JETIR.ORG ISSN: 2349-5162 | ESTD Year : 2014 | Monthly Issue
JOURNAL OF EMERGING TECHNOLOGIES AND INNOVATIVE RESEARCH (JETIR)
An International Scholarly Open Access, Peer-reviewed, Refereed Journal

A SURVEY ON MENTAL HEALTH STATE DETECTION USING OPENCV AND SENTIMENTAL ANALYSIS

[1]Krishna A.K, [2]Krishnaprasad.K, [3]Sushank, [4]Vaishak U.A

[5]Ms Pratheeksha B.R

[1,2,3,4,5], Department Of Computer Science And Business System.
[1,2,3,4,5], Srinivas Institute Of Technology Mangalore, India.

Abstract: This project's approach to mental health detection through facial expression and sentiment analysis is remarkably simple and yet very effective. Our system uses OpenCV to do facial emotion detection in real-time, and employs pre-trained models for text-based sentiment analysis. By analyzing both visual and textual components, our system strives to achieve a fuller picture of a person's emotional state.

The webcam emotion recognition system comprises of facial expression recognition modules for happy, sad, angry, and fearful emotion detection which use Haarcascades and CNN-based approaches. At the same time, the sentiment analysis module uses VADER or TextBlob on user input texts to analyze their emotional value (positivity, negativity, or neutrality). Combination of multiple techniques increases accuracy. The results generated from previous modules can be integrated using weighted scoring or logic into one system allowing users to attain value beyond what is possible through a single method.

The focus of our work was making a simple-to-use prototype with an intuitive design and exploring the use of multimodal emotion recognition technologies aimed at early mental health detection. Our study affirmed Random Forest models for text-based emotion classification. Basic emotions in facial expressions are recognized with reliable accuracy, according to prior research. Despite prior research showing limited correlation between emotions and facial expressions, we attribute the initial deviations from expected outcomes to basic emotions in fvace detection algorithms.

Keywords: Mental Health Detection, Facial Emotion Recognition, Sentiment Analysis, OpenCV, NLP, Multimodal Emotion Recognition, Affective Computing.

I. INTRODUCTION:

Mental health is an important component of general health and social well-being of an individual. The World Health Organization (WHO) has also described mental health as a state of well-being where an individual is capable of fulfilling the usual demands of life, being productive in their occupation, and actively engaging with their communities. However, shared social, economic, and environmental determinants have led to a high and increasing burden of mental disorders all over the world. Depression, in specific, affects nearly 350 million people and accounts for a significant proportion of the total disease burden [1][2].

Traditional diagnostic techniques i.e., self-report questionnaires, in-person interviews, and clinical assessments are generally time-consuming, subjective, and patient disclosure-based [3]. In response, researchers have looked for new, automated, and non-invasive means of identifying mental illness at an early stage. One new developing area uses visual and language signals i.e., facial expression, color use in drawings, and text sentiment to assess psychological status. Color-emotion correspondences are highly documented in the literature. For instance, the Manchester Color Wheel study established that happiness was linked to the color yellow and sadness to the color blue [4]. Jonauskaite et al. also demonstrated that achromatic colors tend to represent negative emotions, while bright colors such as yellow and green represent pleasure and tranquility [5]. Similarly, photographs and paintings, especially in the case of non-artists, can be employed as representations of the mental state of an individual, conveying emotions through the number of strokes, space allocation, and the color employed [6][7]. Coupled with visual cues, text sentiment analysis, especially from online social networks (OSNs), has become an effective tool in measuring mental well-being. Microblogging services like Twitter and Facebook are constantly being used to express oneself, thus yielding immense real-time data amenable for analysis. Existing research has harnessed the use of machine learning and deep learning models for feature extraction in posts and psychopathological diagnosis such as stress, anxiety, and suicidal intent [8][9][10]. Despite vast improvements, the issues persisting include model explainability, data cultural bias, and multimodal fusion needs.

Advances in technology, especially in computer vision and natural language processing, now make it possible to recognize emotions more accurately. OpenCV, as a prominent open-source computer vision library, has been extensively used to analyze facial features and expressions, while sentiment analysis algorithms are utilized to analyze text inputs to classify emotions [11].

