

EmoVIT: Revolutionizing Emotion Insights with Visual Instruction Tuning

Hongxia Xie¹, Chu-Jun Peng², Yu-Wen Tseng²,
Hung-Jen Chen³, Chan-Feng Hsu³, Hong-Han Shuai³, Wen-Huang Cheng²

¹Jilin University

²National Taiwan University

³National Yang Ming Chiao Tung University

hongxiaxie@jlu.edu.cn, {d12922035, ywt, wenhuang}@csie.ntu.edu.tw,

{hjc.eed07g, cfhsu311510211.ee11, hhshuai}@nytu.edu.tw

Abstract

Visual Instruction Tuning represents a novel learning paradigm involving the fine-tuning of pre-trained language models using task-specific instructions. This paradigm shows promising zero-shot results in various natural language processing tasks but is still unexplored in vision emotion understanding. In this work, we focus on enhancing the model’s proficiency in understanding and adhering to instructions related to emotional contexts. Initially, we identify key visual clues critical to visual emotion recognition. Subsequently, we introduce a novel GPT-assisted pipeline for generating emotion visual instruction data, effectively addressing the scarcity of annotated instruction data in this domain. Expanding on the groundwork established by InstructBLIP, our proposed EmoVIT architecture incorporates emotion-specific instruction data, leveraging the powerful capabilities of Large Language Models to enhance performance. Through extensive experiments, our model showcases its proficiency in emotion classification, adeptness in affective reasoning, and competence in comprehending humor. The comparative analysis provides a robust benchmark for Emotion Visual Instruction Tuning in the era of LLMs, providing valuable insights and opening avenues for future exploration in this domain. Our code is available at <https://github.com/aimmemotion/EmoVIT>.

1. Introduction

Visual emotion recognition, a key area within artificial intelligence and computer vision, aims to predict human emotions based on visual cues such as facial expressions and body language. This technology is essential in bridging the gap between human affective states and machine understanding. Its diverse applications [19, 20, 23, 24], spanning

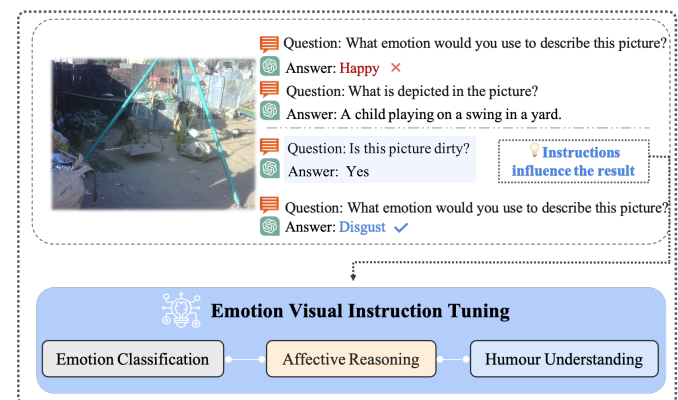


Figure 1. Illustration of the importance of instruction-following ability in visual emotion understanding.

from improving human-computer interaction to aiding in mental health assessment, underscore its significance. Accurate emotion recognition is vital for enhancing user experience and ensuring information security, as it helps prevent emotional manipulation and misinformation [13]. Developing robust emotion recognition models is not only a technical challenge but also a step towards more empathetic and intuitive AI systems, paving the way for more efficient and natural human-computer interactions.

The AI community has recently shown a growing interest in developing foundational vision models, e.g., Flamingo [1], LLaVA [8], BLIP2 [5]. These models excel in open-world visual understanding, tackling several vision tasks such as classification, detection, segmentation, and captioning. In contrast, current large-scale multimodal models are still in its infancy when it comes to emotion perception [7]. As illustrated in Fig. 1, when directly query the GPT-4 [11] about the emotional category of an image, the model tends to provide incorrect responses. However, the model delivers accurate responses when provided with revised instructions. To fully leverage the potential of ex-

isting vision-based large models, our approach is based on the concept of *Instruction Tuning*. This effective strategy is aimed at teaching language models to follow natural language instructions, a technique proven to enhance their generalization performance across unseen tasks [2, 8, 18].

In this work, we focus on *developing the model’s proficiency in understanding and following instructions related to emotional contexts*. This approach highlights the importance of fine-tuning the model’s instruction-following capabilities, enabling it to interpret and respond to emotional content effectively. This is achieved by leveraging its pre-existing knowledge base, thereby eliminating the necessity for an emotion-specific architectural framework.

To address the notable challenges encountered in Instruction Tuning for visual emotion recognition, especially the lack of specific instruction data, we introduce a novel self-generation pipeline explicitly crafted for visual emotion recognition by using GPT-4 [11]. This innovative pipeline excels in generating a diverse array of (image, instruction, output) instances, thereby notably enhancing the dataset with a more extensive and task-oriented variety of examples. This approach not only overcomes the challenge of limited data availability but also reduces the dependence on human labor. Therefore, it streamlines the process, enabling more efficient and effective emotion recognition.

Additionally, Instruction Tuning has been criticized for its emphasis on surface-level features like output patterns and styles, rather than achieving a profound comprehension and assimilation of tasks [26]. To tackle this issue and enhance the diversity and creativity of instruction data, our dataset includes instructions that demand complex reasoning, going beyond basic question-and-answer formats. This is further enriched by incorporating visual cues such as *brightness, colorfulness, scene type, object class, facial expressions, and human actions*. These aspects are pivotal in fostering a nuanced comprehension of visual emotions, thus allowing the model to generate more precise and contextually appropriate interpretations [24].

