











AI-Driven Early Mental Health Screening: Analyzing Selfies of Pregnant Women

Gustavo A. Basílio¹^a, Thiago B. Pereira¹^b, Alessandro L. Koerich²^c, Hermano Tavares³^d,
Ludmila Dias¹^e, Maria G. S. Teixeira⁴^f, Rafael T. Sousa⁵^g, Wilian H. Hisatugu⁴^h,
Amanda S. Mota³ⁱ, Anilton S. Garcia⁶, Marco Aurélio K. Galletta⁷^j and Thiago M. Paixão¹^k

¹Federal Institute of Espírito Santo (IFES), Campus Serra, Serra, Brazil

²École de Technologie Supérieure (ÉTS), Montreal, Canada

³Department of Psychiatry, University of São Paulo Medical School (FMUSP), São Paulo, Brazil

⁴Department of Computing and Electronics, Federal University of Espírito Santo (UFES), Campus São Mateus, São Mateus, Brazil

⁵Federal University of Mato Grosso (UFMT), Barra do Garças, Brazil

⁶Federal University of Espírito Santo (UFES), Campus Goiabeiras, Vitória, Brazil


⁷Department of Obstetrics and Gynecology, University of São Paulo Medical School (FMUSP), São Paulo, Brazil


Keywords: Mobile Health, Mental Health, Pregnancy Healthcare, Affective Computing, Facial Analysis, Convolutional Neural Networks, Visual-Language Models, Deep Learning.


Abstract: Major Depressive Disorder and anxiety disorders affect millions globally, contributing significantly to the burden of mental health issues. Early screening is crucial for effective intervention, as timely identification of mental health issues can significantly improve treatment outcomes. Artificial intelligence (AI) can be valuable for improving the screening of mental disorders, enabling early intervention and better treatment outcomes. AI-driven screening can leverage the analysis of multiple data sources, including facial features in digital images. However, existing methods often rely on controlled environments or specialized equipment, limiting their broad applicability. This study explores the potential of AI models for ubiquitous depression-anxiety screening given face-centric selfies. The investigation focuses on high-risk pregnant patients, a population that is particularly vulnerable to mental health issues. To cope with limited training data resulting from our clinical setup, pre-trained models were utilized in two different approaches: fine-tuning convolutional neural networks (CNNs) originally designed for facial expression recognition and employing vision-language models (VLMs) for zero-shot analysis of facial expressions. Experimental results indicate that the proposed VLM-based method significantly outperforms CNNs, achieving an accuracy of 77.6%. Although there is significant room for improvement, the results suggest that VLMs can be a promising approach for mental health screening.


1 INTRODUCTION


Major depressive disorder (depression) is a prevalent mental health disorder affecting over 280 million people globally and a leading cause of disability worldwide. It contributes significantly to the global disease burden (World Health Organization, 2023). Similarly, anxiety disorders, marked by excessive worry and fear, impact millions of lives, often leading to debilitating effects on daily functioning and quality of life. Particularly, there is a recognized vulnerability to mood disorders during pregnancy, with possible worsening of previous emotional conditions and the emergence of new altered mental states that increase


^a  <https://orcid.org/0009-0008-6329-9966>


^b  <https://orcid.org/0009-0009-6652-2913>


^c  <https://orcid.org/0000-0001-5879-7014>


^d  <https://orcid.org/0000-0002-6632-2745>


^e  <https://orcid.org/0009-0007-6486-6259>


^f  <https://orcid.org/0000-0003-0398-8029>

^g  <https://orcid.org/0000-0001-5998-046X>

^h  <https://orcid.org/0000-0001-8333-0539>

ⁱ  <https://orcid.org/0000-0003-1553-5297>

^j  <https://orcid.org/0000-0002-0016-1803>

^k  <https://orcid.org/0000-0003-1554-6834>

the risk of depression and anxiety (Biaggi et al., 2016). This vulnerability is certainly influenced by the significant hormonal changes of pregnancy. However, it is also important to recognize that this is a time of great physical, psycho-emotional, cultural, and social changes, generating psychic stress and increased anxiety.

Several studies have pointed to the association between depression and anxiety during pregnancy with unfavorable obstetric and neonatal outcomes. These disorders can increase the risk of obstetric complications such as cesarean delivery, preeclampsia, preterm birth, low birth weight, small for gestational age newborns, and newborns with low Apgar scores, indicating lower oxygenation at birth (Li et al., 2021; Nasreen et al., 2019; Dowse et al., 2020; Kurki et al., 2000). These findings underline the importance of making a correct and early diagnosis, allowing for appropriate psychiatric and psychological follow-up during pregnancy to improve both maternal and neonatal outcomes.

For early intervention and improved treatment outcomes, screening processes are essential for identifying individuals who may have conditions such as depression or anxiety disorders before they present symptoms or seek treatment (Thombs et al., 2023). Various methods can be employed for screening, including self-reported questionnaires, clinical interviews with mental health professionals, and observing behaviors and physical symptoms. Additionally, studies in affective computing, a field that explores the interaction between human emotions and computational systems, have demonstrated the potential of artificial intelligence (AI) in screening mental disorders (Kumar et al., 2024). AI-driven technologies can enhance screening by analyzing data from multiple sources, such as text inputs, voice recordings, or facial expressions (Liu et al., 2024c), which is the very focus of this work.

In this work, we address the challenge of ubiquitous depression-anxiety screening from face-centric selfie images in a real-world clinical setting involving high-risk pregnant patients. This effort is part of a broader research project led by our group, aimed at developing mobile applications for mental health assessment. In our application scenario, the user takes a selfie with a smartphone front-facing camera. The application sends the image to a server, where an AI model analyzes it and provides a label indicating whether the patient is normal or has symptoms of depression-anxiety. The server returns the label to the application, providing feedback to the user. To train and evaluate the models, we use image data (selfies) and responses to the Patient Health Questionnaire-4

(PHQ-4), a brief screening tool for anxiety and depression. The PHQ-4 responses are used to derive image labels – normal or abnormal (depression-anxiety) –, which play the role of supervisory signals during training and ground truth for evaluation. To the best of our knowledge, this is the first study that investigates the use of face-centric selfies for screening anxiety and depression in high-risk pregnant patients in a clinical context.

