

Contextual Emotion Recognition using Large Vision Language Models

Yasaman Etesam and Özge Nilay Yalçın and Chuxuan Zhang and Angelica Lim
Simon Fraser University, BC, Canada

yetesam@sfu.ca, oyalcin@sfu.ca, cza152@sfu.ca, angelica@sfu.ca

Abstract—“How does the person in the bounding box feel?” Achieving human-level recognition of the apparent emotion of a person in real world situations remains an unsolved task in computer vision. Facial expressions are not enough: body pose, contextual knowledge, and commonsense reasoning all contribute to how humans perform this emotional theory of mind task. In this paper, we examine two major approaches enabled by recent large vision language models: 1) image captioning followed by a language-only LLM, and 2) vision language models, under zero-shot and fine-tuned setups. We evaluate the methods on the Emotions in Context (EMOTIC) dataset and demonstrate that a vision language model, fine-tuned even on a small dataset, can significantly outperform traditional baselines. The results of this work aim to help robots and agents perform emotionally sensitive decision-making and interaction in the future.

I. INTRODUCTION

Our ability to recognize emotions allows us to understand one another, build successful long-term social relationships, and interact in socially appropriate ways. Equipping virtual agents and robots with emotion recognition capabilities can help us improve and facilitate human-machine interactions [1]. However, emotion recognition systems today still suffer from poor performance [2] due to the complexity of the task. This innate and seemingly effortless capability requires understanding of the causal relations, contextual information, social relationships as well as theory of mind, which are unresolved problems in affective computing research. Many image-based emotion recognition systems focus solely on using facial or body features [3], [4], which can lead to a low accuracy in the absence of contextual information [5], [6].

In the past few years, the affective computing research community has been moving towards creating datasets and building models that include or make use of contextual information. The EMOTIC dataset, for instance, incorporates contextual and environmental factors for apparent emotion recognition in still images [7]. The inclusion of contextual information beyond facial features is found to significantly improve the accuracy of the emotion recognition models [8], [9]. However, using this information to infer the emotions of others requires commonsense knowledge and high-level cognitive capabilities such as reasoning and theory of mind which are missing from traditional emotion recognition models [10].

Another limitation of traditional emotion recognition models is that many of them are trained and tested on the same dataset [11]. This stands in contrast to the challenge of generalization, where robots may perform poorly in novel

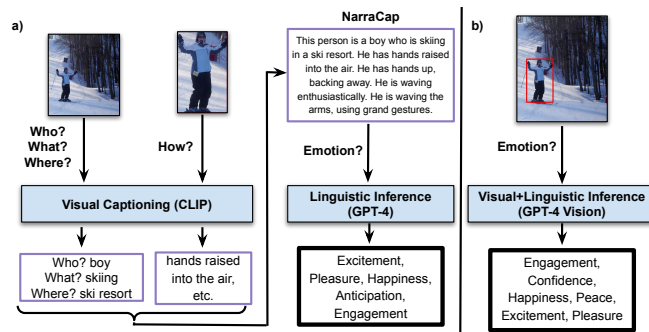


Fig. 1. In this paper, we evaluated two distinct zero-shot approaches: a) using CLIP to generate captions from images, followed by providing these captions alongside a prompt to large language models (LLMs) to obtain emotion labels. b) directly providing the image and a prompt to vision language models (VLMs) and requesting the emotion labels.

situations [12]. In this study, we employ zero-shot models and observe their performance in unseen scenarios. Additionally, we demonstrate how results can be enhanced through fine-tuning. Both LLMs and LVMs are evaluated for this purpose.

Large language models (LLMs) that are based on the transformer architecture [13] have been shown to excel at natural language processing (NLP) tasks [14], [15], offering a way to achieve emotional theory of mind through linguistic descriptors. LLMs gained success in increasing accuracy and efficiency in NLP problems including multimodal tasks such as visual question answering [16] and caption generation [17]. Recently, they have been also used in commonsense reasoning [18], [19], [20], emotional inference [21] and theory of mind [22] tasks, however their capabilities on emotional theory of mind in visual emotion recognition tasks have not been explored.

Vision language models (VLMs) integrate natural language processing with visual comprehension to generate text from visual inputs and are capable of performing a variety of visual recognition tasks. VLMs learn intricate vision-language correlations from large-scale image-text pair datasets, enabling zero-shot predictions across a range of visual recognition tasks [23]. Despite their success in tasks like image classification [24] and object detection [25], their capability in contextual emotion recognition has not yet been explored.

In this paper, we focus on a multi-label, contextual emotional theory of mind task by utilizing the embedded knowledge in large language models (LLMs) and vision language models (VLMs). To the best of our knowledge, this is the first evaluation of VLMs in the contextual emotion recognition task.

The contributions of this paper are as follows:

- Presenting a new state-of-the-art open source method for affect estimation in the wild, using VLM fine-tuning
- Proposing zero-shot approaches for contextual emotion recognition to explore generalizability for robotics
- Evaluating the effectiveness of a) captioning + LLM, versus b) VLM approaches for emotion recognition

II. RELATED WORK

Emotional theory of mind in context. Work in emotional theory of mind (in this paper, also referred to as emotion recognition) has been focusing on the inclusion of contextual information in addition to the facial or posture information in recent years. Early datasets such as HAPPEI [26] proposed, for instance, emotion annotation for groups of people. More recently, the EMOTIC dataset [7] was developed as multilabel dataset containing 26 categories of emotions, 18,316 images and 23,788 annotated people. The related emotion recognition task is to provide a list of emotion labels that matches those chosen by annotators, responding to the question of, "How does the person in the bounding box feel?". In approximately 25% of the person targets, the face was not visible, underscoring the role of context in estimating the emotion of a person in the image. The phrase "emotional theory of mind" is used here to clarify that we are not estimating the sentiment or emotional content of an image, but estimating the emotion of a particular person contained in the image. Note that we do not claim to perform felt emotion recognition, but apparent emotion recognition as perceived by labelers.