After generating the emotion visual instruction data, we propose an Emotion Visual Instruction Tuning (EmoVIT) framework, leveraging the foundation of InstructBLIP [2]. This framework incorporates an emotion-centric, instruction-aware module that proficiently guides Large Language Models (LLMs) in assimilating the nuances of emotion instructions. Our work signifies a paradigm shift, presenting a new era of instruction-based learning for visual emotion understanding that relies less on explicit training data. Remarkably, as shown in Fig. 2, our approach requires almost 50% of the training data typically needed yet exceeds the performance of previous visual emotion recognition methods and popular Visual Instruction Tuning methods.

Our contributions can be summarized as follows:




Figure 2. Performance comparison on EmoSet test set [24] (Accuracy %).

- We explore the potential of the Visual Instruction Tuning paradigm for emotion comprehension and introduce the concept of Emotion Visual Instruction Tuning.
- After thoroughly considering the unique characteristics of visual emotion recognition, we develop a novel GPT-assisted pipeline for generating emotion visual instruction data. This approach effectively bridges the gap in available annotated instruction data within this specific domain.
- Building upon the foundation of InstructBLIP, our EmoVIT architecture integrates emotion domain-specific instruction data, harnessing the robust capabilities of LLMs to boost performance. The extensive experiments demonstrate our model’s proficiency in emotion classification, affective reasoning, and comprehension of humour.

2. Related Work

2.1. Visual Emotion Recognition

A key challenge in visual emotion recognition is bridging the gap between an image’s visual cues and the emotions it portrays [12, 21, 25]. While traditional efforts, *e.g.*, Xu *et al.*’s multi-level dependent attention network [21], focus on visual models for emotional feature learning, recent advancements like EmoSet [24] offer rich emotion-laden datasets with 3.3 million images. The rise of multimodal models, such as the GPT series [11], has further propelled Vision-Language Recognition. However, fully leveraging these models in emotion recognition is an area ripe for exploration. Our work leads the way in utilizing large-scale models for Emotion Visual Instruction Tuning.

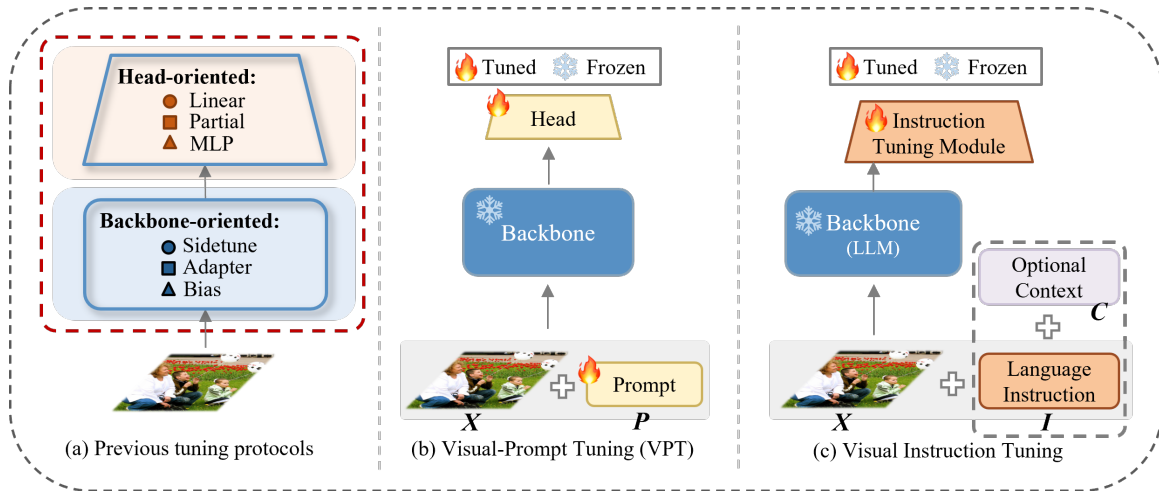


Figure 3. The comparison of different visual tuning paradigms.

2.2. Visual Instruction Tuning

Current Large Language Models (LLMs) have extensive knowledge bases, but their effectiveness depends on accurately interpreting human instructions due to a mismatch between training goals and user expectations. LLMs are trained to minimize prediction errors, whereas users expect helpful and safe instruction-following. Instruction Tuning addresses this by teaching models to follow natural language instructions, enhancing generalization to new tasks. FLAN [18] demonstrated that training a large model on instruction-based datasets improves zero-shot performance. This approach has extended to vision-language tasks, with BLIP2 [5] and LLaVA [8] adapting instruction-tuned LLMs for visual inputs. InstructBLIP [2] introduces instruction-aware visual feature extraction and the Q-Former, enabling more flexible, instruction-driven feature extraction.

As a novel area, visual emotion instruction tuning lacks benchmarks or guidelines for creating emotion instruction data. Our work pioneers the use of large-scale models to develop an emotion instruction data pipeline, overcoming the limitations of manual annotation.

3. Method

3.1. Preliminary of Visual Instruction Tuning

In the deep learning era, **visual tuning** has experienced significant paradigm shifts, as depicted in Fig. 3.

