

# IF-Bench: Benchmarking and Enhancing MLLMs for Infrared Images with Generative Visual Prompting

Tao Zhang<sup>1,2\*</sup> Yuyang Hong<sup>1,2\*</sup> Yang Xia<sup>1,2</sup> Kun Ding<sup>1†</sup> Zeyu Zhang<sup>3</sup>  
Ying Wang<sup>1,3</sup> Shiming Xiang<sup>1,2</sup> Chunhong Pan<sup>1,3</sup>

<sup>1</sup>MAIS, Institute of Automation      <sup>2</sup>School of Artificial Intelligence, UCAS  
<sup>3</sup>Research Center of Aerospace Information, Institute of Automation

## Abstract

Recent advances in multimodal large language models (MLLMs) have led to impressive progress across various benchmarks. However, their capability in understanding infrared images remains unexplored. To address this gap, we introduce **IF-Bench**, the first high-quality benchmark designed for evaluating multimodal understanding of infrared images. IF-Bench consists of 499 images sourced from 23 infrared datasets and 680 carefully curated visual question-answer pairs, covering 10 essential dimensions of image understanding. Based on this benchmark, we systematically evaluate over 40 open-source and closed-source MLLMs, employing cyclic evaluation, bilingual assessment, and hybrid judgment strategies to enhance the reliability of the results. Our analysis reveals how model scale, architecture, and inference paradigms affect infrared image comprehension, providing valuable insights for this area. Furthermore, we propose a training-free generative visual prompting (**GenViP**) method, which leverages advanced image editing models to translate infrared images into semantically and spatially aligned RGB counterparts, thereby mitigating domain distribution shifts. Extensive experiments demonstrate that our method consistently yields significant performance improvements across a wide range of MLLMs. The benchmark and code are available at <https://github.com/casiatao/IF-Bench>.

## 1. Introduction

Recently, multimodal large language models (MLLMs) such as GPT-4o [18], Gemini-2.5-Flash [10], and Qwen3-VL [29] have achieved remarkable progress in image analysis and understanding, consistently setting new records on various benchmarks [16, 33]. However, since these models

\*Equal contribution. †Corresponding author.

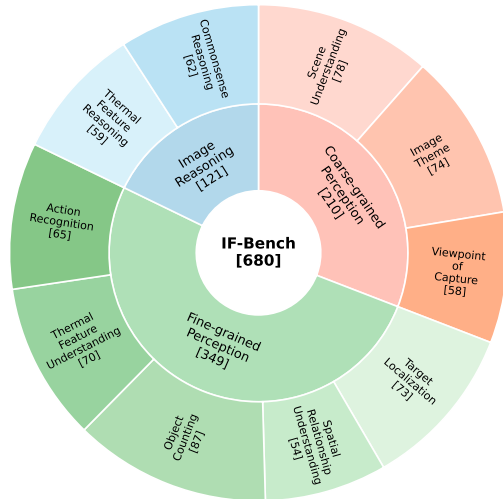


Figure 1. Distribution of questions across dimensions in IF-Bench.

are primarily trained on natural images, existing evaluations mainly focus on natural scenes, leaving their understanding ability on out-of-domain data, such as infrared images, largely unexplored. Infrared imaging offers superior visibility under low illumination and adverse weather conditions, making it widely used in applications such as surveillance [12, 24] and aerial monitoring [23]. This naturally raises a critical question: *How well can current MLLMs understand infrared images?* Previous studies [9, 13] have undertaken preliminary explorations in this direction, but the narrow task coverage, absence of human calibration, and limited model choices leave the actual infrared understanding capability of mainstream MLLMs still unclear.

To address this gap, we propose **IF-Bench**, a high-quality and systematic benchmark for evaluating multimodal understanding of infrared images. As illustrated in Fig. 1, we first decompose infrared image understanding into three major tasks: coarse-grained perception, fine-grained perception, and image reasoning, which are further




Figure 2. The performance of GenViP on IF-Bench.

divided into **10** dimensions to comprehensively cover diverse application scenarios. Subsequently, 1,166 infrared images were selected from 23 datasets, and 4,628 visual question-answer (VQA) pairs were constructed through a hybrid process of manual annotation and automatic generation. After a rigorous two-stage human filtering and calibration process, we obtain the final benchmark consisting of **499** infrared images and **680** high-quality VQA pairs.

Based on IF-Bench, we comprehensively evaluate over **40** open-source and closed-source MLLMs with different architectures, scales, and reasoning paradigms, including Qwen2.5-VL [5], Qwen3-VL [29], InternVL3.5 [27], Gemini-2.5-Pro [10], and Doubao-Seed-Vision [7]. To ensure fair and robust evaluation, we design multiple strategies, including the unified evaluation prompt, circular evaluations, bilingual (English and Chinese) assessments, and hybrid correctness judgments. Experimental results exhibit several key findings: (1) increasing model scale consistently improves infrared understanding performance; (2) Mixture-of-Experts (MoE) architectures achieve a better trade-off between accuracy and inference efficiency; (3) The thinking mode enhances performance on *thermal feature understanding and reasoning* dimensions, but reduces accuracy in fine-grained perception tasks; and (4) Open-source models exhibit comparable performance to closed-source models. Overall, existing MLLMs still struggle with comprehending fine details in infrared images. The primary reasons lie partly in the inherent representation limitations of the models themselves and partly in the distribution shift of the input domain: since most MLLMs are primarily trained on RGB images, feeding infrared inputs often introduces substantial domain mismatch, which may lead to performance degradation. Direct adaptation through supervised fine-tuning is challenging due to the scarcity of high-quality infrared-text data and the need for model-specific adjustments.