When the two are employed together, they offer an integrated concept of an individual's mental state crossing the emotional indications evident in facial data and the psychological indications evident in language. In light of the above discussion, this survey explores contemporary techniques in mental illness detection via OpenCV-based facial and drawing-based emotion recognition, and sentiment analysis for text-based emotion detection. It attempts to explore the extent to which visual aesthetics, color perception, and language usage individually or together can indicate underlying psychological states. Furthermore, the paper explores the effects of age, cultural setting, and technological limitations, and identifies promising research directions for integrated real-time mental health diagnosis.

II. RELATED WORK:

Okay, so here's the lowdown on what's been going on in this area: Tons of nerds (with money, grant proposals, and very little chill) have been geeking out about AI and emotion detection, especially when it comes to figuring out mental health stuff. Way, way back, Ekman and Friesen were basically the Ogs like, the grandpas of "let's look at your face and call it science." They proved that most facial emotions are kind of universal, which still sets the ground rules for all these facial recognition gizmos today.

Fast forward a bit, and suddenly you've got machine learning blowing up. SVMs, CNNs, even LSTMs if it's an acronym, someone's probably built an emotion detector with it. These models seriously stepped up the game, pulling emotions straight off your mugshot way better than whatever clunky code we had before.

But hey, not everyone likes having a camera pointed at their face 24/7, so some clever folks started poking around with text, too. You know VADER and TextBlob? Just your average vocab nerds for hire. They were like, "Let's see if we can spot joy or existential despair in your tweets." Then came the big boys BERT, RoBERTa, XLNet—whose literal job is to read between the lines, even when you're being super vague or moody online. People like Chancellor and De Choudhury basically spent years trawling through Twitter and Reddit, hunting for clues of depression or suicidal feelings, and turns out, it's not all memes and cat videos out there. Some really solid research happened.

And, plot twist, why not just throw everything into one big blender? That's what the cool crowd's doing now—mixing your texts, your face, and even your voice (oh yeah, that's a thing) to guess how you're really holding up. The DAIC-WOZ dataset? That's practically the Holy Grail for these fusion-things. Every study with two or more data types screams about how much better they are than the old "one-trick pony" models. Apparently, if you want to know if someone's depressed, why settle for face or text—grab both (and hey, toss in voice for good measure). Game changer, honestly.

Chatbots for Mental Health Support:

In recent years, the use of AI-driven chatbots in mental health support has grown significantly, offering users accessible and anonymous platforms for psychological assistance. Many of these systems incorporate principles from Cognitive Behavioural Therapy (CBT), enabling scalable delivery of mental health interventions.

Woebot:

Woebot is a leading example that applies CBT principles to help individuals manage symptoms of depression and anxiety. Designed for brief, daily interactions, Woebot fosters a habit of mental health engagement. Through its conversational interface, users are encouraged to reflect on their behaviours and emotions, thereby enhancing emotional resilience.

Wysa:

Wysa, promoted as *"the AI chatbot you hear"*, provides a secure, non-judgmental space for users to explore their thoughts and emotions. With integrated self-help tools, it guides users in reframing negative thoughts, facilitating cognitive restructuring. Wysa emphasizes privacy and safety, with transparent emergency protocols that include contacting emergency services when necessary.

Youper:

Youper uniquely combines CBT with positive psychology to deliver personalized mental health support. It leverages research-backed psychological interventions and adaptive AI to enhance emotional intelligence and provide structured assistance for managing anxiety and depression. Youper is often recognized as a pioneer in the field of digital therapies, offering a guided and dynamic therapeutic experience.

These chatbots exemplify the expanding landscape of AI-based mental healthcare, where conversational agents deliver timely support, promote self-awareness, and serve as valuable supplements to traditional therapeutic approaches.

Emotion Detection Techniques:

Facial emotion recognition plays a critical role in automated mental health assessment by interpreting non-verbal cues that are essential for psychological evaluation. A range of techniques has been explored in the literature, each evaluated based on performance metrics such as accuracy, computational efficiency, and robustness across different datasets and conditions. Among the approaches studied, the combination of Histogram of Oriented Gradients (HOG) with Convolutional Neural Networks (CNNs) has shown promising results. HOG effectively captures facial edge and gradient structures, which are crucial for expression analysis, while CNNs provide powerful deep learning capabilities for feature extraction and classification.