In Fig. 3(a), conventional tuning methodologies encompass *Full fine-tuning*, *Head-oriented*, and *Backbone-oriented* techniques, capitalizing on large-scale pre-trained models. Predominantly, thoroughly fine-tuning these models for specific tasks, conducted end-to-end, is recognized as a highly effective strategy. However, this method requires maintaining separate copies of the backbone parameters for each distinct task, posing challenges in storage and deploy-

ment.

Alternatively, Visual Prompt Tuning (VPT) [3], presents an efficient substitute for full fine-tuning within large-scale vision Transformer models. It achieves this by employing a minimal fraction of trainable parameters in the input space while maintaining a frozen backbone model. The objective function for Visual Prompt Tuning is given by:

$$\min_{\theta_p} \mathcal{L}(f(X, P; \theta_p), Y) \quad (1)$$

where \min_{θ_p} is the minimization over the prompt parameters P , \mathcal{L} is the loss function, f represents the model function with input image X , prompt parameters P , and learnable model parameters θ_p as input, and Y is the target output.

Visual Prompt Tuning focuses on optimizing LLMs using a small set of parameters, whereas Visual Instruction Tuning (VIT) aims to improve the model’s comprehension of instructions, thereby addressing the model’s shortcomings in specific domains. This type of method aims to enhance the model’s proficiency in following instructions, leveraging the capabilities of the latest foundation models, *e.g.*, Llama [17], and BLIP2 [5]. Instructions serve as guiding constraints, shaping the model’s outputs to conform to specific response characteristics and domain-relevant knowledge. This approach enables human monitoring of the model’s behavior, thereby assuring alignment with the desired outcomes. Moreover, Instruction Tuning is computationally efficient, allowing LLMs to swiftly adapt to particular domains without extensive retraining or architectural alterations.

The objective function for Visual Instruction Tuning is given by:

$$\min_{\theta_{\text{tunable}}} \mathcal{L}(g(X, I, C; \theta_{\text{tunable}}), Y) \quad (2)$$

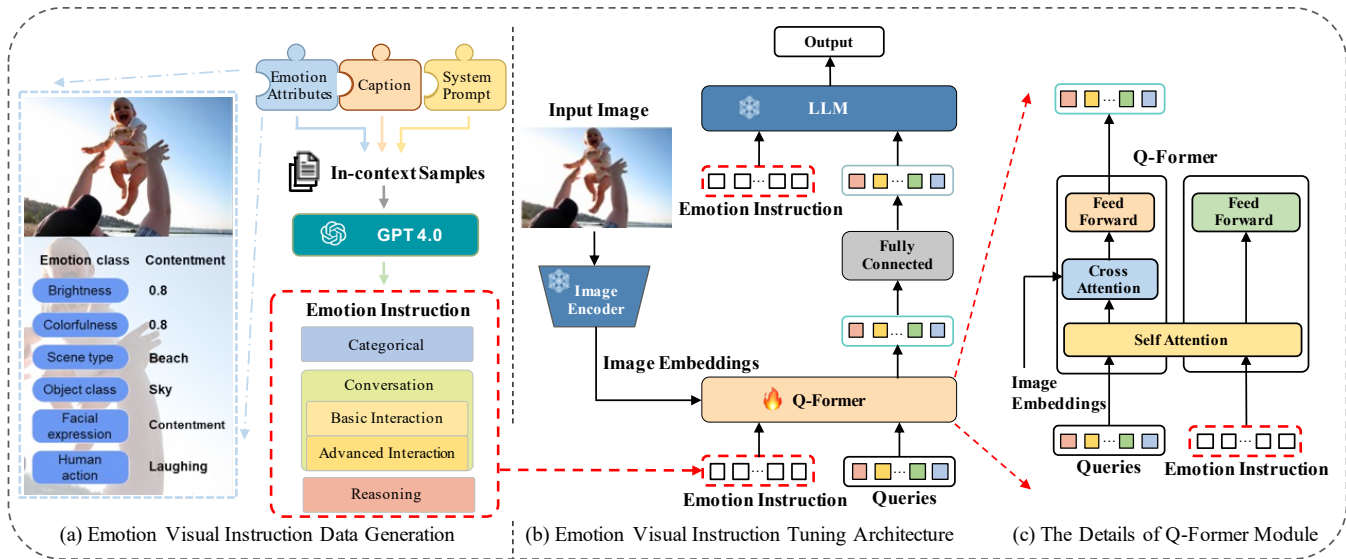


Figure 4. The overall architecture of our proposed method. The *Emotion Instruction* data generated by (a) will be used for *Emotion Visual Instruction Tuning* in (b). During *Emotion Visual Instruction Tuning*, given an input image, the frozen Image Encoder initiates the process by extracting visual features. *Emotion Instruction* generated by (a) are subsequently interacting with *Queries* embedding through the learnable Q-Former. This interaction is key to drawing out image features that are relevant to the task at hand. As a result, the frozen LLM receives visual information conducive to instruction following.

where $\min_{\theta_{\text{tunable}}}$ denotes the minimization over the tunable parameters θ_{tunable} in the Instruction Tuning Module, \mathcal{L} is the loss function, g is the model function with instruction I , image X , other contexts C , and tunable parameters θ_{tunable} , and Y denotes the target output. The optional context C is not just raw data; it encompasses descriptive or directive information guiding the model on how to process input or which task to execute, e.g., image caption. It’s integral to the model’s understanding and execution of tasks based on specific instructions or guidelines.