To overcome these challenges, we propose a training-free **Generative Visual Prompting (GenViP)** method. GenViP leverages an image editing model to convert infrared images into semantically and spatially aligned RGB counterparts, which are then jointly fed with the original infrared

inputs. This design preserves infrared-specific information while effectively reducing domain gaps. The method requires no paired infrared-text data or model fine-tuning, making it applicable to arbitrary MLLMs. Experimental results indicate that using existing editing models can effectively enable GenViP to enhance infrared image understanding. However, open-source models such as Qwen-Edit-2509 [28] still underperform closed-source ones like Seedream 4.0 [21] and Gemini-2.5-Flash [10]. To further boost GenViP’s performance and practical utility with open-source editing models, we fine-tune Qwen-Edit-2509 [28] on 50,000 RGB-T image pairs curated from over 300,000 candidates, allowing it to surpass the aforementioned state-of-the-art (SOTA) closed-source editing models. Extensive results demonstrate that GenViP consistently improves infrared image understanding performance across diverse models, achieving up to **7%** relative performance gain on IF-Bench and even outperforming closed-source understanding models like Doubao-Seed-Vision-1.6-250815 [7] and Gemini-2.5-Pro [10], as shown in Fig. 2.

The main contributions of this work are as follows:

- A high-quality benchmark, IF-Bench, for infrared image understanding is constructed, which fills a critical gap in this field and will be publicly released.
- A comprehensive and reliable evaluation covering over 40 models, along with in-depth analyses, is presented, providing a solid foundation for future research.
- A training-free Generative Visual Prompting (GenViP) method is proposed to effectively enhance the infrared understanding capability of various multimodal models.

## 2. Related Work

**MLLM.** The rapid growth of large language models (LLMs) has driven the development of MLLMs. Early works like Flamingo [2] and LLaVA [15] integrate pre-trained visual encoders with frozen language models for cross-modal alignment. Recent advances further improve visual reasoning and complex semantic understanding. Qwen2.5-VL [5] introduces dynamic resolution processing for multi-scale adaptation. LLaVA-OneVision-1.5 [3] provides an open framework and a large-scale pre-training dataset, achieving competitive performance across diverse tasks. InternVL3.5 [27] leverages cascade reinforcement learning to strengthen visual comprehension and reasoning. Qwen3-VL [29] employs the thinking mode and MoE design to enhance understanding and reasoning in complex tasks. In parallel, closed-source models such as GPT [18], Gemini [10], and Doubao [6] also advance rapidly, demonstrating strong capabilities in diverse multimodal scenarios. Nevertheless, existing MLLMs still remain limited in infrared image understanding. Previous works like Infrared-LLaVA [13] and IRGPT [9] attempt modality-specific adaptation via fine-tuning, but face challenges including scarce

Table 1. **Dimensions and corresponding examples of IF-Bench.** Examples with images are listed in Appendix A.

Task Category	Dimension	Examples
Coarse-grained Perception	Scene Understanding	<i>What type of environment does this image depict?</i> (A) Rural road. (B) Urban highway. (C) Suburban neighborhood. (D) Industrial area.
	Image Theme	<i>What is the most likely theme of this infrared image?</i> (A) Traffic surveillance. (B) Wildlife monitoring. (C) Security surveillance. (D) Agricultural monitoring.
	Viewpoint of Capture	<i>Compared to people in the elevator, from which viewpoint was this image captured?</i> (A) Top-down. (B) Frontal. (C) Side-view. (D) Bottom-up.
Fine-grained Perception	Target Localization	<i>What are the coordinates of the person closest to the straight pole in the image?</i> (Format: (Target Center X, Target Center Y, Target Width, Target Height)) (A) (0.4, 0.8, 0.1, 0.3). (B) (0.54, 0.9, 0.06, 0.2). (C) (0.6, 0.7, 0.05, 0.15). (D) (0.5, 0.95, 0.08, 0.25).
	Spatial Relationship Understanding	<i>What is the spatial relationship between the two cars in the image?</i> (A) The car on the right is behind the car on the left. (B) The cars are side by side. (C) The car on the right is in front of the car on the left. (D) The cars are not visible.
	Object Counting	<i>How many people are visible in the image?</i> (A) Between 20 and 30. (B) Between 10 and 20. (C) Less than 10. (D) More than 30.
	Thermal Feature Understanding	<i>Which area in the image appears to have the highest thermal activity?</i> (A) The crowd of people. (B) The trees on the right. (C) The ground in the foreground. (D) The sky above.
	Action Recognition	<i>What action is the animal in the image performing?</i> (A) Running. (B) Walking. (C) Jumping. (D) Standing still.
Image Reasoning	Thermal Feature Reasoning	<i>What could be the reason for the uniform heat signatures among the individuals?</i> (A) They are all performing different actions. (B) They are standing in a cold environment. (C) The ambient temperature is very high. (D) They are wearing identical clothing materials.
	Commonsense Reasoning	<i>Based on the image, what is the most likely purpose of the barrier in the foreground?</i> (A) To guide traffic flow. (B) To block unauthorized access. (C) To provide shade. (D) To serve as a decorative element.

image-text data, complex adaptation pipelines, and difficulties in keeping pace with rapidly evolving MLLMs. In contrast, our GenViP requires no training or image-text data and can be directly applied to any MLLM, offering superior convenience and generalization.