Vision-based approaches for contextual emotion estimation. A number of computer vision approaches have been developed in response to the release of the EMOTIC dataset. The EMOTIC baseline [7] uses a CNN to extract features from the target human, as well as the entire image. Subsequent fusion methods incorporated body and context visual information [27] at global or local scales [8], investigated contextual videos [28], or worked to improve subsets of the EMOTIC dataset, such as [29] which considered the photos only including two people. In PERI [30], attention is modulated at various levels within their feature extraction network. In [31], relational region-level analysis was employed, while [32] utilized visual relationship between the main target and adjacent objects. To the extent of our knowledge, the current best approach was Emoticon proposed by [9], which explored the use of graph convolutional networks and image depth map. This approach can further improved by adding CCIM [33] plug-in. Overall, the latest results leave room for improvement, and no pretrained models are available for roboticists, with the exception of EMOTIC.

Large language models and theory of mind. Recent investigations into large language models (LLMs) have uncovered some latent capabilities for social intelligence, including some sub-tasks on emotion inference [21]. Emotional theory of mind tasks using language tend to focus on appraisal based reasoning on emotions, inferring based on a sequence of events. For instance, among other social intelligence tasks,

[22] explored how a language model could respond to an emotional event, e.g. "Although Taylor was older and stronger, they lost to Alex in the wrestling match. How does Alex feel?" Their findings suggested some level of social and emotional intelligence in LLMs.

Natural language and emotional theory of mind. Language is a fundamental element in emotion and it plays a crucial role in emotion perception [34], [35], [36]. In this work, in order to apply LLMs, we use the body of work in English literature that discusses how writers use empathy to describe characters, their actions and external cues, and "seek to evoke feelings in readers employing the powers of narrativity." [37]. Writing textbooks such as The Emotion Thesaurus [38] provide sets of visual cues or actions that support emotional evocation, as a guide for readers to imagine the most relevant features of the person in the scene. For example, to evoke curiosity, a writer may narrate, "she tilted the head to the side, leaning forward, eyebrows furrowing" or "she raised her eyebrows and her body posture perked up."

Generalization and zero-shot learning. A major challenge in the field of robotics is generalization, enabling robots to adapt to new and unforeseen scenarios [12]. Effective training on a labeled dataset often results in satisfactory performance only when evaluations are conducted on similar datasets. This is because the emphasis on minimizing training error tends to make machines capture all the correlations present in the training data, rather than understanding the actual causation [11]. In the real world, the distribution of objects across different categories can exhibit a long-tailed pattern, where some categories are represented by a large number of training images, while others have few or no images at all [39], [40], [41]. This disparity makes it challenging to identify all possible correlations. A crucial goal in visual emotion recognition is to improve the ability to recognize instances of previously unseen visual classes, an approach known as zero-shot learning [42], [43]. In our work, we leverage the embedded knowledge within large language and vision language models to address the visual emotional theory of mind task, aiming to perform this task without training on a specific dataset.

III. METHODOLOGY

In this study, we compare two general approaches: a) 2-phased image captioning and large language models, with b) end-to-end vision language models.

A. Image Captioning and Large Language Model

In this first method, we use a two-phased approach to first generate a caption of the image, then use an LLM for linguistic reasoning to perform emotion inference (see Fig. 1). Our captioning method is called Narrative Captioning (NarraCap), and we compare it to a state-of-the-art ExpansionNet [44] captioning.

1) Narrative Captioning: Our zero-shot Narrative Captioning (NarraCap) makes use of templates and the vision language model CLIP [45]. First, given an image with the bounding box of a person, we extract the cropped bounding

box and pass it along with a gender/age category (baby girl, baby boy, girl, boy, man, woman, elderly man, elderly woman) to CLIP to understand **who** is in the picture. Next, we pass the entire image to understand **what** is happening in the image by selecting the action with the maximum probability. The action list comprises 848 different actions extracted from the Kinetics-700 [46], UCF-101 [47], and HMDB datasets [48].

We then add the **how** aspect of the image by passing the cropped bounding box through CLIP, along with 889 signals (available on our website¹) filtered from over 1000 social signals derived from a guide on writing about emotion [38]. Using trial and error, we found that the best approach comes from selecting signals that, when paired with an image in CLIP, returns a probability higher than $mean + 9 * std$ of the class label scores. To provide additional context, we use 224 environmental descriptors from a writer’s guide to urban [49] and rural [50] settings to describe **where** the person in the scene is located. The prompts we selected for CLIP are as follows: ‘A photo of a(n) [gender/age/location]’, and ‘A photo of a person who is(has) [action/physical signals]’. Examples of narrative captions (NarraCap) can be found in Fig. 2.

2) *ExpansionNet Captioning*: We also evaluate a baseline captioning method, ExpansionNet [44], a fast end-to-end trained model for image captioning. The model achieves state of the art performance over the MS-COCO 2014 captioning challenge, and was used as a backbone for a recent approach trained on EMOTIC [51], and serves as a baseline to our NarraCap approach.

3) *Caption to Emotion using LLM (Zero Shot)*: Following captioning, we provide the caption, along with a prompt, to GPT-4. The prompt asks for the top emotion labels understood from the caption: “<caption> From suffering, pain, [...], and sympathy, pick the top labels that this person is feeling at the same time.” We also utilized an open-source LLM, Mistral 7B [52], which incorporates grouped-query attention [53] and sliding window attention [54], [55] techniques to address common LLM limitations such as computational power and memory requirements. Mistral, with its 7 billion parameters, outperforms the best released 34-billion-parameter model, Llama 1 [56], in reasoning tasks. The combination of a relatively low parameter count and high performance makes Mistral a potential option for applications in robotics. The prompt for Mistral is as follows: “<caption> From suffering, pain, [...], and sympathy, the top labels that this person is feeling at the same time are:”

4) *Mistral (Fine-Tuned)*: Fine-tuning LLMs [57], [58], [59] has been demonstrated as an effective strategy to improve their performance. In this paper, using quantization, LoRA [60], [61] and NarraCap, we finetune Mistral on the emotion recognition task. For Mistral, we experiment by fine-tuning on the Emotic validation set and augmentation.

B. End-to-End Vision Language Models

We next explore 3 vision language models (VLMs): CLIP, a closed-source (GPT-4) and an open source (LLaVA) VLM.

We also study the effect of prompt engineering, as well as fine-tuning on the open-source model (LLaVA).