3.2. GPT-assisted Emotion Visual Instruction Data Generation

Previous methodologies commonly employed a consistent template-based set of instructions for every image within a dataset across various specific tasks [2]. For instance, a standard instruction such as “*Briefly describe the content of the image*” was employed uniformly across all images for Image Captioning. In this way, the model may not be able to adequately capture the unique characteristics of each image. Moreover, this *one-size-fits-all* approach often leads to suboptimal performance in emotion recognition tasks that require nuanced perception and differentiation of ambiguous emotion classes.

Since the topic of Emotion Visual Instruction Tuning is still in its infancy, no benchmarks or guidelines have been proposed so far for constructing emotion instruction data. Based on the recent successes of machine-generated instructions demonstrated in LLaVA [8], our work pioneers the use of existing LLMs to create a pipeline for

self-generating emotion instructions. Different from previous template-based and *one-size-fits-all* instruction data, we propose an instance-wise and LLM-assisted visual emotion instruction data pipeline. This methodology transcends the constraints of manual annotation by employing GPT-4 [11] to generate instance-wise, tailored instruction data that dynamically corresponds to visual content.

Prior to the development of instructional data for the visual emotion recognition task, it is imperative to confront a fundamental academic problem: *What types of visual clues are pivotal in identifying emotions?* This necessitates a careful consideration of the unique characteristics inherent to the task, along with a comprehensive understanding of the potential visual cues associated with human emotions. In this work, we propose a novel visual instruction data mechanism to remove the inherent subjectivity and ambiguity in emotional interpretation. Specifically, we integrate a broad spectrum of emotion attributes across multiple levels: *low-level* attributes (e.g., brightness, colorfulness), *mid-level* attributes (e.g., scene type and object class), and *high-level* attributes (e.g., facial expressions and human actions), building upon insights from previous work [24]. This comprehensive strategy not only aligns with the intricate nature of emotions but also significantly enhances the model’s capability to interpret and understand visual emotional cues more accurately and holistically.

The overall pipeline of our proposed emotion visual instruction data is shown in Fig. 4 (a). For an image X_{img} , three types of image-related contexts are essential for GPT-4 to generate emotion instruction data: (i) a caption X_c , (ii)

an emotion attribute list X_{attr} , which includes *emotion class*, *brightness*, *colorfulness*, *scene type*, *object class*, *facial expression*, and *human action*, and (iii) the system prompt, designed to enable GPT-4 to comprehend the specific task requirement¹.

We first manually design a few examples which are used as seed examples for in-context learning to query GPT-4. This operation leverages the model’s ability to extrapolate from given examples, enhancing its understanding and response accuracy based on the principles of few-shot learning [8]. Our generated emotion instruction data includes three types: *Categorical*, *Conversation*, and *Reasoning*. Building upon previous research [8], our generated instruction data adheres to the dialogue format, exemplified in Fig. 5.

Our strategy for generating emotion instruction data adopts a progressive approach *from simple to complex*. Initially, for the *Categorical* data, we transform the associated emotion class of the image into a structured format. This process serves as the foundational component of our emotion instruction data.

For the *Conversation* data, our framework is designed to create dialogues in which the GPT assistant interacts with an inquirer, focusing on the emotion attributes of the image. In this setup, the assistant’s responses are tailored to interpret and describe the image as though it were within its own visual field, thereby providing insights from an observational viewpoint. The scope of questions posed is comprehensive, encompassing the types of objects depicted, their actions, and the dynamics of their interrelationships. The dialogues we generate fall into two categories: (i) *Basic Interaction*, focusing on the provided emotion attribute list with simple, direct characteristics, and (ii) *Advanced Interaction*, which builds on the first type to reach greater conversational complexity and sophistication.

For the *Reasoning* data, our approach extends beyond mere visual content, prompting the model to generate in-depth reasoning questions. To enhance the dialogue’s credibility and structure, detailed examples are incorporated alongside logical reasoning steps, ensuring that the discourse convincingly captures the intricacies of the visual content.

3.3. Emotion Visual Instruction Tuning

After acquiring the emotion visual instruction data as detailed in Sec. 3.2, our goal is to employ this data in enhancing the existing Visual Instruction Tuning model. This enhancement aims to align the LLMs’ existing knowledge with the emotion understanding domain.

As shown in Fig. 4 b, we have developed an Emotion Visual Instruction Tuning (EmoVIT) architecture based on

¹A detailed description of the system prompt is provided in the supplementary materials.