**Multimodal Understanding Benchmark.** Systematic benchmarks are essential for quantifying and comparing model performance. MM-Vet [32] and MM-Bench [16] cover a wide range of tasks, including visual understanding and spatial reasoning. MMMU [33] focuses on domain-specific knowledge and high-level reasoning, challenging models to perform tasks akin to those faced by human experts. MVBench [14] transforms static tasks into dynamic ones to evaluate models’ temporal perception and cognitive abilities. MMSI-Bench [31] examines spatial reasoning, scene reconstruction, and spatial transformation capabilities using over 120,000 multi-view images. CAPTURE [20] assesses the ability to handle occlusion by requiring reasoning about the spatial relationships of partially hidden objects. For infrared image understanding, some works [9, 13] construct evaluation sets automatically from existing annotated datasets. However, their benchmarks suffer from limited task coverage, a lack of manual verification, and a narrow set of evaluated models, making them insufficient for assessing the infrared understanding abilities of current MLLMs. On the contrary, our proposed IF-Bench is a high-quality benchmark designed with diverse tasks and carefully curated questions. It has been systematically evaluated on more than 40 mainstream MLLMs, providing a robust foundation for future research.

### 3. IF-Bench

In this section, we first define the tasks involved in infrared image understanding, then present a detailed description of the construction pipeline of IF-Bench, and finally outline its evaluation protocol, as depicted in Fig. 3.

#### 3.1. Task Definition

To construct a comprehensive evaluation benchmark for infrared image understanding, we first identify three fundamental tasks: coarse-grained perception, fine-grained perception, and image reasoning. Based on this categorization, we further decompose them into ten dimensions. Detailed definitions of these dimensions are provided below, with corresponding examples shown in Tab. 1.

**Coarse-grained Perception.** (1) *Scene Understanding*: Identify the general environment or scene depicted in the image, such as indoor, outdoor, forest, or highway. (2) *Image Theme*: Determine the primary content or application context of the image, such as aided driving, wildlife monitoring, or security surveillance. (3) *Viewpoint of Capture*: Identify the viewpoint from which the image was captured, such as top-down, frontal, or side-view.

**Fine-grained Perception.** (1) *Target Localization*: Identify the spatial location of a target within the image, including precise localization using a bounding box and rough location, such as upper-left or center. (2) *Spatial Relationship Understanding*: Recognize the spatial relationship between multiple objects in the image, such as front/back or left/right. (3) *Object Counting*: Count the number of objects

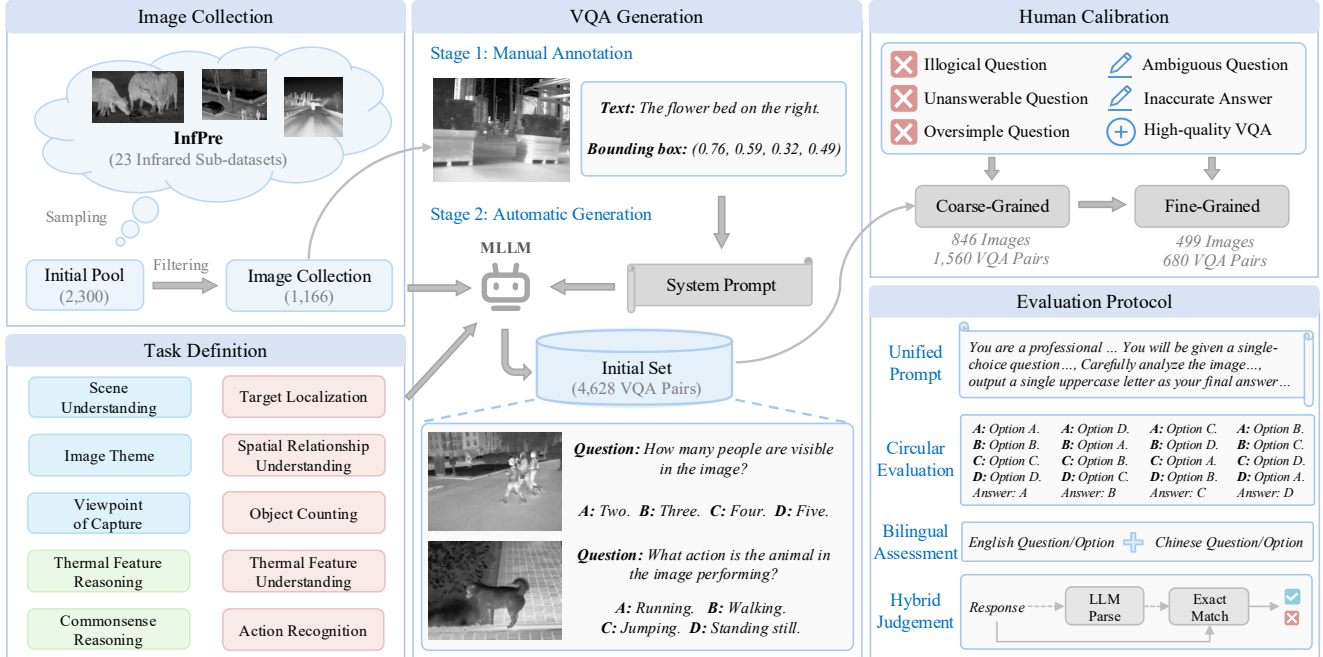


Figure 3. Construction pipeline and evaluation protocol of IF-Bench.

present in the image. (4) *Thermal Feature Understanding*: Assess thermal variations across image regions, such as detecting heat sources or thermal leaks. (5) *Action Recognition*: Determine whether a target is active and identify its action based on infrared contours.

**Image Reasoning.** (1) *Thermal Feature Reasoning*: Infer potential causes of temperature changes using observed thermal patterns. (2) *Commonsense Reasoning*: Deduce the function or intended use of objects based on image content and infer actions that the objects are likely to perform.

