26] and personal user albums [39-41]. The authors wanted to capture the dynamic variance of social media data in particular. Indeed, the social media posts in the dataset are retrieved using temporal information (i.e. posts sequentiality) in order to preserve post sequence continuity. There are two challenges that have been proposed:

• Prediction of popularity: the aim is to predict the popularity of a given image shared by a specific user based on a popularity measure set for the specific social platform.

• Top prediction for tomorrow: given a set of images and data linked to previous photo sharing history, the challenge is to predict the top-n popular posts (i.e. ranking problem over a set of social posts) on the social media platform for the following day.

Unique picture ids (pids) and associated user ids are included in the SMP collection (uid). Almost all of the user and photo-related data accessible on Flickr may be extracted using this information. The SMP dataset aided in the development of time-aware popularity prediction methods, which employ time information to construct new image representation spaces from which the picture popularity score may be inferred at a specific moment or over a range of time scales.

Image Virality:-

Predicting the level of virality of a picture is a new task that is strongly tied to image popularity. The quality of a visual content (i.e. photos or videos) to be quickly and widely distributed on social networks is characterized as image virality [42]. In contrast to popularity, the virality score considers the number of times an image has been resubmitted by various users. As a result, photographs that gain popularity after being posted but not re-posted are not deemed viral [43]. These sometimes feature photos whose substance is unimportant in and of itself, but which are linked to current events that drew attention to the image at a given time, such as a flash news storey or a tragedy. The research in [42] was aimed at determining the impact of picture portions on virality. The authors provided a method for simultaneously detecting and localizing virality in photos, in particular. The detection entails calculating the virality score of an image. The goal of localization is to figure out which parts of a picture are responsible for making it viral, allowing for the creation of a heatmap that highlights the relevant parts of the input image.

Common Sense:-

We also need to encode the 'affective common-sense' in order to close the affective and cognitive gap between images and the sentiments they communicate. An automatic system that analyses image semantics can classify a Halloween picture as a negative image; nevertheless, knowledge about the context (i.e. Halloween) should impact the semantic ideas provided by the picture, and hence its interpretation. In the topic of knowledge representation, which is a sub-field of Artificial Intelligence, this relates to the 'common-sense knowledge problem.' Clearly, such an intelligent programme requires a representation of the knowledge in addition to inferential capabilities. Agrawal et al. [44] used contextual information to determine the sentiment of text after noting how difficult it is to design a Sentiment Analysis system that can be used in any situation with accurate categorization prediction.

Emoticon/ Emoji:-

In this section, we look at how text ideograms, such as emoticons and emoji, can be used to do Sentiment Analysis on both visual and textual information. A face expression is represented by an emoticon, which is a textual shorthand. The emoticons were created to allow writers to communicate their feelings and emotions in response to a text message. It aids in expressing the precise intent of a text sentence and thus improves message comprehension. In written interactions, emoticons are employed to mimic visual cues with the goal of expressing or explicitly clarifying the writer's sentiment. Indeed, visual signals such as facial expressions, posture, and gestures can be used to infer sentiment in real-life dialogues. Visual clues, on the other hand, are absent in text-based dialogues.

The authors of [46] wanted to know if emoticons may help with the textual Sentiment Analysis assignment as well. They looked into the role of emoticons in transmitting emotion, as well as how they may be used in the field of Sentiment Analysis. Weather, celebrations, cuisine, animals, emotions, feelings, and activities, as well as a wide range of facial expressions, are all represented by emojis. They were created with the goal of allowing for more expressive statements. Emojis have gained a lot of traction on social networking platforms and in instant messaging systems. For example, Instagram stated in March 2015 that emoticons appear in nearly half of the texts on its site [47].