1) *CLIP (Zero Shot)*: CLIP jointly trains an image encoder and a text encoder by maximizing the cosine similarity between related (image, text) pairs and minimizing the cosine similarity between the irrelevant pairs:

$$\text{logits} = \text{np.dot}(I_e, T_e^T) * \text{np.exp}(t) \quad (1)$$

where I_e is image feature embeddings and T_e is the text feature embeddings. CLIP can be used to perform zero shot classification by comparing distances between an image and various texts in a multimodal embedding space. We used the images from EMOTIC and compared the distances with each of the emotion labels and selected the six (average number of ground truth labels in validation set) labels with highest probabilities as our labels.

2) *GPT-4 Vision and LLaVA (Zero Shot)*: GPT-4 Vision is a proprietary model from OpenAI that can provide text-based responses given an image and text input. Large Language and Vision Assistant (LLaVA) [62] is an open source multi-purpose multimodal model designed by combining CLIP’s visual encoder [45] and LLAMA’s language decoder [56]. The model is fine-tuned end-to-end on the language-image instruction-following data generated using GPT-4 [63].

3) *LLaVA (Fine-Tuned)*: We use the EMOTIC data to finetune LLaVA with LoRA [60]² on the emotion recognition task. We experiment by fine-tuning LLaVA on the EMOTIC training set (17077 images, and 23706 individuals), EMOTIC validation set (2087 images, and 3330 individuals), and on a small dataset, created by selecting 100 images at random from the validation set. Furthermore, we perform data augmentation by shuffling the ground truth labels for each image in the validation set and using each image 3 times with different shuffled labels.

4) *Prompt Engineering*: We used the images from EMOTIC and a text prompt: “From suffering, pain, [...], and sympathy, pick the top labels that the person in the red bounding box is feeling at the same time.” It has been shown that prompt engineering (e.g. chain of thought [64]) is an effective way to improve results. We also tested a prompt which included definitions of the emotions and specified a number of labels to output.

IV. EXPERIMENTS

Our experiment is focused on the EMOTIC dataset, which covers 26 different labels. The related emotion recognition task is to provide a list of emotion labels that matches those chosen by annotators. The training set (70%) was annotated by 1 annotator, where validation (10%) and test (20%) sets were annotated by 5 and 3 annotators respectively. While previous work on emotion recognition tasks [7], [9], [30], utilize the mean Average Precision (mAP) as a metric, in this work, outputs are textual descriptions indicating the labels the person is feeling, rather than probabilities. Therefore, we could not employ mAP, instead, we used precision, recall,

¹<https://yasaman-etesam.github.io/Contextual-Emotion-Recognition/>

²<https://github.com/haotian-liu/LLaVA>

F1 score, hamming loss, which demonstrates the average rate at which incorrect labels are predicted for a sample, and subset accuracy, which requires the predicted set of labels for a sample to exactly match the actual set of labels. These metrics are implemented using the scikit-learn library [65]. We compare the following methods:

EMOTIC Along with the dataset, [7] introduced a two-branched CNN-based network baseline. The first branch extracts body related features and the second branch extracts scene-context features. Then a fusion network combines these features and estimates the output. For the EMOTIC dataset, using the provided code³ and thresholds calculated from the validation set, we obtained the output labels on the test set and then calculated the target metrics. This approach is the only traditional method with reproducible code.

Emoticon Motivated by Frege’s principle [66], [9] proposed an approach by combining three different interpretations of context. They used pose and face features (context1), background information (context2), and interactions/social dynamics (context3). They used a depth map to model the social interactions in the images. Later, they concatenate these different features and pass it to fusion model to generate outputs [9]. Unfortunately, the code for this project was not made available by the authors, and we could not reproduce the reported results. Consequently, we cannot provide a reliable comparison with this approach.

Random We consider selecting either 6 (average number of labels per person in validation set) emotions randomly from all possible labels (*Rand*) or selecting 6 labels randomly where the weights are determined by the number of times each emotion is repeated in the validation set (*Rand(W)*).

Majority This Majority baseline selects the top 6 most common emotions in the validation set (engagement, anticipation, happiness, pleasure, excitement, confidence) as the predicted labels for all test images (*Maj*).

CLIP For the CLIP model, we employed the *clip-vit-base-patch32*. While utilizing the *clip-vit-large-patch14-336* model did enhance the F1 score for the clip-only method to 19.60, its use significantly increased processing time, particularly for generating NarraCap captions. Therefore, to maintain consistency in our reporting and efficiency in our processing, we present results using the *clip-vit-base-patch32* model. The prompt used here is: "*The person in the red bounding box is feeling {emotion label}*". Additionally, we employed Grad-CAM [67], to generate saliency maps, allowing us to visually highlight the areas within images that significantly influenced the model’s decisions.

Captions+GPT-4 After generating captions, we pass the captions to *gpt-4* (*gpt-4-0613*). We utilized GPT-4 with the temperature parameter set to 0 and the maximum token count set to 256. Additionally, the frequency penalty, presence penalty, and top_p were configured to 0, 0, and 1, respectively. While adjusting these parameters could potentially enhance the model’s performance, we refrained from hyperparameter

tuning for this task due to associated costs.

GPT-4 Vision Using *gpt-4-vision-preview*, we input EMOTIC test images, with parameters set similarly to GPT-4. In this experiment, we tested both the prompt mentioned in Section 3.2.4, and also the inclusion of label definitions provided by EMOTIC to GPT, and requesting the six most likely labels.

LLaVA As GPT4-Vision, we tested both the prompt mentioned in Section 3.2.4, and also the inclusion of label definitions provided by EMOTIC to GPT, and requesting six most likely labels. LLaVA fine-tuning was performed on four A40 48GB GPUs.

Mistral We used huggingface⁴ to run and finetune Mistral on a RTX 3090 Ti GPU. We used maximum new tokens of 256 and repetition penalty equal to 1.15.

V. RESULTS AND DISCUSSION

The results for zero shot methods are shown in Table I, and example images with captions in Fig. 2.