This HOG + CNN approach enables real-time detection of emotional states such as sadness, anger, happiness, and fear. Its effectiveness and efficiency make it a suitable choice for mental health detection systems that rely on facial expression analysis to infer psychological states.

Module	Approaches	Comments	Limitations
Facial Emotion Detection	3D motion-based feature + Hidden Markov Model (HMM).	1) Features are extracted and classified using Gentle boost classifier such as AdaBoost and the result is used to build temporal models of each expression using an HMM. 2) The obtained results show that use of 3D information enhances the accuracy when compared to 2D data.	As HMM is a generative model, it has a problem of Local Optimal solution Trap.
	HMM + SVM	1) The Limitation of the HMM model was improved by hybridizing a discriminant classifier. 2) HMM was hybridized to capture the temporal dynamics and employed SVM as a multiclass classifier of the features on a frame by using frame basis.	It performed better on binary problems rather than multiclass problems like facial emotion recognition.
	Local Binary Pattern (LBP) + Linear Programming	1) It's benefit lies in its simplicity in computation, high tolerance for low image resolution, invariant to illumination changes. 2) It improved the accuracy for the JAFFE database but it could not be generalized as it was implemented only on one database.	It is challenged with factors like rotation, increase in computational complexity, small sample size and limited information representation.
	Hand crafted Preprocessing stages + CNN	1) It achieved a better accuracy than the other approaches mentioned above because of enhancing preprocessing stages like resizing, face detection, cropping, adding noises, data normalization which consist of local, global contrast and histogram equalization. 2) It was trained on CK (posed), JAFFE (posed) and MUG (posed) database and achieved an average accuracy of 97.06%.	It tends to over fit because of overly significance to preprocessing stages.
	Histogram of Oriented Gradient (HOG) + CNN	1) It extracted features from the active facial patches and fed them to the convolutional neural network. 2) It gives better performance than other models and it showed an average accuracy of 95%.	It increases the time of computation as it is computationally intensive.
	Hybrid (LBP and HOG) + SoftMax	1) The fusion of extracted LBP and HOG features reduced the extracted features dimensionality with PCA. 2) The fused feature was permuted on several classifiers and achieved a maximum accuracy of 98.3% using SoftMax classifier. 3) The result is evidence that hybrid features could enhance the performance significantly.	Research is still going on in finding the best combination of features.

III. FACIAL EXPRESSION ANALYSIS USING OPENCV:

Alright, here's the deal with facial expression analysis using OpenCV, let's skip the boring, textbook vibe and talk straight. You ever try to read someone's mood and just… totally miss? Yeah, computers have that problem too, but they're getting better at it, thanks to stuff like OpenCV. Super handy for catching all those sneaky emotions people try to hide, especially since folks are pretty bad at masking their faces when they're not actively trying.

So, the usual move? First you get the computer to find a face in the crowd, right? Old-school methods like Haar cascades or those geeky Dlib landmarks work fast and don't ask for much in terms of hardware, so you could probably run this on your grandpa's basement laptop if you had to.

Once a face is locked in, here come the big guns: convolutional neural networks (CNNs) trained on things like FER-2013 or AffectNet datasets. Basically, people taught these networks to tell when someone looks happy, sad, angry, scared, shocked, grossed out, or just… meh (a.k.a. neutral). And, just so you know, these expressions actually tie back to mental health stuff. Like, if someone's stuck on "sad face" all the time, might be a red flag for depression. Jumping at shadows? Maybe anxiety. You get the picture.

Now, for the cool, nerdy upgrades folks are messing around with 3D CNNs and models that can even pick up on micro-expressions. You know, those brief flashes of real emotion when someone's trying not to react. But honestly, it's not perfect. Throw in some bad lighting, a hand in front of the mouth, or just someone from a different culture who doesn't make the same faces, and the tech starts to wobble.