### 3.2. Construction Pipeline

**Image Collection.** We collect infrared images from InfPre [34], a large-scale pre-training dataset that integrates 23 infrared datasets. Specifically, we randomly sample 100 images from each dataset, yielding an initial pool of 2,300 images. To ensure high visual fidelity, images with either width or height below 200 pixels are first excluded, followed by a manual quality assessment to further refine the image pool. Through this two-stage filtering process, we obtain a collection of 1,166 high-quality infrared images.

**VQA Generation.** To facilitate the accurate assessment, we adopt the form of visual question answering with well-defined, deterministic answers. We employ a two-stage procedure to construct the initial set of VQA pairs. (1) **Manual Annotation**: To provide ground-truth bounding boxes for the *target localization* dimension, we randomly sample 100 images from the image collection. For each image, objects are selected and annotated with a concise textual description and a precise bounding box, while target-free images are

excluded. This results in a final set of 61 annotated images.

(2) **Automatic Generation**: Building upon the curated image set, we establish an automated question-answer generation pipeline. Given annotated images with their bounding box and textual description, as well as unannotated images, Qwen2.5-VL-72B [5] is prompted to generate at most four single-choice questions per image, each accompanied by four options and the ground-truth answer. The specific system prompt is illustrated in Appendix B. Following this pipeline, we obtain 4,628 VQA pairs.

**Human Calibration.** Considering the potential hallucinations introduced in automatic generation, we apply a coarse-to-fine, two-stage manual calibration process. The calibration follows several key criteria: (1) *Plausibility Assessment*: Remove ill-formed or logically inconsistent questions. (2) *Ambiguity Resolution*: Revise ambiguously phrased questions. For instance, clarifying the observer’s viewpoint in the *spatial relationship understanding* dimension. (3) *Answerability Evaluation*: Exclude questions for whose answers cannot be reliably inferred from the image content. (4) *Answer Verification*: Correct inaccurate or mismatched answers. (5) *Difficulty Adjustment*: Filter out highly repetitive or overly simplistic questions to ensure balanced difficulty. (6) *Data Augmentation*: Introduce additional high-quality VQA pairs grounded in images. The two stages both adhere to the above criteria, while the fine-grained filtering stage is performed by domain experts in infrared imaging, who apply more rigorous quality standards.

Through the above three steps, we obtain the final **IF-Bench**, which comprises **499** infrared images and **680** VQA

pairs. Each question is provided in both Chinese and English. The question distribution is shown in Fig. 1. The final image set still maintains a relatively uniform distribution across the 23 infrared datasets, as analyzed in Appendix A. The order of options is randomly shuffled so that the correct answers are evenly distributed among options A-D.

### 3.3. Evaluation Protocol

To ensure the reliability and robustness of evaluation results, we design several evaluation strategies for IF-Bench, as depicted in Fig. 3. (1) **Unified Prompt:** All models are evaluated under an identical system prompt and instructed to output only the correct answer for each question. The specific prompt is provided in Appendix B. (2) **Circular Evaluation:** To mitigate positional bias, we follow the practice in MM-Bench [16]. Specifically, for each question, the four options and the correct answer are cyclically permuted. Each model is evaluated on all permutations, with the final score averaged across them. (3) **Bilingual Assessment:** Each question is presented in both Chinese and English, with the final score averaged across the two languages. Combined with circular evaluation, each question is evaluated eight times, substantially reducing randomness and improving reliability. (4) **Hybrid Correctness Judgment:** A hybrid strategy combining exact answer matching and LLM-based parsing is employed to balance accuracy and efficiency in correctness assessment. A model’s response is first checked for an exact match with the ground-truth answer. If no match is found, we utilize Qwen3-7B [29] to extract the answer from the model’s response without adding any extraneous information, which is then compared to the ground truth. The extraction prompt is listed in Appendix B. This strategy effectively mitigates the impact of non-standard output formats on evaluation accuracy.

## 4. Infrared Understanding Enhancement

### 4.1. Generative Visual Prompting

RGB images form the main training data for current MLLMs. However, when these models are applied to infrared images, distribution shifts between training and inference inputs can degrade their visual understanding capabilities. Fine-tuning on infrared understanding data is a straightforward solution but faces several challenges: (1) limited high-quality infrared image–text datasets; (2) high computational and engineering costs for individually fine-tuning each model, considering the rapid progress of MLLMs. (3) possible performance drops on general tasks.

To overcome these limitations, we propose a training-free **Generative Visual Prompting (GenViP)** method to improve MLLMs’ infrared image understanding. The key idea is to reduce distribution shifts by directly modifying the model inputs during inference. Specifically, given




Figure 4. The illustration of GenViP.

the advanced capabilities of recent image editing models [10, 21, 28], we translate infrared images into spatially and semantically aligned RGB images. Feeding these translated images into MLLMs effectively aligns the input distribution during inference with that of the training phase. However, since the translated natural images lose thermal information, the models struggle to answer questions about thermal features. Therefore, we adopt a composite-input strategy, feeding both the original infrared image and the translated RGB image into MLLMs, as illustrated in Fig. 4. This preserves thermal information while leveraging the model’s strong understanding of RGB images. Consequently, GenViP eliminates the need for high-quality infrared image–text datasets or fine-tuning for each model, making it directly applicable to any existing MLLM.

To handle the dual-image input, we modify the evaluation prompt accordingly, as shown in Fig. 4 and Appendix B. Additionally, we incorporate a brief textual description about the characteristics of infrared images into the prompt (Appendix B), namely **textual prior**, which effectively enhances the model’s ability to interpret infrared thermal features. Extensive experiments subsequently validate the effectiveness of our method in improving the infrared image understanding capabilities of MLLMs.