In line with the state-of-the-art, deep models are employed in facial analysis. Pre-trained models are leveraged using two distinct approaches to address the challenge of limited training data (most participants contributed with only a single photo). In a more traditional approach, we employ a transfer learning strategy where convolutional neural networks (CNNs) originally trained for facial expression recognition are fine-tuned for depression-anxiety detection. In a more innovative approach, we propose using powerful vision-language models (VLMs) as a facial analyzer. While VLMs are not suitable for directly assessing depression-anxiety, as demonstrated in our experiments, they excel in zero-shot detailed descriptions of facial expressions that correspond to basic emotions (e.g., anger, happiness, and sadness). This approach detects depression-anxiety by classifying the generated text with simple neural networks instead of directly classifying the image. Experiments conducted under a rigid *Leave One Subject Out* protocol revealed that our VLM-based approach outperformed the traditional CNNs, achieving an accuracy of 77.6%, which represents a gain of approx. 10.0 percentage points (p.p.) compared to the CNNs, and an F1-score of 56.0%, an improvement of approx. 11.0 p.p.

In summary, the main contributions of this work are:

- A novel VLM-based approach for ubiquitous mental-health screening from selfies.
- A study on the use of AI for depression-anxiety screening in high-risk pregnant patients.
- Collection of a dataset¹ comprising selfies and PHQ-4 responses from high-risk pregnant patients.
- Comprehensive assessment of VLMs and CNNs for depression-anxiety screening with limited data.

The rest of this paper is organized as follows. The next section discusses the related works. Section 3

¹Anonymized descriptions produced by the VLMs and corresponding PHQ-4 responses are publicly available at <https://bit.ly/3E1EPKw>.

describes the two approaches for AI-driven screening methodology. Section 4 presents the experimental setup, while Section 5 discusses the results. Finally, Section 6 concludes the paper and discusses future work opportunities.

2 RELATED WORK

The use of facial features to identify non-basic emotions (depression, stress, engagement, shame, guilt, envy, among others) is becoming increasingly popular in various applications. Nepal et al. (2024) highlight several elements explored in this type of analysis, such as facial expression itself, gaze direction, as well as general image characteristics like luminosity and background settings. In this domain, machine learning (ML) plays an essential role, particularly with the use of deep learning (DL) models and related techniques, such as transfer learning and attention mechanisms. As highlighted by Kumar et al. (2024), DL models have become the primary choice for detecting non-basic emotions since 2020, surpassing traditional machine learning and image processing algorithms. Regarding models, convolutional neural networks (CNNs) are the most popular choice for image-based analysis of non-basic emotions. For instance, Gupta et al. (2023) used CNNs to quantify student engagement in online learning, while Zhou et al. (2018) applied them for detecting depression based on facial images. Transfer learning plays a crucial role as pre-trained models for tasks like basic facial expression recognition can be fine-tuned for specific targets like detecting stress (Voleti et al., 2024). Usually, pre-training leverages larger datasets, which is particularly useful when the target dataset is small, as is often the case in mental health applications. Attention mechanisms are also relevant in this context, as they can help models focus on specific regions of the face that are more informative for the target task (Viegas et al., 2018; Belharbi et al., 2024).

Despite the promising results, most studies on depression detection through facial expressions involve capturing images in controlled environments where individuals follow a predetermined script, which limits the realism of the facial expressions (Liu et al., 2022; Kong et al., 2022) and the broad applicability of pre-trained models. Nonetheless, a few studies focus on natural images captured by smartphone cameras for ubiquitous screening. Darvari et al. (2020) developed an application where users record their emotional state through photographs taken by the rear camera of smartphones. Alternatively, front-facing cameras can facilitate the capture of facial im-

ages, whose analysis can offer valuable insights into a person's emotional state (Wang et al., 2015). In recent work, Nepal et al. (2024) introduced a depression screening approach named MoodCapture, which was evaluated on a large dataset of facial images collected in the wild (over 125,000 photos). A key aspect of MoodCapture is the passive collection of images in multiple shots while users complete a self-assessment depression questionnaire on their smartphones. The authors argue that these images are preferable to traditional selfies as they capture more authentic and unguarded facial expressions.

Similarly to MoodCapture, the proposed work focuses on analyzing selfies captured by smartphone cameras for anxiety and depression screening. However, our study differs from MoodCapture in two major aspects. First, our work targets a specific population: high-risk pregnant patients accompanied by a team of healthcare professionals. This imposes a limitation on the dataset size when compared to crowd-sourced data. Second, our methodology focuses on face-centric selfies rather than analyzing general image aspects. To the best of our knowledge, this is the first study to explore the use of face-centric selfies for anxiety and depression screening in high-risk pregnant patients. Another relevant aspect of our work is the use of vision-language models (VLMs). Unlike traditional models that predict a limited set of basic emotions, VLMs can capture more nuanced emotions such as awe, shame, and emotional suppression (Bian et al., 2024). Despite being used for general emotion analysis, the particular application of VLMs for mental health screening is another contribution of our work.

3 AI-DRIVEN SCREENING VIA SELFIE ANALYSIS

This study relies on selfies taken by pregnant patients paired with their responses to the Patient Health Questionnaire-4 (PHQ-4), a four-item instrument for brief screening of anxiety and depression. The PHQ-4 was chosen for this study due to its efficiency and brevity in screening for both anxiety and depression, making it particularly useful in clinical settings where time is limited. Comprising only four items (Figure 1), it integrates questions from the PHQ-2 and GAD-2, which are validated tools for detecting depression and anxiety, respectively. This dual focus allows for a rapid assessment of two of the most prevalent mental health conditions, while maintaining strong psychometric properties, making it a reliable and practical tool for screening purposes in a clinical sample. The

PHQ-4				
Over the last 2 weeks, how often have you been bothered by the following problems?				
	Not at all	Several days	More than half the days	Nearly every day
Feeling nervous, anxious or on edge	0	1	2	3
Not being able to stop or control worrying	0	1	2	3
Little interest or pleasure in doing things	0	1	2	3
Feeling down, depressed, or hopeless	0	1	2	3

Figure 1: Patient Health Questionnaire-4 (PHQ-4) items and the respective frequency scores.