We observe that fine-tuning LLaVA with an augmented validation set provides the best overall F1 score. In Fig. 2, we can see that LLaVA fine-tuned on the validation set predicts more labels than when trained on the training set. One explanation is that the number of annotators for the training and test sets is 1 and 3, respectively. Since we utilize the combined labels predicted by all annotators, the average number of ground truth labels per person is higher in the test set (4.42) than in the training set (1.96). This discrepancy leads the model trained on the training set to predict fewer labels than what is present in the test ground truth, causing the model to predict cautiously with high precision but miss many labels, resulting in low recall. To address this issue, we attempted fine-tuning on the validation set, which has 5 annotators and an average of 6.157 ground truth labels per person. We also experimented with fine-tuning on a small dataset, selecting 100 images at random from the validation set. This was to demonstrate that using a minimal amount of data can still yield reasonable results with vision language models (VLMs). Future work could try to balance the average number of labels in the training and test set. We also see that simple augmentation of the validation set, by shuffling the labels, improves performance. This may be due to the model learning that label ordering is not an important factor in the text output.

In Fig. 2, we observe that the EMOTIC baseline tends to predict many more labels than the other methods, which reduces its precision and overall F1 score. For an application where choosing a precise emotion label is more important than predicting all possible labels, fine-tuned LLaVA on the training set may be the most useful model.


It is evident that CLIP, which underperforms as indicated in Table I, misinterprets certain images, such as mistakenly attributing the emotion of embarrassment to a woman at the beach (Fig. 2). A deeper analysis using Grad-Cam-generated saliency maps (as seen in Fig. 3.1) offers a plausible explanation: CLIP may inaccurately associate images displaying

³<https://github.com/Tandon-A/emotic>


⁴<https://huggingface.co>




1. GT: Anticipation, Excitement, Happiness, Pleasure, Sympathy
NarraCap: This person is an elderly woman who is petting an animal (not a cat) in a nursing home. She is caretaking.
CLIP: **Doubt/Confusion**, Esteem, Sensitivity, Engagement, Sympathy, Affection
GPT4: Affection, Engagement, Happiness, Peace, Pleasure, Confidence, Sensitivity, Sympathy
GPT4-Vision*: Affection, Engagement, Happiness, Peace, Pleasure, Esteem
LLaVA: Happiness, Engagement, Affection, Sensitivity
LLaVA†: Happiness, Pleasure, Affection, Sympathy, Anticipation, Engagement
EMOTIC: Affection, Anticipation, Engagement, **Fatigue**, Happiness, Peace, Sympathy, Yearning



2. GT: Engagement, Happiness, Peace
NarraCap: This person is a woman who is packing in an airport. She is chest bumping with others. She is clutching at another person for support. She is pulling someone into a side hug. She is putting an arm around someone's shoulders.
CLIP: **Sensitivity**, **Disconnection**, Happiness, **Suffering**, Surprise, Affection
GPT4: Affection, Engagement, Anticipation, Excitement, Happiness, Confidence, Surprise
GPT4-Vision*: Happiness, Affection, Engagement, Confidence, Peace, Excitement
LLaVA: Happiness, Engagement, Excitement, Anticipation
LLaVA†: Excitement, Happiness, Pleasure, Affection, Anticipation
EMOTIC: Anticipation, Engagement, Happiness




3. GT: Fatigue, Peace
NarraCap: This person is a boy who is sleeping in a bookstore. He is curling up to take up less space. He is falling asleep in odd places at odd times. He is flipping through brochures or other reading material at hand. He is sleeping.
CLIP: Sympathy, **Engagement**, **Suffering**, Disquietment, Disconnection, **Fatigue**
GPT4: **Fatigue**, Disquietment, **Doubt/Confusion**, Disconnection, **Yearning**, Sadness
GPT4-Vision*: **Fatigue**, Disquietment, Disconnection, **Peace**, **Engagement**, **Doubt/Confusion**
LLaVA: **Fatigue**, Disconnection
LLaVA†: **Fatigue**, Disconnection, **Peace**, **Suffering**, **Engagement**
EMOTIC: **Annoyance**, **Aversion**, **Disapproval**, Disquietment, **Doubt/Confusion**, **Engagement**, **Fatigue**, **Happiness**, **Peace**, Sensitivity, **Yearning**



4. GT: Confidence, Esteem, Excitement, Happiness, Pleasure
NarraCap: This person is a boy who is skiing in a ski resort. He has hands raised into the air. He has hands up, backing away. He is waving enthusiastically. He is waving the arms, using grand gestures.
CLIP: **Esteem**, **Doubt/Confusion**, **Surprise**, Happiness, **Fear**, Confidence
GPT4: **Excitement**, Happiness, Anticipation, Engagement, Pleasure
GPT4-Vision*: Engagement, Confidence, Happiness, Peace, Excitement, Pleasure
LLaVA: Excitement, Anticipation, Engagement
LLaVA†: Excitement, Happiness, Confidence, Anticipation, Engagement
EMOTIC: Anticipation, Confidence, Engagement, Excitement, Happiness



5. GT: Happiness, Peace, Pleasure
NarraCap: This person is a woman who is windsurfing in a beach. She has chin dipping to the chest. She is participating in relaxing activities.
CLIP: Happiness, **Sympathy**, **Embarrassment**, **Peace**, Confidence, Esteem
GPT4: Confidence, Engagement, Happiness, **Peace**, Pleasure, Excitement, Anticipation
GPT4-Vision*: **Peace**, Pleasure, Confidence, Engagement, Happiness, Anticipation
LLaVA: Excitement, Anticipation, Happiness
LLaVA†: Engagement, Happiness, Anticipation, Excitement, Confidence, Pleasure
EMOTIC: **Affection**, Anticipation, Confidence, **Embarrassment**, Engagement, Excitement, **Fear**, Happiness, Pleasure, **Surprise**



6. GT: Anticipation, Confidence, Doubt/Confusion, Engagement, Excitement, Happiness
NarraCap: This person is a boy who is skijetting in a ski resort. He has shoulders curling forward.
CLIP: **Pain**, **Surprise**, **Doubt/Confusion**, Confidence, **Fear**, **Suffering**
GPT4: Excitement, Anticipation, Engagement, Pleasure, Confidence
GPT4-Vision*: Engagement, Anticipation, Engagement, Disquietment, **Fatigue**, Confidence, **Doubt/Confusion**
LLaVA: Excitement, Anticipation, Engagement
LLaVA†: Confidence, Anticipation, Engagement, Excitement, Happiness
EMOTIC: Anticipation, Confidence, Engagement, Excitement