### 4.2. Editing Model Optimization

Our preliminary investigations show that the translation quality of editing models strongly affects GenViP’s performance: closed-source editing models such as Seedream 4.0 [21] and Gemini-2.5-Flash [10] outperform open-source models like Qwen-Edit-2509 [28]. Considering the cost of using closed-source editing models, we further optimize GenViP for open-source models to improve its practicality.

Initially, we collect over 370,000 RGB-T image pairs and apply rigorous data filtering and sampling as follows: (1) *Data Source Screening:* Manually remove sources where infrared and RGB images are misaligned. (2) *Resolution filtering:* Discard pairs with a width or height below 200 pixels. (3) *Pairing Quality Filtering:* First, remove pairs with excessively low brightness. Next, compute the Canny edges [8] for both the infrared and RGB images, filtering out pairs whose Dice coefficient falls below a threshold. Finally, using Qwen2.5-VL-32B [5] to as-

Table 2. **Evaluation results of various models on IF-Bench.** Column abbreviations: Avg – Average score; SU – Scene Understanding; IT – Image Theme; VC – Viewpoint of Capture; TL – Target Localization; SRU – Spatial Relationship Understanding; OC – Object Counting; TFU – Thermal Feature Understanding; AR – Action Recognition; TFR – Thermal Feature Reasoning; CR – Commonsense Reasoning. Results of closed-source models are marked in gray. The highest average scores within each group are marked in cyan.

Models	Avg	Coarse-grained Perception			Fine-grained Perception					Image Reasoning	
		SU	IT	VC	TL	SRU	OC	TFU	AR	TFR	CR
InternVL3-1B [35]	43.0	72.4	67.4	33.2	27.9	31.0	27.7	32.9	48.3	30.3	58.9
InternVL3.5-1B [27]	<b>56.6</b>	92.9	84.1	67.2	36.3	37.5	35.2	40.0	64.2	36.9	71.6
InternVL3.5-1B-Thinking [27]	52.5	76.6	71.1	60.6	35.3	35.2	39.8	48.2	53.7	42.2	62.7
InternVL3-2B [35]	65.6	91.7	82.1	61.9	56.3	44.7	42.1	55.7	70.0	68.2	83.1
InternVL3.5-2B [27]	62.3	90.5	86.3	61.9	55.7	44.0	38.5	47.7	67.5	57.6	73.0
InternVL3.5-2B-Thinking [27]	<b>65.5</b>	91.0	89.5	62.5	50.2	46.5	44.1	58.9	68.5	69.9	74.0
Qwen2.5-VL-3B [5]	66.0	93.9	79.7	61.2	76.0	45.4	36.9	48.9	62.5	79.7	75.4
InternVL3.5-4B [27]	67.4	92.3	85.6	69.0	58.0	55.3	41.8	58.4	65.2	71.8	76.8
InternVL3.5-4B-Thinking [27]	71.4	92.9	88.0	76.5	61.5	58.3	43.2	64.8	65.2	86.2	77.2
LLaVA-OneVision-1.5-4B-Instruct [3]	75.2	93.3	90.2	69.4	71.9	55.1	50.6	66.6	79.0	90.7	85.3
Qwen3-VL-4B-Instruct [29]	<b>77.4</b>	95.4	92.2	70.3	69.7	61.6	53.2	68.2	82.1	90.3	91.1
Qwen3-VL-4B-Thinking [29]	76.8	93.9	92.6	83.2	59.8	59.7	51.3	66.4	84.0	87.7	89.7
Qwen2.5-VL-7B [5]	71.1	93.3	86.5	74.4	60.3	49.5	42.7	58.9	74.0	88.6	83.1
InternVL3-8B [35]	71.2	95.0	87.8	71.6	66.4	51.9	45.0	57.7	79.0	70.8	87.1
InternVL3.5-8B [27]	69.1	92.6	79.7	66.8	66.7	52.3	45.0	61.8	67.7	82.8	75.6
InternVL3.5-8B-Thinking [27]	72.3	92.8	85.5	65.3	67.3	58.6	47.3	65.0	71.0	91.5	78.8
Keye-VL-1.5-8B [30]	71.5	85.1	82.6	70.7	77.9	61.1	42.1	60.9	71.9	80.9	82.1
Keye-VL-1.5-8B-Auto-Thinking [30]	71.7	85.4	84.5	72.6	73.1	59.5	43.0	61.8	73.1	81.6	82.7
Keye-VL-1.5-8B-Thinking [30]	73.3	86.1	85.3	77.4	72.1	63.0	45.8	64.8	69.2	89.0	80.6
LLaVA-OneVision-1.5-8B-Instruct [3]	75.9	95.0	89.7	70.9	73.1	54.2	54.6	62.0	79.2	91.9	88.3
Qwen3-VL-8B-Instruct [29]	<b>78.8</b>	93.9	93.1	79.7	74.1	65.5	60.2	68.6	75.8	88.3	88.9
Qwen3-VL-8B-Thinking [29]	76.5	94.6	91.2	74.8	62.0	65.3	58.0	69.3	76.7	89.8	83.1
GLM-4.1V-9B-Thinking [25]	77.2	94.4	92.1	77.8	67.8	59.7	55.0	69.3	77.7	89.4	88.9
InternVL3-14B [35]	73.8	91.3	91.0	73.7	64.9	59.3	48.4	65.9	74.2	82.4	86.3
InternVL3.5-14B [27]	73.6	94.1	92.1	66.4	74.0	56.9	48.6	70.2	67.3	89.9	76.8
InternVL3.5-14B-Thinking [27]	<b>74.6</b>	93.9	93.1	71.3	69.9	58.6	46.0	72.7	69.6	92.6	78.6
InternVL3.5-20B-A4B [27]	69.7	90.5	89.4	68.8	63.9	55.1	42.2	60.9	67.9	76.1	81.9
InternVL3.5-20B-A4B-Thinking [27]	69.8	90.7	91.2	63.4	58.0	56.0	40.7	66.8	64.2	87.1	80.0
InternVL3.5-30B-A3B [27]	74.4	94.9	88.2	73.7	75.0	59.0	45.0	69.0	68.1	85.8	85.7
InternVL3.5-30B-A3B-Thinking [27]	75.4	93.4	90.0	74.1	69.0	62.0	48.4	71.4	68.8	92.4	84.5
Qwen3-VL-30B-A3B-Instruct [29]	<b>82.3</b>	96.2	95.3	83.4	79.6	62.5	61.9	69.1	85.2	98.1	91.9
Qwen3-VL-30B-A3B-Thinking [29]	79.6	95.2	94.4	84.9	68.8	60.6	60.2	67.9	77.5	95.6	90.7
Qwen2.5-VL-32B [5]	75.4	94.9	89.7	75.6	73.5	59.7	45.1	63.0	71.5	96.0	84.5
InternVL3-38B [35]	78.9	93.8	92.4	76.3	76.0	61.1	58.3	69.1	78.5	94.3	89.7
InternVL3.5-38B [27]	79.0	93.4	93.2	77.8	80.5	57.6	59.6	70.0	77.7	91.5	88.7
InternVL3.5-38B-Thinking [27]	<b>80.1</b>	94.1	94.4	78.4	75.2	62.5	65.8	76.3	72.5	92.4	89.7
Qwen2.5-VL-72B [5]	78.1	93.9	93.1	80.6	72.4	64.1	50.0	66.8	76.9	94.9	88.3
InternVL3-78B [35]	80.2	94.2	94.1	78.7	74.1	62.3	61.9	73.6	80.4	95.1	87.9
GLM-4.5V-106B-A12B-Thinking [25]	<b>80.3</b>	95.2	95.3	80.6	82.0	61.8	57.6	69.1	82.1	93.9	85.7
Qwen3-VL-235B-A22B-Instruct [29]	83.7	95.8	94.9	84.5	83.2	65.3	68.4	71.6	84.0	97.5	92.1
Qwen3-VL-235B-A22B-Thinking [29]	82.8	95.5	95.4	81.9	79.6	67.1	64.7	75.4	80.2	94.9	92.9
InternVL3.5-241B-A28B [27]	83.9	95.5	95.3	82.1	88.0	66.7	70.7	73.8	82.3	95.8	89.1
InternVL3.5-241B-A28B-Thinking [27]	83.9	95.4	95.9	81.5	85.3	65.3	69.4	81.1	82.3	93.9	88.9
Doubao-Seed-1.6-251015 [6]	79.9	92.5	92.9	83.4	77.6	67.6	54.3	74.6	78.3	91.1	86.9
Doubao-Seed-Vision-1.6-250815 [7]	<b>84.2</b>	92.3	91.9	87.9	87.7	68.5	63.2	78.6	83.1	96.6	91.9
Gemini-2.5-Flash [10]	79.8	94.2	94.6	75.9	76.0	67.6	62.6	70.7	71.5	93.2	91.1
Gemini-2.5-Pro [10]	82.0	94.9	89.9	85.3	80.8	63.9	63.2	76.4	80.0	94.1	91.9