Summary:-

The key topics and strategies connected to Visual Sentiment Analysis have been explored in this work. The present state-of-the-art has been thoroughly examined, with merits and cons for each approach and dataset highlighted. Despite the fact that this work has been studied for many years, the area is still in its early stages. Due to a variety of issues addressed in this study, visual sentiment analysis is a difficult undertaking. The findings addressed in this article, such as [24], concur that the semantic content of a picture has a significant impact on its emotional impact. Images with comparable color histograms and textures may elicit very varied emotional responses. As a result, an image representation that expresses both the overall appearance of the image as well as the intrinsic semantics of the seen scene is required. Earlier Visual Sentiment Analysis approaches attempted to bridge the so-called affective divide by creating visual representations. Other approaches use conventional Sentiment Analysis systems that work on textual contents to compute the polarity of the text associated with the photos (e.g. post messages, tags, and comments) and then try to learn ML systems that can infer that polarity from the associated visual content [23, 45].

The goal of this work was to provide a comprehensive review of the Visual Sentiment Analysis topic, related challenges, and state-of-the-art techniques. Relevant points with real business applications that would benefit from Sentiment Analysis on visual content studies have also been explored. This survey covers all areas of Visual Sentiment Analysis and includes important references, datasets, and results for each of them. It also offers a critical viewpoint on each topic under consideration. The paper is presented as a structured and critical review of previous works, as well as a description of the available datasets, features, and techniques, with the goal of serving as a starting point for researchers interested in tackling tasks in the field of Visual Sentiment Analysis and related challenges. It also seeks to pave the way for new solutions by recommending fresh and innovative methodologies and information sources that might be used to solve the problem. Although several methodologies and psychological foundations (e.g. emotional models) are frequently applied to both images and videos, the current work is primarily focused on still image analysis. Future developments of this paper could focus on a review of the state-of-the-art in video-based Visual Sentiment Analysis.

Reference:-

[1] Stuti Jindal and Sanjay Singh “Image Sentiment Analysis using Deep Convolution Neural Networks with Domain Specific Fine Tuning”, International Conference on Information Processing (ICIP) 2015.

[2] Jyoti Islam, Yanqing Zhang “Visual Sentiment Analysis for Social Images Using Transfer Learning Approach” IEEE International Conferences on Big Data and Cloud Computing 2016.

[3] Malhar Anjaria, Ram Mahana Reddy Guddeti “Influence Factor Based Opinion Mining of Twitter Data Using Supervised Learning” 978-1-4799-3635-9/14 2014 IEEE.

[4] Y. Wang, Y. Hu, S. Kambhampati, and B. Li, “Inferring sentiment from web images with joint inference on visual and social cues: A regulated matrix factorization approach.” in ICWSM, 2015, p. 21.

[5] Yuhai Yu1, Hongfei Lin, Jiana Meng,† and Zhehuan Zhao “Visual and Textual Sentiment Analysis of a Micro blog Using Deep Convolution Neural Networks.”

[6] Q. You, J. Luo, H. Jin, and J. Yang, “Robust image sentiment analysis using progressively trained and domain transferred deep networks,” 2015.

[7] Y. Yang, J. Jia, S. Zhang, B. Wu, Q. Chen, J. Li, and J. Tang, “How do your friends on social media disclose your emotions?,” in Proc. AAAI Conf. Artificial Intelligence (AAAI), 2014, pp. 306–312.

[8] Frome, G. S. Corrado, J. Shlens, S. Bengio, J. Dean, M. A. Ranzato, and T. Mikolov, “DeViSE: A deep visual-semantic embedding model,” in Proc. Advances in Neural Information Processing Systems (NIPS), 2013, pp. 2121–2129.

[9] Yilin Wang and Baoxin Li “Sentiment Analysis for Social Media Images.”

[10] Alameda-Pineda, X., Pilzer, A., Xu, D., Sebe, N., and Ricci, E. (2017). Viraliency: Pooling local virality. In 2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pages 484–492.

[11] Altwaijry, H. and Belongie, S. (2013). Relative ranking of facial attractiveness. In IEEE Workshop on Applications of Computer Vision (WACV).

[12] Battiato, S., Farinella, G. M., Milotta, F. L., Ortis, A., Ad- desso, L., Casella, A., D’Amico, V., and Torrisi, G. (2016). The social picture. In Proceedings of the 2016 ACM on International Conference on Multimedia Retrieval, pages 397–400. ACM.