Even so, people love OpenCV for this stuff, especially in mental health apps and virtual doctor visits. It's quick, does the job in real time, and gives therapists an extra tool when they can't sit face-to-face with someone. Not perfect, but definitely cooler than guessing based on a few awkward Zoom calls.

IV. ARCHITECTURE:

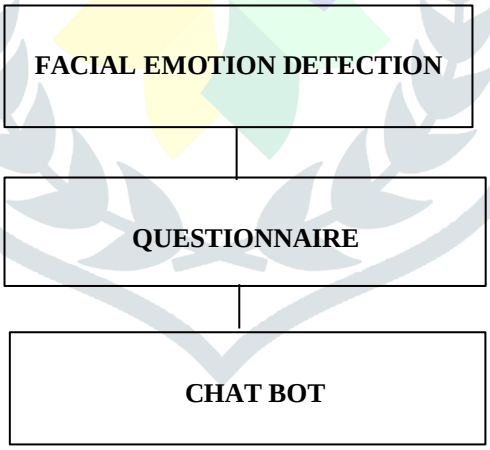

Block Diagram

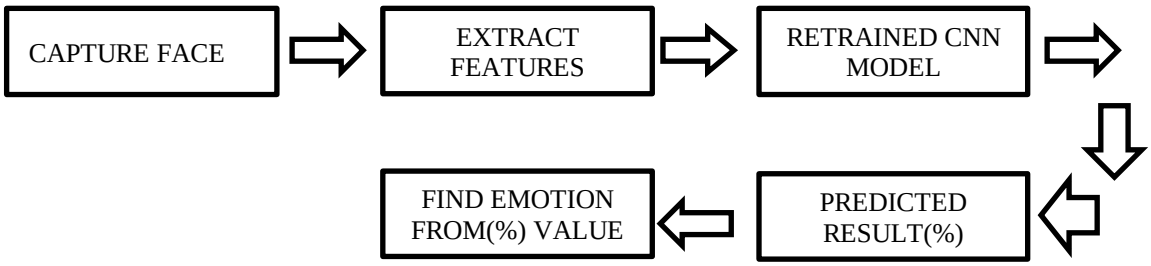

Facial Emotions

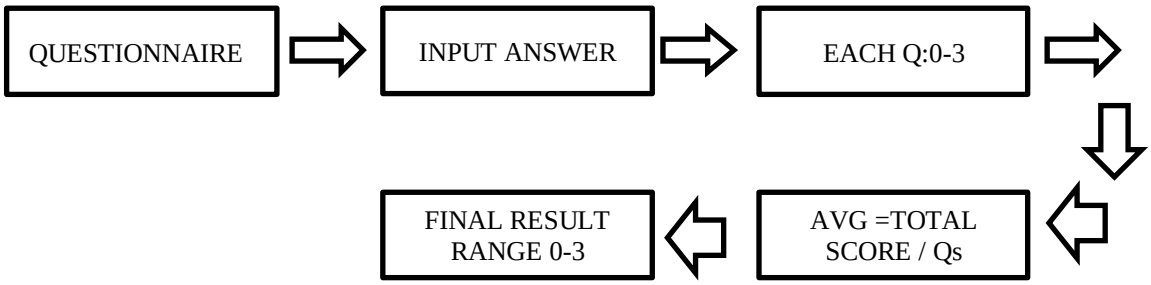

Questionnaire

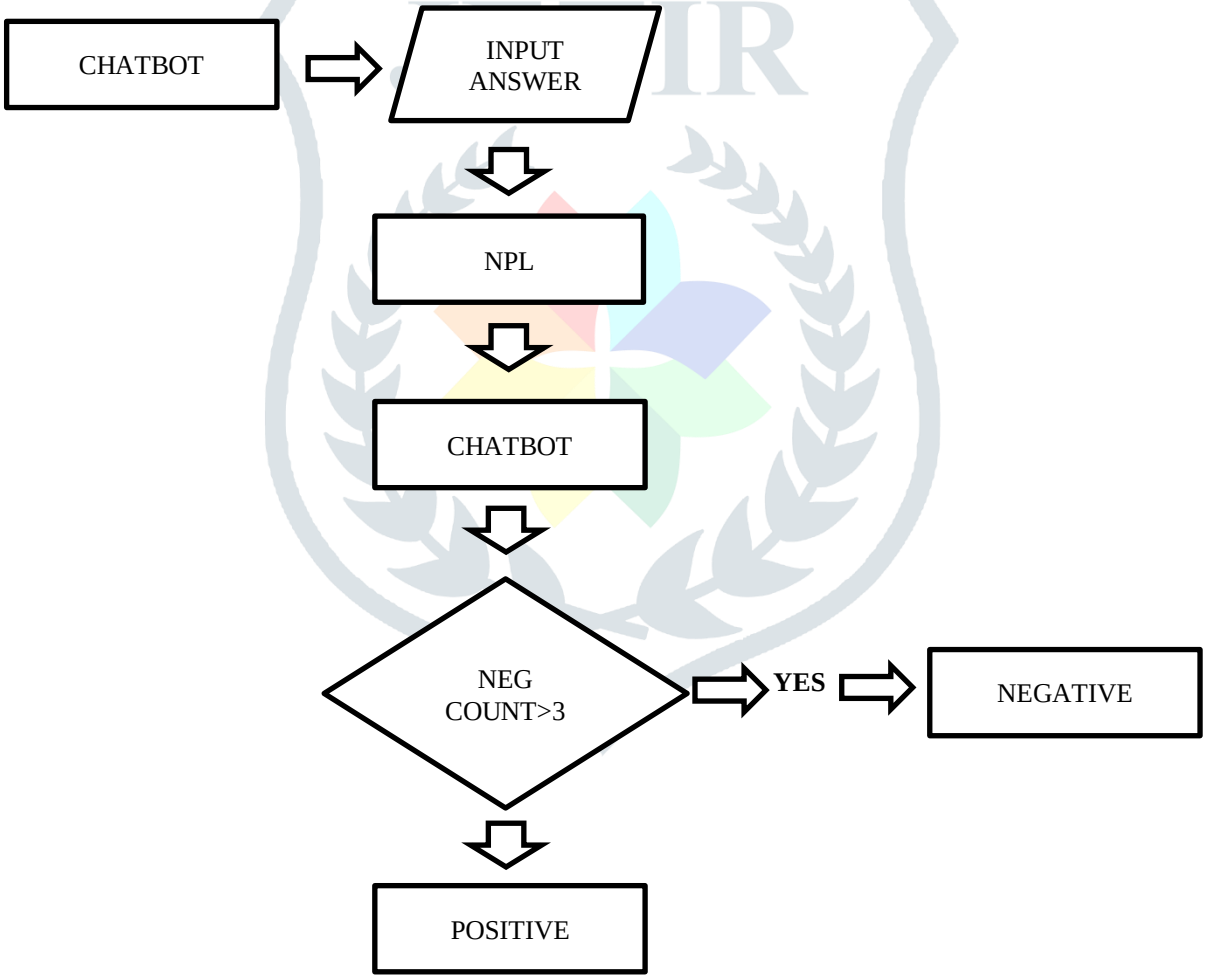

CHATBOT

V. MULTIMODAL MENTAL HEALTH STATE DETECTION:

Alright, let's get real for a second sentiment analysis isn't just about checking if someone's ranting or raving online. It's more like eavesdropping on the emotional static woven through their words, picking up on vibes that scream anxiety or sadness even when the surface seems chill. Turns out, psychology folks have cracked this wide-open people struggling with depression or anxiety toss out all kinds of tiny language breadcrumbs: way too many sad words, me-me-me stuff, or weirdly extreme phrases like "always" and "never." Honestly, it's like their texts light up in neon if you know what to look for. Now, classic tools like VADER or TextBlob? They're pretty decent if you just wanna know "good mood, bad mood, meh," but they can totally trip up when the feelings get messy like, who can blame 'em? Life is messy! So, researchers have souped up the old models, using stuff like CLPsych datasets and actual machine learning magic. These beefed-up systems can sniff out joy, fear, rage, not just your basic happy-sad binary. It's like going from black-and-white TV to full-on IMAX. And, oh man, the real MVPs? Deep learning beasts like LSTMs, GRUs, and Transformers (no, not the ones that turn into trucks). These models actually "remember" context so if someone's been cranking out dark or hopeless posts every night for two months, the system picks up on the pattern. They're practically therapists with better memory. Mash up all this semantic and syntactic wizardry, and suddenly the software gets deep insight into someone's headspace. Kind of mind-blowing, actually, and while a little creepy it's a huge win for spotting folks who need a hand before things really spiral...