sess the pairing quality of the remaining data. (4) *Feature Extraction and Deduplication*: Extract visual features using DINOv3 [22], and remove pairs whose infrared features closely match those in IF-Bench to prevent data leakage. (5) *Hierarchical Clustering and Balanced Sampling*: After the above steps, 156,330 high-quality pairs remain. To ensure a balanced distribution across scenes, we perform hierar-

chical clustering [26] based on the RGB DINOv3 features, followed by balanced sampling [26]. Finally, a high-quality dataset of 50,000 RGB-T image pairs is constructed.

Based on this dataset, we fine-tune Qwen-Edit-2509 [28] and obtain **Qwen-Edit-2509-FT**, which achieves superior infrared-to-RGB translation quality and even outperforms closed-source editing models on IF-Bench.

Table 3. **Correlation coefficients between model scale and scores on IF-Bench.** We first compute the correlation coefficients within each model family, and then average the results across families. For MoE models, we use their total parameter count.

	Avg	SU	IT	VC	TL	SRU	OC	TFU	AR	TFR	CR
Corr	0.76	0.42	0.67	0.69	0.58	0.70	0.86	0.77	0.63	0.63	0.62

## 5. Experiments

### 5.1. Evaluation Models.

We systematically evaluate over 40 mainstream MLLMs on IF-Bench, covering diverse architectures, parameter scales, and reasoning paradigms, including InternVL3 [35], InternVL3.5 [27], Qwen2.5-VL [5], Qwen3-VL [29], LLaVA-OneVision [3], Keye-VL-1.5 [30], GLM-4.1V [25], GLM-4.5V [25], Gemini-2.5-Pro [10], and Doubao-Seed-Vision-1.6 [7]. These models range from 1B to 241B parameters, encompassing both dense and (MoE) architectures, with and without explicit reasoning capabilities. Collectively, they provide a representative and comprehensive snapshot of the current multimodal model landscape.

### 5.2. Evaluation Results of IF-Bench.

The complete evaluation results are presented in Tab. 2. We conduct an in-depth analysis in the following aspects.