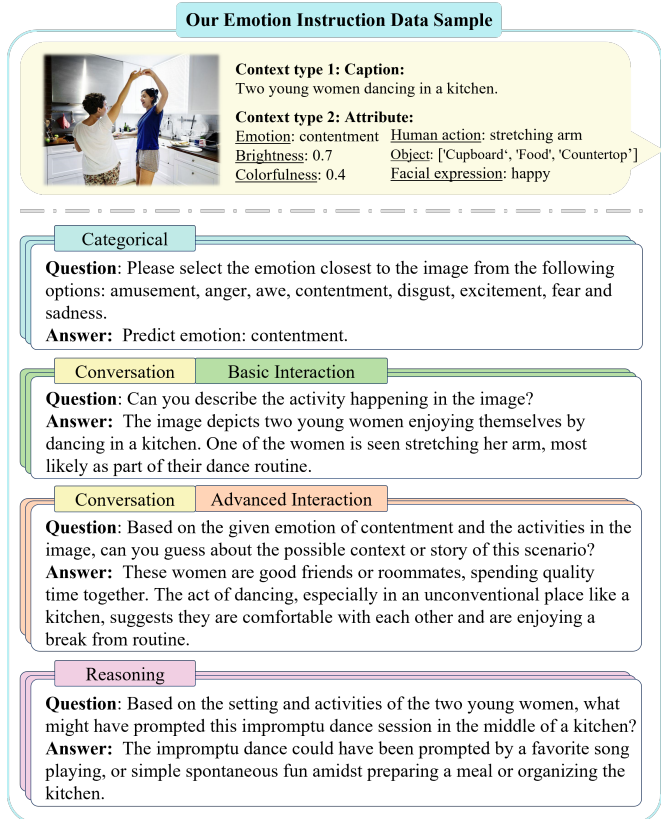


Figure 5. The sample of our generated visual emotion instruction data.

InstructBLIP [2]. This architecture specifically leverages its Instruction-aware Q-Former Module, as depicted in Fig. 4 c, for emotion-centric instructional tasks.

Specifically, the Instruction-aware Q-Former Module takes in the *emotion instruction tokens*, *queries*, and *image embeddings* as input. The image embeddings are extracted by a frozen image encoder. The learnable queries are initially produced by the pre-trained Q-Former of InstructBLIP. During training, the Instruction-aware module enhances task-specific feature extraction. It does this by integrating emotion instruction and query embeddings within self-attention layers, aligning visual information with the LLM’s instruction-following requirements. Our approach adopts cross-entropy loss, tailoring it to the intricacies of visual emotion recognition tasks, thus ensuring precise and contextually relevant model training outcomes.

We note that the data generated by our approach is not confined to a single model but can also be applied to other Visual Instruction Tuning models, such as LLaVA [17]. Notably, when LLaVA is fine-tuned with our data, it exhibits a significant enhancement in emotion recognition capabilities, as detailed in Sec. 4.2. In this way, we demonstrate not only the effectiveness but also the transferability of our generated data, showing its broad applicability and impact.

4. Experimental Results

4.1. Implemental Details

Our implementation is based on the LAVIS library [4]. Our EmoViT starts with a pre-trained InstructBLIP baseline and proceeds to fine-tune exclusively the Q-Former module, whilst keeping both the image encoder and the language model frozen. The parameters for our training adhere to the default settings established by InstructBLIP.

Datasets. We evaluate our framework on ten benchmark datasets annotated under different scenarios and class number, namely EmoSet [24], WEBEmo [12], Emotion6 [14], the Flickr and Instagram (FI) [25], Artphoto [9], IAPS [10], Abstract [9], EmotionROI [15], UnbiasedEmo [12], and OxfordTVG-HIC [6].

Held-in Pretraining. Following previous work [2], we divide our dataset into two categories: held-in for pretraining and held-out for evaluation². Considering the EmoSet dataset’s comprehensive inclusion of emotion attributes for each image, it has been chosen as the primary resource for our held-in pretraining phase. Simultaneously, for a broader assessment, we perform held-out evaluations using the test sets from various other datasets.

For the generation of emotion visual instruction data, we initially employ the BLIP2 model for image captioning, followed by leveraging the GPT-4 API to generate emotion instruction data. In total, our collection comprises *Categorical*, *Conversation*, and *Reasoning* instruction data, derived from 51,200 unique images. This represents less than 50% of the entire EmoSet.

4.2. Held-out Evaluation

As shown in Tab. 1, our proposed methodology exhibits a marked superiority in performance relative to the burgeoning Visual Instruction Tuning Methods. Since they have been pre-trained on dozens of large-scale datasets, it is evident that our generated emotion visual instruction data is particularly effective for emotional understanding. Our results signify a paradigm shift, heralding a new era of model training that relies less on explicit supervision and more on the robustness of emotion instruction-driven learning.

The Effectiveness of Our Proposed Emotion Visual Instruction Data. As the first to introduce the concept of emotion visual instruction data, our study seeks to evaluate the generalizability of this newly generated instruction data. Our goal is to test its efficacy not only with InstructBLIP but also across other Visual Instruction Tuning model, to understand its broader applicability. As depicted in Fig. 6, we employ two Visual Instruction Tuning models, LLaVA and InstructBLIP, which were fine-tuned on our specially gen-

²Unlike the setup in InstructBLIP, our dataset exclusively comprises emotion-related content. Consequently, our held-out evaluation does not constitute a strict zero-shot evaluation in the conventional sense.




Figure 6. The improvement of our proposed emotion visual instruction tuning data tuning on LLaVA [8] and InstructBLIP [2].

erated emotion visual instruction data. Subsequent testing across five distinct datasets reveals notable improvements in both models, substantiating the efficacy of our generated data. Notably, InstructBLIP demonstrated a more substantial overall enhancement compared to LLaVA. This can be attributed to InstructBLIP’s specialized Instruction-aware Q-Former Module, which adeptly extracts the salient features of our emotion instructions and synergizes them effectively with the corresponding images, thereby yielding improved performance.