[13] Borth, D., Ji, R., Chen, T., Breuel, T., and Chang, S.-F. (2013). Large-scale visual sentiment ontology and detectors using adjective noun pairs. In Proceedings of the 21st ACM international conference on Multimedia, pages 223–232. ACM.
[14] Bradley, M. M. (1994). Emotional memory: A dimensional analysis. Emotions: Essays on emotion theory, pages 97–134.
[15] Campos, V., Jou, B., and i Nieto, X. G. (2017). From pixels to sentiment: Fine-tuning CNNS for visual sentiment prediction. Image and Vision Computing, 65:15 – 22. Multimodal Sentiment Analysis and Mining in the Wild Image and Vision Computing.
[16] Campos, V., Salvador, A., Giro´-i Nieto, X., and Jou, B. (2015). Diving deep into sentiment: Understanding fine-tuned CNNS for visual sentiment prediction. In Proceedings of the 1st International Workshop on Affect & Sentiment in Multimedia, ASM '15, pages 57–62, New York, NY, USA. ACM.
[17] Colombo, C., Del Bimbo, A., Pala, P.: 'Semantics in visual information retrieval', IEEE Multimedia, 1999, 6, (3), pp. 38–53.
[18] Schmidt, S., Stock, W.G.: 'Collective indexing of emotions in images. a study in emotional information retrieval', J. Am. Soc. Inf. Sci. Technol., 2009, 60, (5), pp. 863–876.
[19] Wei-ning, W., Ying-lin, Y., Sheng-ming, J.: 'Image retrieval by emotional semantics: a study of emotional space and feature extraction', IEEE Int. Conf. on Systems, Man and Cybernetics, Taipei, Taiwan, 2006, vol. 4, pp. 3534– 3539.
[20] Zhao, S., Yao, H., Yang, Y., et al.: 'Affective image retrieval via multi-graph learning', Proc. of the 22nd ACM Int. Conf. on Multimedia, Orlando, FL, 2014, pp. 1025–1028.
[21] McParlane, P.J., Moshfeghi, Y., Jose, J.M.: 'Nobody comes here anymore, it's too crowded; predicting image popularity on flickr', Proc. of Int. Conf. on Multimedia Retrieval, Glasgow, UK, 2014, p. 385.
[22] Siersdorfer, S., Minack, E., Deng, F., et al.: 'Analyzing and predicting sentiment of images on the social web', Proc. of the 18th ACM Int. Conf. on Multimedia, Firenze, Italy, 2010, pp. 715–718.
[23] Esuli, A., Sebastiani, F.: 'Sentiwordnet: a publicly available lexical resource for opinion mining', Proc. of LREC, Genoa, Italy, 2006, vol.6, pp. 417–422.
[24] Machajdik, J., Hanbury, A.: 'Affective image classification using features inspired by psychology and art theory', Proc. of the 18th ACM Int. Conf. on Multimedia, Firenze, Italy, 2010, pp. 83–92.
[25] Mikels, J.A., Fredrickson, B.L., Larkin, G.R., et al.: 'Emotional category data on images from the international affective picture system', Beh. Res. Meth., 2005, 37, (4), pp. 626–630.
[26] Borth, D., Ji, R., Chen, T., et al.: 'Large-scale visual sentiment ontology and detectors using adjective noun pairs', Proc. of the 21st ACM Int. Conf. on Multimedia, Barcelona, Spain, 2013, pp. 223–232.
[27] Peng, K.-C., Chen, T., Sadovnik, A., et al.: 'A mixed bag of emotions: model, predict, and transfer emotion distributions', Proc. of the IEEE Conf. on Computer Vision and Pattern Recognition, Boston, MA, USA, 2015, pp. 860– 868.
[28] Campos, V., Salvador, A., Giró-i Nieto, X., et al.: 'Diving deep into sentiment: understanding fine-tuned cnns for visual sentiment prediction', Proc. of the 1st Int. Workshop on Affect & Sentiment in Multimedia (ASM '15), New York, NY, USA, 2015, pp. 57–62.
[29] Katsurai, M., Satoh, S.: 'Image sentiment analysis using latent correlations among visual, textual, and sentiment views', IEEE Int. Conf. on Acoustics, Speech and Signal Processing, Shanghai, China, 2016, pp. 2837–2841.
[30] Yang, J., Sun, M., Sun, X.: 'Learning visual sentiment distributions via augmented conditional probability neural network', AAAI, San Francisco, CA, USA, 2017, pp. 224–230.
[31] Jufeng Yang, M.S., She, D.: 'Joint image emotion classification and distribution learning via deep convolutional neural network', Proc. of the Twenty-Sixth Int. Joint Conf. on Artificial Intelligence, (IJCAI-17), Melbourne, Australia, 2017, pp. 3266–3272.
[32] Campos, V., Jou, B., Nieto, X.G.i.: 'From pixels to sentiment: fine-tuning CNNS for visual sentiment prediction', Image Vis. Comput., 2017, 65, pp. 15– 22, Multimodal Sentiment Analysis and Mining in the Wild Image and Vision Computing.
[33] Zhao, S., Ding, G., Huang, Q., et al.: 'Affective image content analysis: a comprehensive survey', Int. Joint Conf. on Artificial Intelligence, 2018, pp. 5534–5541.
[34] Soleymani, M., Garcia, D., Jou, B., et al.: 'A survey of multimodal sentiment analyses, Image Vis. Comput., 2017, 65, pp. 3–14.
[35] Li, Z., Fan, Y., Jiang, B., et al.: 'A survey on sentiment analysis and opinion mining for social multimedia', Multimedia Tools Appl., 2019, 78, (6), pp. 6939–6967.
[36] Goeldi, A.: Website network and advertisement analysis using analytic measurement of online social media