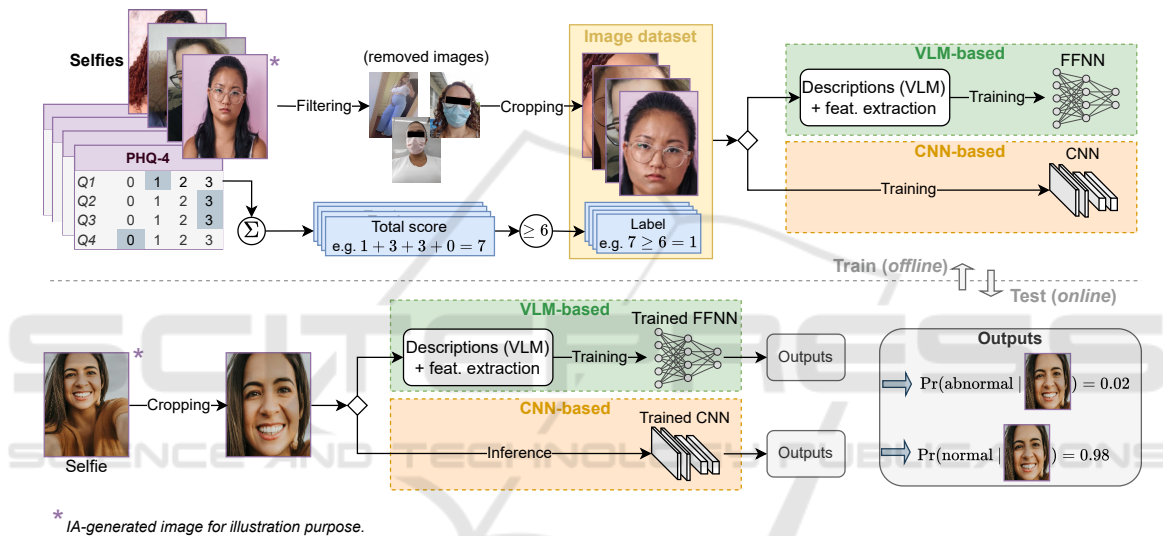


Figure 2: Overview of the proposed methodology for depression-anxiety detection using selfies and PHQ-4 responses. The training pipeline (top flow) involves filtering invalid selfies, cropping face regions with MTCNN, and labeling images based on PHQ-4 scores. The model (CNN or FFNN) is trained either directly from images (CNN-based) or text descriptions (VLM-based). The test pipeline (bottom flow) uses the trained model to classify new selfies.

total score, calculated as the sum of individual scores, ranges from 0 to 12, with higher scores indicating greater severity of anxiety and depression symptoms. A score of 6 or higher on the PHQ-4 is typically used as a cut-off point for identifying cases where either anxiety or depression (or both) may be present and warrant further clinical evaluation (Caro-Fuentes and Sanabria-Mazo, 2024). This cut-off point helps to identify individuals with moderate to severe symptoms of either condition, ensuring efficient screening, including mixed anxiety-depression states, which are fairly common among pregnant women. It allows for initial assessment without the need to differentiate between the two disorders (Javadekar et al., 2023).

To enable AI-driven screening, we propose a

methodology that leverages selfies and PHQ-4 responses to train machine learning (ML) models for depression-anxiety detection. Figure 2 provides a joint overview of the VLM- and CNN-based approaches addressed in this work. There are two pipelines: the training pipeline (top flow) and the test pipeline (bottom flow). The training pipeline leverages image and PHQ-4 data to train an ML model for depression-anxiety detection. The training requires an image dataset comprising faces and the respective labels (0-normal, 1-abnormal). To assemble this dataset, the first step is manually filtering out invalid selfies, which are those taken with the rear-facing camera or with faces covered by a mask. The face region of the remaining samples is subsequently

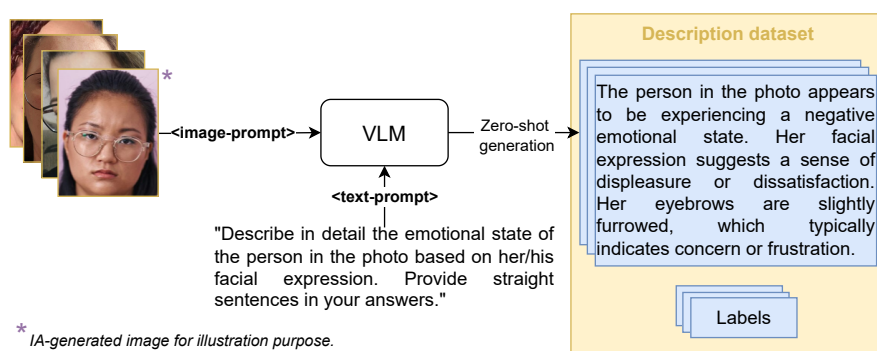


Figure 3: Zero-shot description generation with VLMs. The VLM prompt consists of an image (cropped face from a selfie) and a text instruction (text prompt). A description is generated for each face image in the image dataset. The label from each source image is transferred to the respective generated description, giving rise to an annotated dataset of textual descriptions.

cropped by using a multi-task cascaded convolutional network (MTCNN) (Zhang et al., 2016). Selfies with multiple detected faces are discarded, as the PHQ-4 responses are individual. As previously explained, a label is derived for each face image by thresholding the PHQ-4 overall score.