4.3. Effectiveness of Different Instruction Data

4.3.1 Ablation Study of Different Instruction Data

The ablation study outlined in Tab. 2 provides a comprehensive analysis of the impact that different instructional data types have on model performance, specifically concerning accuracy metrics on the EmoSet test set. Initially, the model, referred to as InstructBLIP [2], operates without the integration of the three types of instructional data and attains a baseline accuracy of 42.20%. This foundational performance is significantly enhanced with the inclusion of *Categorical* data, which alone contributes to a substantial increase in accuracy. The introduction of *Conversation* data further amplifies this effect, underscoring the value of conversational context in improving the model’s predictive capabilities. The addition of *Reasoning* data notably boosts performance, achieving a peak accuracy of 83.36%. This indicates that the model significantly benefits from the nuanced cues in reasoning, aiding in understanding complex emotional instructions. The gradual improvements with each data type support the idea that a diverse approach to instructional data markedly enhances model comprehension and performance.

Method	WebEmo	FI	Emotion6	Abstract	ArtPhoto	IAPSa	EmotionROI	EmoSet
Number of Classes	25	8	6	8	8	8	6	8
Flamingo [1]	9.36	14.91	21.67	3.57	17.5	10.13	21.72	29.59
LLaVA [8]	12.55	56.04	49.44	19.54	36.25	42.43	46.46	44.03
BLIP2 [5]	20.10	57.72	50.00	28.57	36.25	39.24	50.51	46.79
InstructBLIP [2]	12.80	37.97	46.11	21.42	26.25	34.18	46.13	42.20
Ours*	21.12	68.09	57.81	32.34	44.90	44.13	53.87	83.36

Table 1. Held-out performance comparison on visual emotion datasets (%).

Categorical	Conversation	Reasoning	Accuracy (%)
-	-	-	42.20
✓	-	-	80.90 (+38.70)
✓	✓	-	81.95 (+39.75)
✓	✓	✓	83.36 (+41.16)

Table 2. Ablation study of three types of instruction data. Accuracy (%) on EmoSet test set.

4.3.2 Instruction Sensitivity

This work is dedicated to the creation of a varied corpus of visual emotion instruction data, alongside the development of a robust instruction-based model. Our objective is for the model to demonstrate stability, producing consistent results in the face of minor variations in instruction phrasing, provided the core objective of the task persists unchanged. To this end, we employ the *Sensitivity* evaluation metric, as introduced by [22], to assess the model’s fidelity in generating uniform outcomes irrespective of instructional nuances.

We employ two semantically similar instructions as input prompts for the model, testing their impact on the Sensitivity score across three visual emotion datasets for different Visual Instruction Tuning models. The first instruction is: “From the given options: `cls_1`, `cls_2`, `cls_3`, etc., identify the emotion that most accurately reflects the image. Ensure your selection is confined to the listed options. Respond in the format: Predicted emotion:” The second one states: “Please choose the emotion that best corresponds to the image from the following options: `cls_1`, `cls_2`, `cls_3`, etc. (Do not provide answers beyond the provided candidates.) Please reply in the following format: Predict emotion:”

As illustrated in Fig. 7, our approach, along with BLIP2, exhibited exceptionally low Sensitivity values, demonstrating robustness in understanding the instructions. Conversely, Flamingo and InstructBLIP displayed a higher degree of sensitivity, indicating a relative susceptibility to variations in instruction wording.

4.4. Robustness

Given that current emotion recognition datasets often exhibit category imbalances and labeling biases, our aim is




Figure 7. The sensitivity score comparison (the lower the better).

to evaluate the generalization ability of various learning strategies more impartially. Hence, we selected the UnBiasedEmo test set [12], which is uniquely suited for recognizing intricate emotions, such as those associated with identical objects or scenes, *e.g.*, landscapes, crowds, families, babies, and animals, where the emotional undertones can be particularly subtle and complex.

As depicted in Tab. 3, our proposed methodology demonstrates superior performance when benchmarked against conventional supervised emotion recognition techniques, thereby underscoring the efficacy of our approach in more accurately discerning complex emotional contexts.

Method	Accuracy (%)
Direct Learning [12]	71.64
Self-Directed Learning [12]	72.45
Joint Learning [12]	71.64
Curriculum Learning [12]	74.27
Ours*	74.72

Table 3. Performance comparison on UnbiasedEmo dataset.

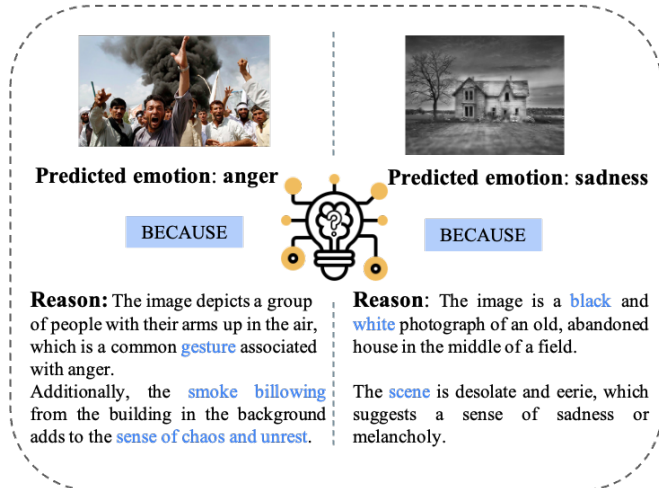


Figure 8. The sample of our generated explanation.