**Model Scale and Architecture.** (1) **Significant performance gaps are observed between models of different architectures or series**, particularly for smaller models (<30B). For instance, LLaVA-OneVision-1.5-4B-Instruct and Qwen2.5-VL-3B differ by 9.2 in average score, while Qwen3-VL-8B-Instruct and InternVL3.5-8B are nearly 10.0 points apart. As model scale increases, the performance differences between architectures gradually diminish. Qwen3-VL-235B-A22B-Instruct and InternVL3.5-241B-A28B-Thinking differ by only 0.2 points (83.7 vs 84.2). (2) **Smaller models of the Qwen3-VL series demonstrate outstanding performance**, highlighting their practical value for infrared image understanding. For example, Qwen3-VL-8B achieves an overall score of 78.8, surpassing that of Qwen2.5-VL-72B and approaching InternVL3-78B, using only approximately one-tenth of the parameters. (3) **MoE models provide a favorable trade-off between performance and inference efficiency.** InternVL3.5-30B-A3B outperforms the fully dense InternVL3.5-14B (74.4 vs 73.6), while only activating fewer than one-quarter of parameters. (4) To further evaluate the relationship between model size and performance, we compute their Pearson correlation coefficients, as shown in Tab. 3. The results indicate that **model size is positively correlated with performance across all dimensions**, with particularly strong correlations observed on more challenging tasks, such as *object counting* and *thermal feature understanding*.




Figure 5. The average performance change after using thinking.




Figure 6. Box plot of scores across dimensions in IF-Bench.

**Thinking vs Non-Thinking.** To explore the impact of the thinking mode, we analyze 14 paired thinking and non-thinking variants in Tab. 2. Interestingly, the effectiveness of thinking varies considerably across model families. In the Qwen3-VL series, enabling thinking results in a noticeable performance drop, whereas all InternVL3.5 models from 2B to 38B benefit from thinking. As shown in Fig. 5, thinking enhances performance in *thermal feature understanding* (+4.2) and *thermal feature reasoning* (+4.42), but reduces accuracy in dimensions such as *target localization* (-4.90) and *action recognition* (-1.68). Overall, the thinking mode yields only a modest average improvement of 0.34. Considering the substantial increase in reasoning overhead it introduces, **incorporating thinking is not an effective approach to enhancing infrared image understanding.**

**Open-Source vs Closed-Source.** The performance gap between open-source and closed-source models is minimal. Although Doubao-Seed-Vision-1.5-250815 [7] achieves the highest score of 84.2, surpassing all open-source models, it exceeds the best open-source model, InternVL3.5-241B-A28B, by only 0.3 points. Moreover, the other three closed-source models, including Gemini-2.5-Pro [10], perform worse than Qwen3-VL-235B-A22B and InternVL3.5-241B-A28B, underscoring that **open-source models have significant practical value for infrared image understanding.**

**Evaluation Dimension.** The accuracy distribution across different dimensions is illustrated in Fig. 6. Overall, *scene understanding* and *image theme* in coarse-grained perception tasks, as well as the two reasoning dimensions, are relatively easy, with average scores exceeding 80. In contrast, the five fine-grained perception dimensions and the *viewpoint of capture* dimension exhibit notably lower performance. In particular, for *object counting* and *spatial relationship understanding*, the average accuracies drop below

Table 4. Performance of GenViP on IF-Bench across various models.

Model	GenViP	Avg	SU	IT	VC	TL	SRU	OC	TFU	AR	TFR	CR
InternVL3.5-1B [27]	✗	56.6	92.9	84.1	67.2	36.3	37.5	35.2	40.0	64.2	36.9	71.6
	✓	<b>60.8 (+4.2)</b>	94.9	85.8	67.5	34.6	41.7	38.8	34.3	69.0	60.6	80.6
InternVL3.5-2B-Thinking [27]	✗	65.5	91.0	89.5	62.5	50.2	46.5	44.1	58.9	68.5	69.9	74.0
	✓	<b>67.9 (+2.4)</b>	95.7	87.7	68.5	49.3	51.4	35.2	53.8	76.2	83.9	77.8
Qwen2.5-VL-3B [5]	✗	66.0	93.9	79.7	61.2	76.0	45.4	36.9	48.9	62.5	79.7	75.4
	✓	<b>70.8 (+4.8)</b>	97.1	82.1	68.5	76.0	44.9	40.9	51.3	78.5	87.9	81.0
InternVL3.5-4B [27]	✗	67.4	92.3	85.6	69.0	58.0	55.3	41.8	58.4	65.2	71.8	76.8
	✓	<b>70.7 (+2.7)</b>	96.5	87.8	72.8	51.4	54.9	44.0	49.8	76.7	86.9	86.3
Qwen2.5-VL-7B [5]	✗	71.1	93.3	86.5	74.4	60.3	49.5	42.7	58.9	74.0	88.6	83.1
	✓	<b>74.2 (+3.1)</b>	94.9	90.4	75.0	57.4	55.8	48.3	60.4	80.4	93.0	86.1
Qwen3-VL-8B-Instruct [29]	✗	78.8	93.9	93.1	79.7	74.1	65.5	60.2	68.6	75.8	88.3	88.9
	✓	<b>80.3 (+1.5)</b>	95.2	94.9	78.7	68.2	66.9	59.2	70.2	82.7	96.4	90.3
GLM-4.1V-9B-Thinking [25]	✗	77.2	94.4	92.1	77.8	67.8	59.7	55.0	69.3	77.7	89.4	88.9
	✓	<b>79.1 (+1.9)</b>	96.6	93.9	72.4	69.3	62.7	61.1	70.7	80.0	95.3	88.9
Qwen3-VL-235B-A22B-Instruct [29]	✗	83.7	95.8	94.9	84.5	83.2	65.3	68.4	71.6	84.0	97.5	92.1
	✓	<b>84.4 (+0.7)</b>	96.0	96.1	85.1	83.6	66.2	66.8	74.1	85.6	97.5	92.9

60, with the highest scores only around 70. These findings indicate that [current MLLMs still face challenges in capturing fine-grained details in infrared image understanding](#).