With a valid set of face images and labels, an ML model can be trained. In the CNN-based approach, the face images are directly input into the classification model. The VLM-based approach, by contrast, involves generating descriptions of emotional states through face analysis and then extracting features from these textual descriptions. These features are subsequently used as inputs for the classification head, specifically a feed-forward neural network (FFNN). Once the model (CNN or FFNN) is trained, it can be used to classify an input selfie, as illustrated in the bottom flow of Figure 2. The face region in the selfie is cropped using the MTCNN, as in the top flow. The face image (or text features) is then input to the trained model, which outputs class probabilities. The depression-anxiety condition is verified if, and only if, $\Pr(\text{abnormal}|\text{sample}) > 0.5$. More details of the two detection approaches are provided in the following sections.

3.1 CNN-Based Approach

This approach employs CNNs for predicting depression-anxiety directly from image data, as illustrated in Figure 2. Training from scratch is unfeasible in our context due to the reduced number of selfies/PHQ-4 responses: 147 samples from a total of 108 participants. To circumvent this issue, we start with CNN models pre-trained on the ImageNet dataset and fine-tune them on large-scale facial expression recognition (FER) datasets, which are closely related to our target task. A second

fine-tuning is performed to adapt the FER pre-trained models to our task.

This work investigates the adaptation of models pre-trained on the FER2013 (Goodfellow et al., 2013), and RAF-DB (Li et al., 2017) datasets, both of which encompass seven basic emotions (classes): anger, disgust, fear, happiness, sadness, surprise, and neutral state. Four CNN architectures were investigated: EfficientNetV2 (Tan and Le, 2021), ResNet-18 and ResNet-50 (He et al., 2016), and VGG11 (Simonyan and Zisserman, 2014). In the second fine-tuning stage, the classifier consists of a CNN backbone pre-trained on FER or RAF-DB appended with a 2-output fully connected layer (classification head). The backbone is frozen during training to prevent overfitting, ensuring that only the classification head is trained.

3.2 VLM-Based Approach

This approach leverages large generative vision-language models (VLMs) for facial analysis. Modern VLMs (Bordes et al., 2024) combine visual encoders with large language models (LLMs) and can simultaneously learn from images and text. This enables them to perform various tasks, such as answering visual questions and captioning images. In particular, we benefit from the VLM zero-shot instruction-following ability to produce high-quality descriptions when provided with an image and text prompt. Three modern VLMs were investigated in this work: LLaVA-NeXT (Liu et al., 2024a), which is an improvement on LLaVA (Liu et al., 2024b), Kosmos-2 (Peng et al., 2023), and the proprietary GPT-4o.

Figure 3 illustrates the intended usage of VLMs in this work. Instead of using the image-labeled dataset directly, textual descriptions (in the form of sentences) are generated by following the instruction




Figure 4: Data distribution after removing invalid samples. Most subjects contributed with a single sample, while one contributed with nine. The imbalance is evident, with a bias towards negative samples (PHQ-4 < 6).

in the input prompt: “Describe in detail the emotional state of the person in the photo based on her/his facial expression. Provide straight sentences in your answer.”. It is worth mentioning that the prompt does not address the target task directly, i.e., detecting anxiety and/or depression. Nonetheless, experimental results (Section 5) reveal that the pre-trained VLMs were unable to directly predict the depression-anxiety condition. This motivated using VLMs as analysts rather than judges, delegating the final decision to a secondary model, specifically an FFNN.

To build the classification model, features are extracted from the text descriptions by using a pre-trained Sentence-BERT (Reimers and Gurevych, 2019) model called `all-MiniLM-L6-v2`, which is designed to embed sentences and paragraphs into a 384-dimensional representation. Two FFNN architectures (default and alternative) are investigated in this work, as expressed in the Equation (1):

$$f_{net}(\mathbf{x}) = \begin{cases} \mathbf{W}\mathbf{x} + \mathbf{b} & \text{(default)} \\ \mathbf{W}_2(\text{ReLU}(\mathbf{W}_1\mathbf{x} + \mathbf{b}_1) + \mathbf{b}_2 & \text{(alternative)} \end{cases} \quad (1)$$

where \mathbf{W} . and \mathbf{b} . denote a weight matrix and a bias vector, respectively. The default model is a simple linear classifier, while the alternative model includes a hidden layer and is explored in a sensitivity analysis in the experiments. The logits of the network, $f_{net}(\mathbf{x})$, are used to calculate the class probabilities through softmax normalization: $\Pr(\cdot|\mathbf{x}) = \text{softmax}(f_{net}(\mathbf{x}))$.

4 EXPERIMENTAL METHODOLOGY

This section outlines the elements of the experimental methodology: the data collection procedure involving selfies and PHQ-4 responses, the experiments conducted, and, finally, the hardware and software setup.

4.1 Data Collection

Selfies and PHQ-4 responses constitute a subset of data collected for a research project (local ethics

board consent: CAAE 64158717.9.0000.0065) involving high-risk pregnant patients from the Clinical Hospital of the University of São Paulo (São Paulo, Brazil). The research focused on pregnant women aged 18 and above who could provide informed consent, had at least an elementary education, and owned a smartphone. Data collection utilized smartphones, with the participants capturing natural photos (selfies) in such a way that their faces were visible. They were responsible for uploading image data to the REDCap platform (Harris et al., 2009) instance hosted at the University of São Paulo, which was also utilized for completing the PHQ-4 questionnaire.

For this study, 108 participants were recruited and instructed to submit bi-weekly reports, each consisting of the PHQ-4 questionnaire² and a selfie. Figure 4 illustrates the data distribution after manually removing invalid selfies. The left chart displays a total of 147 selfies from the 108 participants. Although each participant was expected to contribute multiple samples (11 in total), most participants submitted only a single selfie. The right-side chart highlights the imbalance in the data, with a bias towards negative samples (PHQ-4 < 6): 106 negative versus 41 positive samples.