7. GT: Engagement, Sadness, Suffering, Sympathy
NarraCap: This person is a man who is crying in a refugee camp. He is crying with the person. He is crying, wailing, begging for help. He is crying.
CLIP: Anger, **Fatigue**, **Affection**, Sensitivity, Sadness, Sympathy
GPT4: **Suffering**, Pain, **Fear**, Disquietment, **Doubt/Confusion**, Sadness, **Yearning**, Disapproval, Annoyance, **Fatigue**
GPT4-Vision*: **Suffering**, Sadness, Disquietment, Pain, Sympathy, **Fatigue**
LLaVA: Sadness, Pain, **Embarrassment**
LLaVA†: Sadness, **Suffering**, Pain, Disapproval, Disquietment, Disconnection, Sensitivity, **Fear**
EMOTIC: **Affection**, Anticipation, **Aversion**, Disquietment, **Doubt/Confusion**, **Embarrassment**, Engagement, **Fear**, **Yearning**, Pleasure, Sadness, Sensitivity, **Surprise**, Sympathy, **Happiness**



8. GT: Happiness, Pleasure
NarraCap: is a man who is shaking hands in a board room. He is engaging excitedly with others. He is putting one's hands in one's pockets. He is using a strong, businesslike handshake.
CLIP: **Suffering**, **Disconnection**, Sensitivity, Esteem, Confidence, Engagement
GPT4: Confidence, Engagement, Happiness, Excitement, Anticipation, Esteem
GPT4-Vision*: Confidence, Engagement, Happiness, **Peace**, Pleasure, Esteem
LLaVA: Happiness, Engagement, Confidence
LLaVA†: Confidence, Happiness, Esteem, Pleasure, Excitement, Engagement, Affection
EMOTIC: Affection, **Annoyance**, Anticipation, Confidence, Disapproval, **Doubt/Confusion**, Engagement, Esteem, Happiness, Pleasure, **Sympathy**



9. GT: Affection, Engagement, Happiness, Pleasure
NarraCap: This person is a girl who is grooming a dog in a pet store. She is displaying affection with friends or loved ones. She is gripping one's own throat. She is pulling loved ones close to protect them.
CLIP: **Sadness**, Engagement, **Peace**, Happiness, Pleasure, Affection
GPT4: Affection, Confidence, Engagement, Happiness, **Peace**, Pleasure, Anticipation, **Surprise**.
GPT4-Vision*: Affection, Happiness, **Peace**, Pleasure, Engagement, Excitement
LLaVA: Affection, Happiness, Engagement
LLaVA†: Happiness, Affection, Pleasure, Sympathy, Engagement, Excitement
EMOTIC: Affection, **Annoyance**, Disquietment, **Doubt/Confusion**, Engagement, **Fatigue**, Happiness, **Peace**, Pleasure, Sensitivity, **Suffering**, **Surprise**, Sympathy

Fig. 2. Qualitative analysis of EMOTIC images showcasing ground truth (GT) labels, NarraCap-generated captions, and inferred labels using (a) zero-shot methods: CLIP, GPT4 (NarraCap+GPT4), GPT4-Vision, LLaVA; and (b) trained methods: LLaVA finetuned on the validation set with augmentation (LLaVA[†]), and EMOTIC.

bare skin with embarrassment. Additionally, CLIP seems to exhibit spurious correlations of body language, predicting emotions like surprise and fear in response to raised arms (Fig.3.2), or sadness from the positioning of hands near the

	Model	↑Precision (%)	↑Recall (%)	↑F1 Score (%)	↓Hamming (%)	↑S-acc (%)
Zero shot	Majority	11.41 ± 0.05	23.08 ± 0.00	15.01 ± 0.05	17.24 ± 0.09	0.76 ± 0.09
	Rand6	16.90 ± 0.13	23.12 ± 0.34	14.90 ± 0.15	32.29 ± 0.07	0.00 ± 0.00
	Rand6-weighted	17.02 ± 0.15	23.17 ± 0.20	19.45 ± 0.17	22.67 ± 0.08	0.01 ± 0.01
	ExpNet + GPT4	24.94 ± 0.58	23.57 ± 0.27	22.29 ± 0.30	17.27 ± 0.09	1.79 ± 0.13
	NarraCap + GPT4	25.50 ± 0.32	33.37 ± 0.42	26.67 ± 0.30	21.26 ± 0.11	0.82 ± 0.09
	NarraCap + Mistral	17.56 ± 0.09	64.53 ± 0.59	23.89 $^{0.14\pm}$	52.70 ± 0.24	0.01 ± 0.01
	CLIP	21.77 ± 0.19	28.58 ± 0.35	16.97 ± 0.18	31.72 ± 0.07	0.00 ± 0.00
	LLaVA	33.78 ± 0.86	21.38 ± 0.30	22.86 ± 0.32	15.02 ± 0.08	0.99 ± 0.10
	LLaVA*	27.77 ± 0.63	18.51 ± 0.25	19.58 ± 0.27	16.04 ± 0.08	0.78 ± 0.09
	GPT-4 vision	29.07 ± 0.37	27.48 ± 0.37	26.12 ± 0.28	16.72 ± 0.12	1.90 ± 0.14
GPT-4 vision*	37.48 ± 0.79	38.35 ± 0.36	34.47 ± 0.35	16.95 ± 0.08	0.67 ± 0.08	
Trained	EMOTIC	25.02 ± 0.28	35.07 ± 0.49	28.83 ± 0.33	19.35 ± 0.14	2.73 ± 0.17
	Mistral-F (val set augmented)	18.01 ± 0.09	78.40 ± 0.45	26.41 ± 0.13	54.16 ± 0.19	0.00 ± 0.00
	LLaVA-F (train set)	54.27 ± 1.42	16.81 ± 0.30	22.73 ± 0.37	13.17 ± 0.07	1.72 ± 0.13
	LLaVA-F (val set)	32.55 ± 0.55	42.95 ± 0.42	34.42 ± 0.34	17.14 ± 0.09	0.78 ± 0.09
	LLaVA-F (val set augmented)	38.71 ± 0.55	39.52 ± 0.42	36.83 ± 0.37	14.13 ± 0.08	2.90 ± 0.17
	LLaVA-F (val set 100 samples)	31.36 ± 0.37	40.41 ± 0.40	33.85 ± 0.33	17.30 ± 0.09	1.28 ± 0.12

TABLE I

PERFORMANCE METRICS OF VARIOUS MODELS USING MACRO AVERAGE. "S-ACC" REPRESENTS SUBSET ACCURACY, "-F" INDICATES THAT THIS IS A FINE-TUNED VERSION, AND "*" INDICATES PROMPT ENGINEERING (DEFINITIONS + REQUEST 6 LABELS, AVE. IN VALIDATION SET)



Fig. 3. CLIP saliency maps using Grad-Cam [67]. It identifies the regions in an input image that most significantly influence the classification score by leveraging the gradients of the score relative to the last convolutional layer's feature map.

face (Fig. 3.3).