VI. CHALLENGES AND LIMITATIONS:

Ugh, where do I even start? These fancy new multimodal mental health detectors sound cool on paper, but in real life? It's kind of a mess. First off, privacy yeah, it actually matters. Who wants their cringiest, saddest selfies and angsty texts floating around some server? Not me. It's extra tricky too, 'cause you gotta get this whole informed consent jazz right, stash the data somewhere safe, and not break a million ethical rules in the process. Talk about a headache.

Then there's the whole "one size fits all" myth. Sorry, but emotion is super messy. What works for some college kid in New York won't work for, I dunno, a grandma in Tokyo. We're all out here showing feelings different ways. So if you train your little robot-brain model on, like, a tiny sample of people, it's gonna flop in the wild. Guaranteed.

And don't even get me started on how easily these systems can be faked out. Like, how often do people smile when they're actually miserable? Or toss sarcasm around till no one knows what's real anymore? Good luck making a deep learning model pick that apart. Spoiler: it probably can't. Also, these cutting-edge algorithms suck up more computer power than a NASA launch, so if you want this thing running live, hope you've got a supercomputer handy. Most folks just don't.

Oh, before I forget: none of this is even properly proven in clinics yet. That's right lots of hype, not nearly enough real-world testing. You wouldn't want your therapist trusting your diagnosis to a half-baked chatbot, would you? Yeah, me neither.

VII. FUTURE DIRECTIONS:

What's next? Well, honestly, if AI's gonna be useful for mental health stuff, it desperately needs to get better at a few things. For starters, let's quit acting like everyone speaks English or shares the same cultural quirks. We need models that actually "get" people from all over otherwise, you're missing the point. Think about all those little signals we give off: not just what we type, but how we say it, how we type, the rhythm, the weird late-night message patterns. That juicy, real-life context right now, AI basically misses most of it.

Oh, and don't even get me started on black-box algorithms. Therapists wouldn't give you advice if they couldn't explain themselves, so why should an AI? Let's make these things less mysterious bring in some actual explainable AI, so people (and, y'know, doctors) can understand why it spits out certain results.

Personalized systems are the real MVPs here, too. Instead of just guessing at "average sadness," AI should learn your vibe over time. If you're generally chill but suddenly go full-on emoji-sad-face for a week, that's what we care about—not generic data points. And in the end, this all needs to hook into real health systems, like those clunky health record portals and whatever your therapist uses to Zoom-call you. Otherwise, it's just another flashy tech demo nobody trusts. Bottom line: we need AI tools that actually make sense for real people, in real life, not just on paper.

VIII. CONCLUSION:

Honestly, it feels like we're living in sci-fi with all these fancy tools trying to read your mind through your face or text messages. But let's be real multimodal mental health systems have a bunch of hurdles to jump. First off, privacy? Huge red flag. We're talking about seriously personal stuff here, and most folks probably aren't lining up to hand over their private messages or videos just so an algorithm can analyze them. Getting actual consent, keeping everything locked down, and following the endless rules? Not exactly easy-peasy.

And don't even get me started on "generalizability." Just because a model works with one group doesn't mean it won't totally flop with another. Emotions aren't one-size-fits-all. Cultural differences, age, gender everyone's got their own way of showing stuff. You can't just train an algorithm on, say, college kids in one town and expect it to nail the vibe for, like, retired folks in another country. Plus, let's admit it humans are confusing as hell. Someone grins when they're anxious, throws in a sarcastic "I'm fine" that fools the computer, and suddenly your shiny AI model is toast. Machines don't get irony, not really. Oh, and then there's the nerdy, techy bit real-time systems? They chug power like it's going out of style. All those cutting-edge deep learning models gobble up resources, and most hospitals or clinics just don't have a NASA-level setup running in the back office.