4.2 Comparative Evaluation

The comparative evaluation aims to assess the performance of the CNN- and VLM-based approaches for detecting depression-anxiety in selfies. Recall that while a CNN architecture is trained in the first approach, a FFNN architecture is trained in the latter approach. The performance of pre-trained VLMs for zero-shot screening was also evaluated to serve as a baseline for the VLM-based approach. To conduct the experiments, the collected data was processed to create an image dataset consisting of cropped faces and labels, as discussed in Section 3. The CNN and FFNN models were trained and tested under a *Leave One Subject Out* (LOSO) cross-validation protocol to avoid performance overestimation. This protocol ensures that each subject’s data is used as a unique test set while the remaining data forms the training set, providing a robust measure of model performance across different individuals.

In each training-testing session, the samples of a single subject are classified, and the predictions (normal or abnormal) are recorded. Once all predictions are available, performance metrics are calculated. Traditional metrics in binary classification are utilized: precision, recall, F1-score, area under

²The Portuguese version of the PHQ-4 used in this work is provided by Pfizer (Copyright©, 2005 Pfizer Inc).

ROC curve (AUC), and accuracy. F1-score is particularly relevant because it accounts for the dataset imbalance, a characteristic of our data. Specific implementation details for each depression-anxiety detection approach are provided below.

CNN-Based Approach. The pre-training on both FER2013 and RAF-DB was conducted for 30 epochs with a fixed learning rate of 10^{-4} and a ℓ_2 -regularization factor of 10^{-5} . In the default settings, fine-tuning was conducted for 30 epochs, with a learning rate of 10^{-4} , and an ℓ_2 -regularization factor of 10^{-5} . Different parameters were used based on empirical evidence of overfitting in the training. ResNet-18 (RAF-DB): for 50 epochs, a learning rate of 2×10^{-5} , and a weight decay of 10^{-6} . ResNet-50 (RAF-DB): it included a dropout regularization with a probability of 50%. Pre-training and fine-tuning utilized Adam optimization.

VLM-Based Approach. VLMs were configured for deterministic inference, employing a strategy that selects the most probable next token at each step of the generation process. The default (linear) model (Equation 1) was adopted as the classifier. To address the dataset imbalance, the descriptions associated with positive samples (minority class) were up-sampled to match the number of negative samples. For each VLM, the classification head underwent training using the Adam optimizer for 15 epochs. Additional training parameters included a learning rate of 10^{-4} , a batch size of 2, and a ℓ_2 -regularization factor of 10^{-4} . During each training session, a randomly selected subset comprising 10% of the training data was reserved for validation purposes. The best-epoch checkpoint was determined based on achieving the highest F1-score on the validation set.

Zero-Shot Screening with VLMs. A natural question arises when using powerful models such as VLMs: “Are pre-trained VLMs capable of zero-shot screening depression-anxiety?” To answer this question, we conducted an experiment in which VLMs were asked to classify an input image based on the following instruction prompt: “Describe in detail the emotional state of the person in the photo based on her/his facial expression. Provide straight sentences in your answers. Based on your description, classify the emotional state as either ‘normal’, ‘anxiety’, or ‘depression’. The output must be exactly one of these words. Follow the template: Output: {result}”. In preliminary tests, only GPT-4o followed strictly the instruction prompt, while LLaVA-NeXT and Kosmos-2 generated longer and more complex

descriptions. Therefore, this experiment was performed only for GPT-4o. GPT-4o classified each of the 147 samples into one of the three classes: normal, anxiety, or depression. The samples classified as anxiety or depression were considered positive, while the normal samples were considered negative.

4.3 Sensitivity Analysis

This additional experiment explores the effects of incorporating a hidden layer into the FFNN of the VLM-based approach (alternative model, Equation 1). The model was evaluated with various hidden units: $h = 4, 8, 16, \dots, 256$. This investigation was motivated by the promising results observed in the VLM-based approach, as elaborated further in Section 5. Furthermore, understanding the impact of the classification head is relevant because it is the only trainable component in the VLM-based pipeline. The LOSO protocol was also employed in this investigation. Additionally, during model training, a dropout layer was incorporated after the ReLU activation function to prevent overfitting.

4.4 Hardware-Software Setup

Hardware: Intel(R) Xeon(R) CPU @ 2.20GHz with 32GB of RAM, running Linux Ubuntu 22.04.4 LTS, and equipped with an NVIDIA L4 GPU with 24GB of memory. Software: The source code was written in Python 3.10, mostly using PyTorch 2.3 for model training and inference. The Sentence-BERT (Reimers and Gurevych, 2019) is implemented in SentenceTransformers library³. The LLaVA-NeXT and Kosmos-2 models – `llava-hf/llava-v1.6-mistral-7b-hf` and `microsoft/kosmos2-patch14-224`, respectively, are available at the HuggingFace⁴, while GPT-4o was accessed via the OpenAI API.

5 RESULTS AND DISCUSSION

This section presents the results of the conducted experiments: comparative evaluation and sensitivity analysis. Limitations and challenges are also discussed at the end of this section.

³<https://sbnet.net>

⁴<https://huggingface.co>

Table 1: Overall results for the comparative evaluation (%).

Model	Pre-train	Prec.	Rec.	F1-score	AUC	Acc.
CNN-based						
ResNet-18	FER2013	36.2	51.2	42.4	62.7	61.2
ResNet-50		36.5	46.3	40.9	60.1	62.6
VGG11		39.0	39.0	39.0	57.1	66.0
EfficientNetV2		39.6	51.2	44.7	65.7	64.6
ResNet-18	RAF-DB	37.9	53.7	44.4	64.5	62.6
ResNet-50		28.0	34.1	30.8	52.1	57.1
VGG11		28.6	43.9	34.6	54.8	53.7
EfficientNetV2		41.7	48.8	45.0	68.4	66.7
VLM-based						
GPT-4o*	-	61.8	51.2	56.0	72.6	77.6
Kosmos-2*		40.0	53.7	45.8	72.0	64.6
LLaVA-NEXT*		41.2	51.2	45.7	66.2	66.0
Zero-shot Screening						
GPT-4o	-	53.9	34.2	41.8	61.4	73.5

Bold values indicate the highest metric value within each approach.

*VLM with (default) linear classifier model.