In the captioning combined with GPT-4 analysis, NarraCap proves to be more effective than ExpNet in aiding GPT-4's understanding of emotions. However, it ranks as the second-best zero-shot approach. GPT-4 Vision with prompt engineering emerges as the top performer among zero-shot methods, surpassing EMOTIC, which was trained on the EMOTIC training set.

How does captioning + LLM compare to the end-to-end VLM approach? In addition to our experiments in Table I, we performed an additional study on a smaller test set. Yang et al. [68] recruited an annotator fluent in North American English to manually generate captions for 387 images, encompassing 14 negative emotion categories: suffering, annoyance, pain, sadness, anger, aversion, disquietment, doubt/confusion, embarrassment, fear, disapproval, fatigue, disconnection, and sensitivity. This focus on negative emotions stemmed from their comparatively poor recognition across all methods tested, relative to positive emotions. The outcomes for all methodologies applied to this dataset are detailed in Table II. On this smaller, challenging test set, our best proposed zero-shot captioning + LLM approach (NarraCap+GPT-4) resulted

	Zero-shot				Trained	
	NC+GPT4	HNC+GPT4	GPT-Vis	LLaVA	LLaVA-f	EMOTIC
F1	26.19	34.17	35.79	27.08	42.14	26.50

TABLE II

HUMAN-GENERATED NARRATIVE CAPTIONS + GPT-4 (HNC) VS. GPT-4 VISION (GPT-VIS) ON NEGATIVE EMOTION SUBSET [68]

in an F1 score of **26.19**. The GPT4-Vision zero-shot VLM approach attained an F1 score of **35.79**. This disparity appears large; however, leveraging human-generated captions with GPT-4 (LLM) achieved an F1 score of **34.17**. This indicates that while the *automatic captioning + LLM* method does not reach the VLM performance, human-level captioning when coupled with LLMs provides nearly comparable performance and outperforms the traditional EMOTIC baseline.

How do different prompts affect the results? Selecting an appropriate prompt for LLMs and VLMs is crucial for optimizing their performance. However, the same prompt can affect different models in varied ways. As shown in Table I, incorporating label definitions and requesting the top 6 labels significantly enhances the results for GPT-Vision, yet it adversely impacts LLaVA's performance. Thus, tailoring prompts to the specific characteristics and capabilities of each model may be necessary to achieve the best outcomes. Furthermore, following [45], we adjusted the phrasing of the CLIP's input prompt from "*The person in the red bounding box is [emotion label]*" to "*A photo of a person in a red bounding who is [emotion label]*." Interestingly, this modification led to a decrease in the CLIP approach's F1 score to 13.76.

How does the number of people in the image impact the emotion recognition outcome? We evaluate different

		CLIP	NarraCap	GPT4-V	LLaVA	LLaVA-f	EMOTIC
P	1	22.07	26.44	37.53	32.06	41.27	27.72
	2	21.52	26.06	36.64	36.10	38.07	23.36
	>2	20.97	23.18	34.31	31.30	35.02	22.71
R	1	28.76	32.94	39.07	22.79	41.41	41.62
	2	28.11	35.02	38.67	22.15	41.21	35.09
	>2	27.45	31.74	36.54	18.39	35.66	28.51
F1	1	16.20	26.53	34.75	23.53	38.42	32.50
	2	16.67	27.61	34.28	23.38	37.35	27.60
	>2	16.81	25.01	32.57	19.77	33.34	24.60

TABLE III

TEST SET PRECISION(P), RECALL(R), AND F1 SCORE WHEN SEPARATED INTO 1, 2, AND MULTIPLE PEOPLE IN THE IMAGE. HERE, NARRACAP IS PAIRED WITH GPT-4

methods based on the number of people in the image: one, two, or multiple people. As shown in Table III, the precision, recall, and F1 score tend to decrease as the number of people increases. This reduction in performance can be attributed to the more complex situations that arise when there are more people in a scene [69]. Furthermore, NarraCap does not account for human interactions, and vision language models (VLMs) struggle with identifying the specific individual referred to in a prompt. This challenge is partly due to the models' limitations in interpreting visual markers (bounding boxes), which are crucial for distinguishing among multiple subjects in an image [70]. For the EMOTIC model, which was trained on the training set, it was observed that while it surpassed the performance of NarraCap, a zero-shot approach, for images featuring a single person, it was less effective in handling images with two or more people. This discrepancy suggests that NarraCap demonstrates superior performance in more complex scenarios involving multiple individuals.

How do age, gender, activity, environment, and physical signals affect the results? One of the advantages of the NarraCap approach is that it provides a way to explicitly select image details to include for inference and perform ablations using the text representation. To assess the effect of gender, instead of using specific gender labels such as "a baby boy," "a baby girl," etc. we only utilized the labels "a female" and "a male." Furthermore, to examine gender, we modified the label list to "a baby," "a kid," "an adult," and "an elderly person." To investigate the impact of activity, environment, and physical signals, we excluded those components from the captions. The findings from each study, conducted on a set of 1000 images randomly selected from the validation set, are summarized in Table IV. This table reveals that the *action* depicted in an image had the most significant impact on the outcome, followed by the *environment*. These insights suggest that future research on image caption generation may focus on understanding image actions and contexts, as they are keys to create accurate and relevant captions.