content, 5 July 2011. US Patent.
[37] Hanna, R., Rohm, A., Crittenden, V.L.: 'We're all connected: the power of the social media ecosystem', Bus. Horiz., 2011, 54, (3), pp. 265–273.
[38] Trusov, M., Bucklin, R.E., Pauwels, K.: 'Effects of word-of-mouth versus traditional marketing: findings from an internet social networking site', J. Mark., 2009, 73, (5), pp. 90–102.
[39] Wu, B., Cheng, W.-H., Zhang, Y., et al.: 'Time matters: multi-scale temporalization of social media popularity', Proc. of the 2016 ACM on Multimedia Conf. (ACM MM), Amsterdam, The Netherlands, 2016.
[40] Wu, B., Cheng, W.-H., Zhang, Y., et al.: 'Sequential prediction of social media popularity with deep temporal context networks', Int. Joint Conf. on Artificial Intelligence (IJCAI), Melbourne, Australia, 2017.
[41] Wu, B., Mei, T., Cheng, W.-H., et al.: 'Unfolding temporal dynamics: predicting social media popularity using multi-scale temporal decomposition', Proc. of the Thirtieth AAAI Conf. on Artificial Intelligence (AAAI), Phoenix, AZ, USA, 2016.
[42] Alameda-Pineda, X., Pilzer, A., Xu, D., et al.: 'Viraliency: 'pooling local virality', 2017 IEEE Conf. on Computer Vision and Pattern Recognition (CVPR), Honolulu, HI, USA, July 2017, pp. 484–492.
[43] Deza, A., Parikh, D.: 'Understanding image virality', 2015 IEEE Conf. on Computer Vision and Pattern Recognition (CVPR), Boston, MA, USA, June 2015, pp. 1818–1826.
[44] Agarwal, B., Mittal, N., Bansal, P., et al.: 'Sentiment analysis using common- sense and context information', Comput. Intell. Neurosci., 2015.
[45] Wilson, T., Wiebe, J., Hoffmann, P.: 'Recognizing contextual polarity in phrase-level sentiment analysis', Proc. of the Conf. on Human Language Technology and Empirical Methods in Natural Language Processing, 2005, pp. 347–354.
[46] Hogenboom, A., Bal, D., Frasincar, F., et al.: 'Exploiting emoticons in sentiment analysis', Proc. of the 28th Annual ACM Symp. on Applied Computing, Las Vegas, NV, USA, 2013, pp. 703–710.
[47] Dimson, T.: Emojineering part 1: Machine learning for emoji trends. https:// instagram-engineering.com/emojineering-part-1-machine-learning-for-emoji-trendsmachine-learning-for-emoji-trends-7f5f9cb979ad, 2015.