5.1 Comparative Evaluation

Table 1 shows the results of the comparative evaluation experiment. Overall, EfficientNetV2 outperformed the compared models among the CNNs for both FER2013 and RAF-DB pre-training, while GPT-4o yielded the best performance among the VLMs. Pre-training on RAF-DB improved most metrics for EfficientNetV2 and ResNet-18, although a notable decrease in performance was observed for ResNet-50 and VGG11. Notably, ResNet-18 yielded better performance than ResNet-50, despite its reduced size. This suggests that larger models might require more data to achieve better performance.

Regarding the F1-score, using RAF-DB resulted in only a 0.3 p.p. improvement over FER2013 for EfficientNetV2, while recall decreased by 2.4 p.p.. In medical applications, lower recall indicates the risk of missing a significant number of actual cases, which can lead to undiagnosed conditions and potentially severe consequences for patient health and safety. In this context, a remarkable gain was observed for ResNet-18, whose recall raised from 46.3 to 53.7% with RAF-DB. Despite its lower precision compared to EfficientNetV2, ResNet-18 presents itself as a viable alternative given the critical importance of the recall metric in this context. Moreover, its F1-score is only marginally lower by less than 1 p.p. compared to EfficientNetV2. The use of CNNs represents a more traditional way to address this problem. The obtained results for these models reveal how challenging the task is for the addressed scenario, specifically when using a small dataset.

As an alternative, we proposed using pre-trained VLMs due to their ability to analyze faces. Table 1 shows the results for the VLM-based approach with the classification head (FFNN) in its default configu-





Figure 5: Sensitivity analysis. The FFNN classifier (VLM-based) was evaluated with various hidden units: $h = 4, 8, 16, \dots, 256$. The dashed line in each chart represents the highest value for the respective metric as reported in Table 1.

ration. Overall, the F1-score for the three evaluated VLMs surpassed the CNN-based models, with GPT-4o achieving the highest value. GPT-4o outperformed the open-source models, Kosmos-2 and LLaVA-NeXT, across all metrics, except for the recall metric, where Kosmos-2 achieved the same 53.7% as ResNet-18. LLaVA-NeXT and Kosmos-2 showed a similar F1-score, with the most significant difference observed in the AUC metric: nearly 6 p.p. in favor of Kosmos-2. A particularly notable result is the F1-score obtained with GPT-4o, which surpasses its competitors by a large margin, almost 10 p.p. higher, being the only model to perform above 50% in this metric. In summary, GPT-4o yielded the best performance considering both CNN- and VLM-based approaches.

The last row of Table 1 shows the results for the zero-shot screening with GPT-4o. While the accuracy in this scenario was superior to that obtained with LLaVA-NeXT and Kosmos-2, the recall and F1-score demonstrated the opposite trend. Remarkably, the recall in zero-shot screening was 17 p.p. lower than that obtained with the VLM-based classification models. By comparing with the GPT-4o in the VLM-based approach, we conclude that the proposed formulation based on description generation and classification is crucial for achieving high performance.

5.2 Sensitivity Analysis

While Table 1 focused on the default FFNN model (VLM-based), this experiment investigated the alternative (single hidden layer) FFNN model (Equation (1)). More specifically, it was analyzed the impact of increasing the model complexity by varying the number of hidden units ($h = 4, 8, 16, \dots, 256$). Figure 5 shows the results of this experiment. The dashed line in each chart represents the highest value for the respective metric as reported in Table 1. Notably, the F1-score achieved with GPT-4o and the default classification model (56.0%) – indicated by the dashed



(a) PHQ-4=8 (b) PHQ-4=2 (c) PHQ-4=7 (d) PHQ-4=2

Figure 6: Challenging cases that resulted in misclassification. The Grad-CAM attention maps highlight the areas around the mouth as the most influential for EfficientNetV2’s prediction in (a) and (b).

line in the third chart (from left to right) – demarcates an empirical upper bound for this metric.

The F1-score curves follow a similar pattern to the recall curves, both exhibiting a ‘V’ shape from $h = 4$ to 64. Kosmos-2 outperformed the other VLMs in recall in most cases, achieving a maximum recall of 85.3% for $h = 4$ (with an F1-score of 52.2%). The high recall indicates a solid ability to flag more potential cases, which is highly desirable for this type of application. Kosmos-2 (with $h = 256$) and LLaVA-NExT (with $h = 32$) were able to match the F1-score of GPT-4o with the default model, however, with significantly higher recall: 65.3 and 68% for Kosmos-2 and LLaVA-NExT, respectively, compared to 51.2% achieved by GPT-4o. This shows that open-source VLMs can yield competitive performance when combined with a more complex classification model.

5.3 Limitations and Challenges

Learning the relationship between facial expressions and self-reported PHQ-4 scores from single selfies can be challenging with small datasets. When using CNNs, prediction is mainly influenced by features around the mouth, as evidenced by the Grad-CAM (Selvaraju et al., 2017) attention maps for EfficientNetV2 in Figures 6a and 6b. In Figure 6a, the slight smile led the model to classify a positive sample ($\text{PHQ-4} \geq 6$) as negative. However, facial cues around the eyes might suggest a posed (non-Duchenne) smile, which is not necessarily a sign of happiness or well-being. This issue could be mitigated by incorporating facial action units (AUs) into the learning process, enhancing attention guidance to other critical regions, as recently proposed in a study on basic emotions (Belharbi et al., 2024).

Approximately 12% of the errors with GPT-4o involve positive samples misclassified as negative where the individual is smiling. Figure 6c shows an example of a genuine (Duchenne) smile associated with a PHQ-4 score of 7 (close to the threshold) that was misclassified as negative. The description generated by GPT-4o includes terms such as “The per-

son in the photo appears to be happy”, “She is smiling broadly”, and “Her eyes are slightly squinted”, which accurately reflect the visible facial expressions. This situation is potentially misleading even for experienced face analysts, as the individual’s mood is positive, but the PHQ-4 score is close to the threshold.