VI. LIMITATIONS

The current study evaluating emotional theory of mind within the EMOTIC image dataset has its limitations. It primarily focused on various OpenAI models, including the zero-shot classifier CLIP, the large language model GPT-4, and the vision-language model GPT-4 Vision. Additionally,

Ablation Settings						F1	diff.	SE
age	gender	environment	action	physical signals				
✓	✓	✓	✓	✓	✓	29.67	-	0.78
-	✓	✓	✓	✓	✓	28.41	-1.26	0.73
✓	-	✓	✓	✓	✓	29.25	-0.42	0.78
✓	✓	-	✓	✓	✓	27.27	-2.4	0.78
✓	✓	✓	-	✓	✓	23.67	-6	0.70
✓	✓	✓	✓	-	-	29.47	-0.2	0.85

TABLE IV

ABLATION STUDY USING NARRACAP+GPT4 ON 1000 RANDOMLY SELECTED IMAGES FROM THE VALIDATION SET. HERE, WE PRESENT F1 SCORES ALONG WITH THEIR CORRESPONDING STANDARD ERRORS.

two open-source methods were examined: LLaVA (VLM) and Mistral (LLM). Although state-of-the-art zero-shot techniques were employed, fine-tuning some models (GPT-4, GPT-4 Vision) was not feasible due to their proprietary, closed-source nature. In addition, it is not possible to know the extent that the vision language models (i.e. GPT-4 Vision and LLaVA) were exposed to the EMOTIC test set.

Another limitation is the absence of certain traditional emotion recognition models (e.g., emoticon) from our study, as their code was not made publicly available by the authors and our attempts to re-implement these models failed to replicate the reported results.

In addition, we noticed that the EMOTIC dataset, while being one of the most challenging image emotion datasets including context, also has small imperfections, including some bounding boxes that contain 2 people instead of only one (Fig. 2.2). Although the case described is rare, a future study could evaluate on other datasets, e.g. one person emotion expression datasets without context, which is considered to be a simpler task. For NarraCap, captions did not describe the social interactions or interactions with objects, which if added may increase performance. Moreover, for the activity and environment detection using CLIP, we performed a standard evaluation with a limited set of classes without a "don't know" or null class, resulting in some mis-captioning.

VII. CONCLUSION

In this study, we delve into the potential of vision language models (VLMs) and large language models (LLMs) for assessing visual emotional theory of mind. Our findings reveal that the zero-shot approach of GPT-4 Vision with prompt engineering could outperform the trained EMOTIC model. Furthermore, our success in enhancing the performance beyond conventional methods by fine-tuning LLaVA, an open-source VLM, on a modest dataset underscores the profound potential of these models to comprehend human emotions. Further research could explore improving the narrative captions by adding other contextual factors to the caption, such as human-object interactions and relationships with other people in the image. Further studies on the characteristics of emotionally comprehensive captioning, coupled with GPT or with human evaluators could be done. Additionally, enhancing the ability of VLMs to accurately recognize visual markers, such as bounding boxes, could significantly boost their performance.