4.4.1 Affective Reasoning

In the domain of visual emotion recognition, where ambiguity and subjectivity are pervasive, the advent of an interpretable model is of considerable value. Such a model elucidates its cognitive processes, enhancing its trustworthiness and practicality in scenarios requiring a delicate grasp of emotional subtleties.

Leveraging Visual Instruction Tuning, our model transcends mere categorization of emotions; it articulates the underlying rationale for its classifications. The executing commands for identifying emotions and elucidating the decision basis is illustrated below:

```
Predicted emotion: [emotion].
Reason: [explanation].
```

Our model delineates the visual features influencing its determinations, thereby addressing the complexities inherent in discerning and explaining emotion-related nuances.

The explanations provide us with visual clues contained within the images, as exemplified in Fig. 8. It provides interpretable visual indicators that inform the model’s outputs, as demonstrated in our example, by disambiguating the often abstract emotional categories.

4.5. Scaling Law

Pretraining data. As demonstrated in Tab. 4, there is a clear correlation between the size of the pre-training dataset and improved performance. Consequently, we anticipate that an increase in training data in the future could enhance the effectiveness of Emotion Visual Instruction Tuning.

4.6. Humour Caption Generation

The comprehension of humor is intricately linked to the understanding of emotions. Leveraging our generative language model, we conduct a caption generation task without

5%	10%	30%	50%
79.00	81.00	79.34	83.36

Table 4. Ablation study of different portion of pre-training data. Accuracy (%) on EmoSet test set.

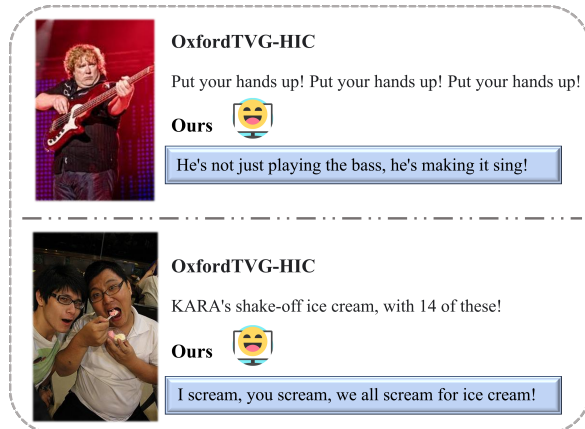


Figure 9. The sample of our generated humour caption vs human writing humour caption from OxfordTVG-HIC.

modifying the model’s architecture, specifically testing the model’s proficiency in generating humorous captions. For this purpose, we select 50 images from the OxfordTVG-HIC dataset [6] and generate corresponding captions using our model. Subsequently, the captions produced by our model are compared with manually annotated captions from the dataset in a user study. Thirty participants were asked to vote on which captions were more humorous. Our model-generated captions receive 60% of the votes, demonstrating its effective humor generation capabilities. One sample is visualized in Fig. 9.

5. Conclusion

In our study, drawing upon the distinctive visual cues key to visual emotion recognition, we present a GPT-assisted pipeline specifically designed for generating emotion visual instruction data. The developed EmoVIT model incorporates emotion-specific instructions, leveraging LLMs for enhanced performance. Our comprehensive experiments validate its effectiveness in emotion classification, affective reasoning, and humor understanding. This comparative analysis sets a benchmark for Emotion Visual Instruction Tuning with LLMs, providing valuable insights and directions for future research in this field.

Acknowledgement

This work is partially supported by the National Science and Technology Council, Taiwan under Grant, NSTC-112-2628-E-002-033-MY4, NSTC-112-2634-F-002-002-MBK, NSTC-112-26234-F-002-005, NSTC-112-2221-E-A49-059-MY3, and NSTC-112-2221-E-A49-094-MY3.