Nearly 18% of the errors with GPT-4o are related to negative samples with PHQ-4 score lower than 3 described by the model as “neutral”. The sample in Figure 6d (PHQ-4=2) was described by GPT-4o as “neutral or somewhat weary expression”, “corners of their mouth are slightly downturned”, “lack of enthusiasm”. This suggests that the negative elements critically influenced the positive response despite the neutral emotional state described by the VLM. From a data perspective, multiple captures or fusion with complementary data modalities could yield a significant improvement in the overall performance for both smiling and neutral expressions scenarios.

AI systems are notorious for showing errors and biases in different racial groups (Nazer et al., 2023). Despite the significance of this issue and its implications for fairness and equity, this sensitive topic was not explored in the present study.

6 CONCLUSION

This work addressed the AI-driven mental health screening in mobile applications using face-centric selfies. The scope of the study included collecting a dataset of selfies paired with responses to a self-reported questionnaire (PHQ-4) and evaluating two depression-anxiety detection approaches. In the comparative evaluation, the proposed VLM-based approach yielded better results than the typical transfer learning with CNNs or zero-shot screening with GPT-4o. Notably, GPT-4o achieved the best F1-score (56%) and accuracy (77.6%), while Kosmos-2 attained the highest recall (53.7%).

The sensitivity analysis demonstrated that open-source VLMs can yield competitive performance, nearing the F1-score of GPT-4o. When combined with more complex FFNNs, Kosmos-2 and LLaVA-NExT achieved 56% of F1-scores but with significantly higher recall, which is particularly important in this context. Specifically, Kosmos-2’s recall increased to 85.3% (with an F1-score of 52.2%) by adding a 4-unit hidden layer.

Future work will explore two main directions. From a data perspective, we plan to collect a larger dataset with more reliable data by developing a smartphone application for multiple-shot passive capture of face images. From a methodological perspective,

smile analysis and action unit information will be investigated to address the limitations of the current approach. Furthermore, fusion with complementary data modalities, such as audio and text transcriptions, will be investigated to enhance the screening performance.

ACKNOWLEDGEMENTS

We acknowledge the support of FAPES and CAPES (Process 2021-2S6CD, FAPES No. 132/2021) through the PDPG (Programa de Desenvolvimento da Pós-Graduação – Parcerias Estratégicas nos Estados). We sincerely thank Vera Lucia Carvalho Tess, Kelly dos Santos Prado and the team at the Clinical Outpatient Department for Pregnancy and Postpartum Care at the Institute of Psychiatry at HCFMUSP for their invaluable assistance in the recruitment and evaluation of participants. Our gratitude also extends to Ana Maria S. S. Oliveira and Luciana Cristina Ranhel for their support in participant recruitment, and to Lucas Batini Araujo and Herbert Dias de Souza for their help in facilitating communication with participants and acquiring the necessary data.

REFERENCES

- Belharbi, S., Pedersoli, M., Koerich, A. L., Bacon, S., and Granger, E. (2024). Guided interpretable facial expression recognition via spatial action unit cues. In *2024 IEEE 18th International Conference on Automatic Face and Gesture Recognition (FG)*, pages 1–10.
- Biaggi, A., Conroy, S., Pawlby, S., and Pariante, C. M. (2016). Identifying the women at risk of antenatal anxiety and depression: A systematic review. *Journal of affective disorders*, 191:62–77.
- Bian, Y., Küster, D., Liu, H., and Krumhuber, E. G. (2024). Understanding naturalistic facial expressions with deep learning and multimodal large language models. *Sensors*, 24(1):126.
- Bordes, F., Pang, R. Y., Ajay, A., Li, A. C., Bardes, A., Petryk, S., Mañas, O., Lin, Z., Mahmoud, A., Jayaraman, B., et al. (2024). An introduction to vision-language modeling. *arXiv preprint arXiv:2405.17247*.
- Caro-Fuentes, S. and Sanabria-Mazo, J. P. (2024). A systematic review of the psychometric properties of the patient health questionnaire-4 (phq-4) in clinical and non-clinical populations. *J. of the Academy of Consultation-Liaison Psychiatry*.
- Darvariu, V.-A., Convertino, L., Mehrotra, A., and Mulesi, M. (2020). Quantifying the relationships between everyday objects and emotional states through deep learning based image analysis using smartphones. *Proc. of the ACM on Interactive, Mobile, Wearable and Ubiquitous Tech.*, 4(1):1–21.
- Dowse, E., Chan, S., Ebert, L., Wynne, O., Thomas, S., Jones, D., Fealy, S., Evans, T.-J., and Oldmeadow, C. (2020). Impact of perinatal depression and anxiety on birth outcomes: a retrospective data analysis. *Maternal and child health journal*, 24:718–726.
- Goodfellow, I. J., Erhan, D., Carrier, P. L., Courville, A., Mirza, M., Hamner, B., Cukierski, W., Tang, Y., Thaler, D., Lee, D.-H., et al. (2013). Challenges in representation learning: A report on three machine learning contests. In *Neural Inf. Process.: 20th Int. Conf., Daegu, Korea, November 3-7, 2013. Proc., Part III 20*, pages 117–124. Springer.
- Gupta, S., Kumar, P., and Tekchandani, R. K. (2023). Facial emotion recognition based real-time learner engagement detection system in online learning context using deep learning models. *Multimedia Tools and Applicat.*, 82(8):11365–11394.
- Harris, P. A., Taylor, R., Thielke, R., Payne, J., Gonzalez, N., and Conde, J. G. (2009). Research electronic data capture (redcap)—a metadata-driven methodology and workflow process for providing translational research informatics support. *J. of Biomed. Informat.*, 42(2):377–381.
- He, K., Zhang, X., Ren, S., and Sun, J. (2016). Deep residual learning for image recognition. In *Proc. of the IEEE Conf. on Comput. Cision and Pattern Recognition*, pages 770–778.
- Javadekar, A., Karmarkar, A., Chaudhury, S., Saldanha, D., and Patil, J. (2023). Biopsychosocial correlates of emotional problems in women during pregnancy and postpartum period. *Industrial Psychiatry Journal*, 32(Suppl 1):S141–S146.
- Kong, X., Yao, Y., Wang, C., Wang, Y., Teng, J., and Qi, X. (2022). Automatic identification of depression using facial images with deep convolutional neural network. *Medical Science Monitor: Int. Med. J. of Experimental and Clinical Research*, 28:e936409–1.
- Kumar, P., Vedernikov, A., and Li, X. (2024). Measuring non-typical emotions for mental health: A survey of computational approaches. *arXiv preprint arXiv:2403.08824*.
- Kurki, T., Hiihesmaa, V., Raitasalo, R., Mattila, H., and Ylikorkala, O. (2000). Depression and anxiety in early pregnancy and risk for preeclampsia. *Obstetrics & Gynecology*, 95(4):487–490.
- Li, H., Bowen, A., Bowen, R., Muhajarine, N., and Balbuena, L. (2021). Mood instability, depression, and anxiety in pregnancy and adverse neonatal outcomes. *BMC Pregnancy and Childbirth*, 21:1–9.
- Li, S., Deng, W., and Du, J. (2017). Reliable crowdsourcing and deep locality-preserving learning for expression recognition in the wild. In *Proc. of the IEEE Conf. on Comput. Vision and Pattern Recognition*, pages 2852–2861.
- Liu, D., Liu, B., Lin, T., Liu, G., Yang, G., Qi, D., Qiu, Y., Lu, Y., Yuan, Q., Shuai, S. C., et al. (2022). Measuring depression severity based on facial expression and body movement using deep convolutional neural network. *Frontiers in psychiatry*, 13:1017064.