REFERENCES

- [1] R. W. Picard, *Affective computing*. MIT press, 2000.
- [2] L. Barrett et al., “Emotional expressions reconsidered: Challenges to inferring emotion from human facial movements,” *Psychological science in the public interest*, vol. 20, no. 1, pp. 1–68, 2019.
- [3] M. Pantic and L. J. Rothkrantz, “Expert system for automatic analysis of facial expressions,” *Image and Vision Computing*, vol. 18, no. 11, pp. 881–905, 2000.
- [4] K. Schindler et al., “Recognizing emotions expressed by body pose: A biologically inspired neural model,” *Neural networks*, vol. 21, no. 9, pp. 1238–1246, 2008.
- [5] L. F. Barrett, B. Mesquita, and M. Gendron, “Context in emotion perception,” *Current directions in psychological science*, vol. 20, no. 5, pp. 286–290, 2011.
- [6] L. F. Barrett, *How emotions are made: The secret life of the brain*. Pan Macmillan, 2017.
- [7] R. Kosti et al., “Context based emotion recognition using emotic dataset,” *PAMI*, vol. 42, no. 11, pp. 2755–2766, 2019.
- [8] N. Le et al., “Global-local attention for emotion recognition,” *Neural Computing and Applications*, vol. 34, no. 24, pp. 21 625–21 639, 2022.
- [9] T. Mittal et al., “Emoticon: Context-aware multimodal emotion recognition using frege’s principle,” in *CVPR*, 2020, pp. 14 234–14 243.
- [10] D. C. Ong et al., “Computational models of emotion inference in theory of mind: A review and roadmap,” *Topics in cognitive science*, vol. 11, no. 2, pp. 338–357, 2019.
- [11] D. Lopez-Paz, “From dependence to causation,” *arXiv*, 2016.
- [12] E. Jang et al., “Bc-z: Zero-shot task generalization with robotic imitation learning,” in *CoRL*, 2022.
- [13] A. Vaswani et al., “Attention is all you need,” *NeurIPS*, vol. 30, 2017.
- [14] T. Brown et al., “Language models are few-shot learners,” *Advances in neural information processing systems*, vol. 33, pp. 1877–1901, 2020.
- [15] A. Chowdhery et al., “Palm: Scaling language modeling with pathways,” *arXiv*, 2022.
- [16] S. Antol et al., “Vqa: Visual question answering,” in *ICCV*, 2015, pp. 2425–2433.
- [17] O. Vinyals et al., “Show and tell: A neural image caption generator,” in *CVPR*, 2015, pp. 3156–3164.
- [18] M. Sap et al., “Socialliqa: Commonsense reasoning about social interactions,” *arXiv*, 2019.
- [19] Y. Bisk et al., “Piqa: Reasoning about physical commonsense in natural language,” in *AAAI*, vol. 34, no. 05, 2020, pp. 7432–7439.
- [20] X. L. Li et al., “A systematic investigation of commonsense knowledge in large language models,” in *EMNLP*, 2022, pp. 11 838–11 855.
- [21] R. Mao et al., “The biases of pre-trained language models: An empirical study on prompt-based sentiment analysis and emotion detection,” *IEEE Transactions on Affective Computing*, 2022.
- [22] M. Sap et al., “Neural theory-of-mind? on the limits of social intelligence in large lms,” *arXiv*, 2022.
- [23] J. Zhang et al., “Vision-language models for vision tasks: A survey,” *PAMI*, 2024.
- [24] S. Pratt et al., “What does a platypus look like? generating customized prompts for zero-shot image classification,” in *ICCV*, 2023, pp. 15 691–15 701.
- [25] Y. Long et al., “Fine-grained visual–text prompt-driven self-training for open-vocabulary object detection,” *IEEE Transactions on Neural Networks and Learning Systems*, 2023.
- [26] A. Dhall et al., “Finding happiest moments in a social context,” in *ACCV*. Springer, 2013, pp. 613–626.
- [27] Y. Huang et al., “Emotion recognition based on body and context fusion in the wild,” in *ICCV*, 2021, pp. 3609–3617.
- [28] J. Lee et al., “Context-aware emotion recognition networks,” in *ICCV*, 2019, pp. 10 143–10 152.
- [29] S. Thuseethan et al., “Boosting emotion recognition in context using non-target subject information,” in *IJCNN*. IEEE, 2021, pp. 1–7.
- [30] A. Mittel and S. Tripathi, “Peri: Part aware emotion recognition in the wild,” in *ECCV 2022 Workshops*. Springer, 2023, pp. 76–92.
- [31] W. Li et al., “Human emotion recognition with relational region-level analysis,” *IEEE Transactions on Affective Computing*, vol. 14, no. 1, pp. 650–663, 2021.
- [32] M.-H. Hoang et al., “Context-aware emotion recognition based on visual relationship detection,” *IEEE Access*, vol. 9, pp. 90 465–90 474, 2021.
- [33] D. Yang et al., “Context de-confounded emotion recognition,” in *CVPR*, 2023, pp. 19 005–19 015.
- [34] L. et al., “The role of language in emotion: Predictions from psychological constructionism,” *Frontiers in psychology*, vol. 6, p. 444, 2015.
- [35] M. D. Lieberman et al., “Putting feelings into words,” *Psychological science*, vol. 18, no. 5, pp. 421–428, 2007.
- [36] K. A. Lindquist and M. Gendron, “What’s in a word? language constructs emotion perception,” *Emotion Review*, vol. 5, no. 1, pp. 66–71, 2013.
- [37] S. Keen and S. Keen, “Narrative emotions,” *Narrative Form: Revised and Expanded Second Edition*, pp. 152–161, 2015.
- [38] B. Puglisi and A. Ackerman, *The emotion thesaurus: A writer’s guide to character expression*. JADD Publishing, 2019, vol. 1.
- [39] W. Wang et al., “A survey of zero-shot learning: Settings, methods, and applications,” *TIST*, 2019.
- [40] S. Liu et al., “Generalized zero-shot learning with deep calibration network,” *NeurIPS*, 2018.
- [41] F. Pourpanah et al., “A review of generalized zero-shot learning methods,” *PAMI*, 2022.
- [42] Z. Abderrahmane et al., “Haptic zero-shot learning: Recognition of objects never touched before,” *Rob. Auton. Syst.*, 2018.
- [43] R. Socher et al., “Zero-shot learning through cross-modal transfer,” *NeurIPS*, 2013.
- [44] J. Hu et al., “Expansionnet v2: Block static expansion in fast end to end training for image captioning,” *arXiv*, 2022.
- [45] A. Radford et al., “Learning transferable visual models from natural language supervision,” in *ICML*. PMLR, 2021, pp. 8748–8763.
- [46] L. Smaira et al., “A short note on the kinetics-700-2020 human action dataset,” *arXiv*, 2020.
- [47] K. Soomro et al., “Ucf101: A dataset of 101 human actions classes from videos in the wild,” *arXiv*, 2012.
- [48] H. Kuehne et al., “Hmdb: a large video database for human motion recognition,” in *ICCV*. IEEE, 2011, pp. 2556–2563.
- [49] B. Puglisi and A. Ackerman, *The Urban Setting Thesaurus: A Writer’s Guide to City Spaces*. JADD Publishing, 2016, vol. 5.
- [50] —, *The Rural Setting Thesaurus: A Writer’s Guide to Personal and Natural Places*. JADD Publishing, 2016, vol. 4.
- [51] W. de Lima Costa et al., “High-level context representation for emotion recognition in images,” in *CVPR Workshops*, June 2023, pp. 326–334.
- [52] A. Q. Jiang et al., “Mistral 7b,” *arXiv*, 2023.
- [53] J. Ainslie et al., “Gqa: Training generalized multi-query transformer models from multi-head checkpoints,” *arXiv*, 2023.
- [54] R. Child et al., “Generating long sequences with sparse transformers,” *arXiv*, 2019.
- [55] I. Beltagy et al., “Longformer: The long-document transformer,” *arXiv*, 2020.
- [56] H. Touvron et al., “Llama: Open and efficient foundation language models,” *arXiv*, 2023.
- [57] S. Min et al., “Metaicl: Learning to learn in context,” *arXiv*, 2021.
- [58] L. Ouyang et al., “Training language models to follow instructions with human feedback, 2022,” *arXiv*, vol. 13, p. 1, 2022.
- [59] H. Liu et al., “Few-shot parameter-efficient fine-tuning is better and cheaper than in-context learning,” *NeurIPS*, vol. 35, pp. 1950–1965, 2022.
- [60] E. J. Hu et al., “Lora: Low-rank adaptation of large language models,” *arXiv*, 2021.
- [61] T. Dettmers et al., “Qlora: Efficient finetuning of quantized llms,” *NeurIPS*, vol. 36, 2024.
- [62] H. Liu et al., “Visual instruction tuning,” *arXiv*, 2023.
- [63] OpenAI, “Gpt-4 technical report,” 2023.
- [64] J. Wei et al., “Chain-of-thought prompting elicits reasoning in large language models,” *NeurIPS*, vol. 35, pp. 24 824–24 837, 2022.
- [65] F. Pedregosa et al., “Scikit-learn: Machine learning in Python,” *Journal of Machine Learning Research*, vol. 12, pp. 2825–2830, 2011.
- [66] M. D. Resnik, “The context principle in frege’s philosophy,” *Philosophy and Phenomenological Research*, vol. 27, no. 3, pp. 356–365, 1967.
- [67] R. R. Selvaraju et al., “Grad-cam: Visual explanations from deep networks via gradient-based localization,” in *ICCV*, 2017, pp. 618–626.
- [68] V. Yang et al., “Contextual emotion estimation from image captions,” *arXiv*, 2023.
- [69] E. A. Veltmeijer et al., “Automatic emotion recognition for groups: a review,” *IEEE Transactions on Affective Computing*, vol. 14, no. 1, pp. 89–107, 2021.
- [70] D. Wan et al., “Contrastive region guidance: Improving grounding in vision-language models without training,” *arXiv*, 2024.