References

- [1] Jean-Baptiste Alayrac, Jeff Donahue, Pauline Luc, Antoine Miech, Iain Barr, Yana Hasson, Karel Lenc, Arthur Mensch, Katherine Millican, Malcolm Reynolds, et al. Flamingo: a visual language model for few-shot learning. In *Advances in Neural Information Processing Systems*, pages 23716–23736, 2022.
- [2] Wenliang Dai, Junnan Li, Dongxu Li, Anthony Meng Huat Tiong, Junqi Zhao, Weisheng Wang, Boyang Li, Pascale Fung, and Steven Hoi. Instructblip: Towards general-purpose vision-language models with instruction tuning. In *Advances in Neural Information Processing Systems*, 2023.
- [3] Menglin Jia, Luming Tang, Bor-Chun Chen, Claire Cardie, Serge Belongie, Bharath Hariharan, and Ser-Nam Lim. Visual prompt tuning. In *Proceedings of the European Conference on Computer Vision*, pages 709–727. Springer, 2022.
- [4] Dongxu Li, Junnan Li, Hung Le, Guangsen Wang, Silvio Savarese, and Steven CH Hoi. Lavis: A library for language-vision intelligence. *arXiv preprint arXiv:2209.09019*, 2022.
- [5] Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. In *Proceedings of the International Conference on Machine Learning*, 2023.
- [6] Runjia Li, Shuyang Sun, Mohamed Elhoseiny, and Philip Torr. Oxfordtv-g-1c: Can machine make humorous captions from images? In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 20293–20303, 2023.
- [7] Zheng Lian, Licai Sun, Mingyu Xu, Haiyang Sun, Ke Xu, Zhuofan Wen, Shun Chen, Bin Liu, and Jianhua Tao. Explainable multimodal emotion reasoning. *arXiv preprint arXiv:2306.15401*, 2023.
- [8] Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. Visual instruction tuning. In *Advances in Neural Information Processing Systems*, 2023.
- [9] Jana Machajdik and Allan Hanbury. Affective image classification using features inspired by psychology and art theory. In *Proceedings of the 18th ACM international conference on Multimedia*, pages 83–92, 2010.
- [10] Joseph A Mikels, Barbara L Fredrickson, Gregory R Larkin, Casey M Lindberg, Sam J Maglio, and Patricia A Reuter-Lorenz. Emotional category data on images from the international affective picture system. *Behavior research methods*, 37:626–630, 2005.
- [11] OpenAI. Gpt-4 technical report. Technical report, 2023.
- [12] Rameswar Panda, Jianming Zhang, Haoxiang Li, Joon-Young Lee, Xin Lu, and Amit K Roy-Chowdhury. Contemplating visual emotions: Understanding and overcoming dataset bias. In *Proceedings of the European Conference on Computer Vision*, pages 579–595, 2018.
- [13] EU Parliament. Eu ai act: first regulation on artificial intelligence. Accessed June, 25:2023, 2023.
- [14] Kuan-Chuan Peng, Tshuan Chen, Amir Sadovnik, and Andrew C Gallagher. A mixed bag of emotions: Model, predict, and transfer emotion distributions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 860–868, 2015.
- [15] Kuan-Chuan Peng, Amir Sadovnik, Andrew Gallagher, and Tshuan Chen. Where do emotions come from? predicting the emotion stimuli map. In *2016 IEEE international conference on image processing (ICIP)*, pages 614–618. IEEE, 2016.
- [16] Dongyu She, Jufeng Yang, Ming-Ming Cheng, Yu-Kun Lai, Paul L Rosin, and Liang Wang. Wscnet: Weakly supervised coupled networks for visual sentiment classification and detection. *IEEE Transactions on Multimedia*, 22(5): 1358–1371, 2019.
- [17] Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*, 2023.
- [18] Jason Wei, Maarten Bosma, Vincent Zhao, Kelvin Guu, Adams Wei Yu, Brian Lester, Nan Du, Andrew M Dai, and Quoc V Le. Finetuned language models are zero-shot learners. In *International Conference on Learning Representations*.
- [19] Hongxia Xie, Hua Chung, Hong-Han Shuai, and Wen-Huang Cheng. Learning to prompt for vision-language emotion recognition. In *2023 11th International Conference on Affective Computing and Intelligent Interaction Workshops and Demos (ACIIW)*, pages 1–4, 2023.
- [20] Hongxia Xie, Ming-Xian Lee, Tzu-Jui Chen, Hung-Jen Chen, Hou-I Liu, Hong-Han Shuai, and Wen-Huang Cheng. Most important person-guided dual-branch cross-patch attention for group affect recognition. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 20598–20608, 2023.
- [21] Liwen Xu, Zhengtao Wang, Bin Wu, and Simon Lui. Mdan: Multi-level dependent attention network for visual emotion analysis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9479–9488, 2022.
- [22] Zhiyang Xu, Ying Shen, and Lifu Huang. Multiinstruct: Improving multi-modal zero-shot learning via instruction tuning. *arXiv preprint arXiv:2212.10773*, 2022.
- [23] Jingyuan Yang, Jie Li, Xiumei Wang, Yuxuan Ding, and Xinbo Gao. Stimuli-aware visual emotion analysis. *IEEE Transactions on Image Processing*, 30:7432–7445, 2021.
- [24] Jingyuan Yang, Qirui Huang, Tingting Ding, Dani Lischinski, Danny Cohen-Or, and Hui Huang. EMOSet: A large-scale visual emotion dataset with rich attributes. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 20383–20394, 2023.
- [25] Quanzeng You, Jiebo Luo, Hailin Jin, and Jianchao Yang. Building a large scale dataset for image emotion recognition: The fine print and the benchmark. In *Proceedings of the AAAI conference on artificial intelligence*, 2016.
- [26] Shengyu Zhang, Linfeng Dong, Xiaoya Li, Sen Zhang, Xiaofei Sun, Shuhe Wang, Jiwei Li, Runyi Hu, Tianwei Zhang, Fei Wu, et al. Instruction tuning for large language models: A survey. *arXiv preprint arXiv:2308.10792*, 2023.
- [27] Wei Zhang, Xuanyu He, and Weizhi Lu. Exploring discriminative representations for image emotion recognition with

cnn. *IEEE Transactions on Multimedia*, 22(2):515–523, 2019.

- [28] Sicheng Zhao, Zizhou Jia, Hui Chen, Leida Li, Guiguang Ding, and Kurt Keutzer. Pdanet: Polarity-consistent deep attention network for fine-grained visual emotion regression. In *Proceedings of the 27th ACM International Conference on Multimedia*, pages 192–201, 2019.