- Liu, H., Li, C., Li, Y., Li, B., Zhang, Y., Shen, S., and Lee, Y. J. (2024a). Llava-next: Improved reasoning, ocr, and world knowledge.
- Liu, H., Li, C., Wu, Q., and Lee, Y. J. (2024b). Visual instruction tuning. *Advances in Neural Inf. Proc. Syst.*, 36.
- Liu, Y., Wang, K., Wei, L., Chen, J., Zhan, Y., Tao, D., and Chen, Z. (2024c). Affective computing for healthcare: Recent trends, applications, challenges, and beyond. *arXiv preprint arXiv:2402.13589*.
- Nasreen, H. E., Pasi, H. B., Rifin, S. M., Aris, M. A. M., Rahman, J. A., Rus, R. M., and Edhborg, M. (2019). Impact of maternal antepartum depressive and anxiety symptoms on birth outcomes and mode of delivery: a prospective cohort study in east and west coasts of malaysia. *BMC pregnancy and childbirth*, 19:1–11.
- Nazer, L. H., Zatarah, R., Waldrip, S., Ke, J. X. C., Moukheiber, M., Khanna, A. K., Hicklen, R. S., Moukheiber, L., Moukheiber, D., Ma, H., et al. (2023). Bias in artificial intelligence algorithms and recommendations for mitigation. *PLOS Digital Health*, 2(6):e0000278.
- Nepal, S., Pillai, A., Wang, W., Griffin, T., Collins, A. C., Heinz, M., Lekkas, D., Mirjafari, S., Nemesure, M., Price, G., et al. (2024). Moodcapture: Depression detection using in-the-wild smartphone images. In *Proc. of the CHI Conf. on Human Factors in Comput. Syst.*, pages 1–18.
- Peng, Z., Wang, W., Dong, L., Hao, Y., Huang, S., Ma, S., and Wei, F. (2023). Kosmos-2: Grounding multimodal large language models to the world. *arXiv preprint arXiv:2306.14824*.
- Reimers, N. and Gurevych, I. (2019). Sentence-bert: Sentence embeddings using siamese bert-networks. *arXiv preprint arXiv:1908.10084*.
- Selvaraju, R. R., Cogswell, M., Das, A., Vedantam, R., Parikh, D., and Batra, D. (2017). Grad-cam: Visual explanations from deep networks via gradient-based localization. In *Proc. of the IEEE Int. Conf. on Comput. Vision*, pages 618–626.
- Simonyan, K. and Zisserman, A. (2014). Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*.
- Tan, M. and Le, Q. (2021). EfficientNetV2: Smaller Models and Faster Training. In Meila, M. and Zhang, T., editors, *Proc. of the 38th Int. Conf. on Machine Learning*, volume 139 of *Proc. of Mach. Learning Research*, pages 10096–10106. PMLR.
- Thombs, B. D., Markham, S., Rice, D. B., and Ziegelstein, R. C. (2023). Screening for depression and anxiety in general practice. *BMJ*, 382.
- Viegas, C., Lau, S.-H., Maxion, R., and Hauptmann, A. (2018). Towards independent stress detection: A dependent model using facial action units. In *2018 Int. Conf. on Content-based Multimedia Indexing*, pages 1–6. IEEE.
- Voleti, S., NagaRaju, M. S., Kumar, P. V., and Prasanna, V. (2024). Stress detection from facial expressions using transfer learning techniques. In *2024 Int. Conf. on Distrib. Comput. and Optimization Tech.*, pages 1–6. IEEE.
- Wang, R., Campbell, A. T., and Zhou, X. (2015). Using opportunistic face logging from smartphone to infer mental health: challenges and future directions. In *Adjunct Proc. of the 2015 ACM Int. Joint Conf. on Pervasive and Ubiquitous Comput. and Proc. of the 2015 ACM Int. Symp. on Wearable Comput.*, pages 683–692.
- World Health Organization (2023). Depressive disorder (depression). Accessed: 2024-04-07.
- Zhang, K., Zhang, Z., Li, Z., and Qiao, Y. (2016). Joint face detection and alignment using multitask cascaded convolutional networks. *IEEE Signal Processing Letters*, 23(10):1499–1503.
- Zhou, X., Jin, K., Shang, Y., and Guo, G. (2018). Visually interpretable representation learning for depression recognition from facial images. *IEEE Trans. on Affective Comput.*, 11(3):542–552.