But wait, here's the kicker: most of these AI tools haven't even been properly tested in real clinics. Health pros want proof, not just some tech hype. Without that clinical rubber stamp, who's actually going to trust these things in the wild? So, yeah, cool tech, lots of potential, but we're not exactly ready to hand over the keys just yet..

IX. REFERENCES:

[1] A. B. R. Shatte, D. M. Hutchinson, and S. J. Teague, "Machine learning in mental health: A scoping review of methods and applications," Psychological Medicine, vol.49, pp.14261448, 2019.doi:10.1017/S0033291719000151.

[2] J. Gao, P. Zheng, Y. Jia, H. Chen, Y. Mao, S. Chen et al., "Mental health problems and social media exposure during COVID-19 outbreak," PLOS One, vol. 15, 2020. doi: 10.1371/journal.pone. 0231924..

[3] S. Hinduja, M. Afrin, S. Mistry, and A. Krishna, "Machine learning-based proactive social-sensor service for mental health monitoring using Twitter data," Int. J. Inf. Manag. Data Insights, vol. 2, p. 100113, 2022. doi: 10.1016/j.jjimei.2022.100113.

[4] S. Graham et al., "Artificial intelligence for mental health and mental illnesses: An overview," Curr. Psychiatry Rep., vol. 21, 2019. doi: 10.1007/s11920-019-1094-0.

[5] H. Herdiansyah, R. Roestam, R. Kuhon, and A. S. Santoso, "Their post tell the truth: Detecting social media users' mental health issues with sentiment analysis," Procedia Computer Science, vol. 216, pp. 971–978, 2022. doi: 10.1016/j.procs.2022.12.185.

[6] M. Wongkar and A. Angdresey, "Sentiment Analysis Using Naive Bayes Algorithm of the Data Crawler: Twitter," in Proc. Int. Conf. Informatics and Computing (ICIC), 2019, pp. 1–5. doi: 10.1109/ICIC47613.2019.8985884..

[7] A. R. Isnain, J. Supriyanto, and M. P. Kharisma, "Implementation of K-Nearest Neighbor (K-NN) Algorithm for Public Sentiment Analysis of Online Learning," Indonesian Journal of Computing and Cybernetics Systems, vol. 15, pp. 121–132, 2021. doi: 10.22146/ijccs.65176.

[8] J. Sharma and V. Tomer, "Depression detection using sentiment analysis of so cial media data," AIP Conference Proceedings, vol. 2481, p. 020044, 2022. doi: 10.1063/5.0104192.

[9] G. Thakkar, S. Hakimov, and M. Tadi´ c, "M2SA: Multimodal and Multilingual Model for Sentiment Analysis of Tweets," Procedia Computer Science, vol. 245, pp. 971–978, 2024.

[10] S. Elbagir and J. Yang, "Analysis Using Natural Language Toolkit and VADER Sentiment," Procedia Computer Science, 2024.

[11] D. M. Low et al., "Natural Language Processing Reveals Vulnerable Mental Health Support Groups and Heightened Health Anxiety on Reddit During COVID 19: Observational Study," J. Med. Internet Res., vol. 22, p. e22635, 2020. doi: 10.2196/22635.

[12] APPLICATION OF MACHINE LEARNING METHOD SINMENTAL HEALTH DETECTION. [Rohizah ABD Rahman, Khairuddin Omar,Shahrul Azman Mohd Noah(2020) https://ieeexplore.ieee.org/document/9214815

[13] SENTIMENTAL ANALYSIS AND AFFECTIVE COMPUTING FOR DEPRESSION MONITORING. Chiara zucco,Barbara calabrese ,Mario cannataro(2021). 2017 IEEE International Conference on Bioinformatics and Biomedicine (BIBM) [3] DEPRESSION DETECTION USINGCNN IEEE International Conference on Signal Processing, Information, Communication & Systems (SPICSCON).

[14] MULTI-MODE EDUCATION DATA FUSION FOR STUDENT MENTAL HEALTH DETECTION. IEEE Conference paper by T. Guo, W. Zhao, M. Alrashoud, A. Tolba, S. Firmin and F. Xia