### 5.3. Results of GenViP

**Main Results.** As shown in Tab. 4, GenViP consistently improves performance across various models. It achieves over 4 points improvement on InternVL3.5-1B and Qwen2.5-VL-3B. Qwen3-VL-8B-Instruct with GenViP matches the performance of InternVL3-78B and GLM-4.5V-106B-A12B-Thinking in Tab. 2, while using only one-tenth of their parameters. Moreover, Qwen3-VL-235B-A22B-Instruct with GenViP attains a score of 84.4, outperforming all other models in Tab. 2, including closed-source ones. The consistent performance gains across diverse models highlight the effectiveness and generalization ability of GenViP. However, the improvement diminishes as the model scale increases. This trend may stem from two factors: (1) Larger models already achieve high baseline performance, leaving limited room for further improvement. (2) Larger models may possess stronger robustness to input distribution shifts, reducing their reliance on domain adaptation techniques like GenViP.

**Editing Models.** Tab. 5 compares the performance of various image editing models employed in GenViP. All models enable GenViP to effectively enhance infrared image understanding. However, the open-source Qwen-Edit-2509 lags behind top-performing closed-source models such as Seedream 4.0 [21] and Gemini-2.5-Flash [10]. After fine-tuning on our high-quality RGB-T paired dataset, Qwen-Edit-2509-FT effectively bridges this performance gap and even surpasses the closed-source models on this specific task. We provide a visual comparison of the translation quality across different models in Appendix D.

**Inference Input.** Tab. 6 examines the effect of different input configurations on infrared image understanding perfor-

Table 5. Impact of editing models using Qwen2.5-VL-7B [5].

Edit Model	Avg	SU	IT	VC	TL	SRU	OC	TFU	AR	TFR	CR
Vanilla	71.1	93.3	86.5	74.4	60.3	49.5	42.7	58.9	74.0	88.6	83.1
<i>Closed-Source Models</i>											
Seedream 4.0 [21]	73.8	94.2	89.7	77.4	56.9	53.0	50.0	60.0	79.8	92.8	84.1
Gemini-2.5-Flash [10]	73.6	95.7	91.6	79.5	57.4	48.6	50.7	58.8	79.4	92.8	82.1
<i>Open-Source Models</i>											
Qwen-Edit-2509 [28]	72.7	92.6	88.5	75.4	61.3	51.4	46.6	62.5	72.7	92.8	83.5
Qwen-Edit-2509-FT	<b>74.2</b>	94.9	90.4	75.0	57.4	55.8	48.3	60.4	80.4	93.0	86.1

Table 6. Impact of inference input with Qwen2.5-VL-7B [5], where RGB images are generated by Qwen-Edit-2509-FT. Abbreviations: IF – Infrared Image. Text – Textual Prior.

Inference Input	Avg	SU	IT	VC	TL	SRU	OC	TFU	AR	TFR	CR
(a) IF	71.1	93.3	86.5	74.4	<b>60.3</b>	49.5	42.7	58.9	74.0	88.6	83.1
(b) IF + Text	71.2	92.5	85.8	74.8	55.1	51.6	42.7	<b>62.5</b>	71.9	<b>93.0</b>	81.9
(c) RGB	69.2	<b>96.3</b>	76.4	69.2	58.0	54.2	46.1	50.0	<b>80.8</b>	77.8	83.7
(d) IF + RGB	73.2	96.2	<b>91.2</b>	73.5	58.0	55.6	<b>49.3</b>	55.0	80.4	87.3	85.7
(e) IF + RGB + Text	<b>74.2</b>	94.9	90.4	<b>75.0</b>	57.4	<b>55.8</b>	48.3	60.4	80.4	<b>93.0</b>	<b>86.1</b>

mance. Incorporating the textual prior (b) improves *thermal feature understanding and reasoning* compared with using only infrared images (a), but degrades performance on other dimensions, resulting in negligible overall improvement. Using only the translated RGB image (c) causes substantial drops in *image theme, thermal feature understanding and reasoning* due to the loss of thermal information. Nevertheless, by mitigating domain shift, it improves tasks that rely less on infrared-specific cues, such as *spatial relationship understanding, object counting, and action recognition*. Combining both infrared and RGB images (d) consistently enhances performance across most dimensions. Further adding the textual prior yields additional gains in *thermal understanding and reasoning*, leading to the best overall performance on IF-Bench.

## 6. Conclusion

In this work, we present IF-Bench, a comprehensive benchmark for evaluating multimodal understanding of infrared images, along with several evaluation strategies designed to ensure reliable assessment. Through an extensive evaluation of over 40 MLLMs, we uncover several key findings: (1) model size and architecture substantially influence performance; (2) the thinking mode does not effectively improve overall results; (3) open-source and closed-source models exhibit comparable performance; and (4) current MLLMs still perform poorly on fine-grained infrared image understanding. We believe these insights can inspire future advances in this domain. Furthermore, we introduce GenViP to enhance the infrared understanding capability of existing MLLMs. By leveraging generative models to produce information-rich visual prompts, GenViP significantly improves model generalization, offering evidence that generation can effectively facilitate visual understanding. We hope our work provides new perspectives and inspiration for the broader multimodal research community.

## Acknowledgments

This research was supported by the National Natural Science Foundation of China under Grant 62306310 and 62433003.

## References

- [1] Bisma Abidi. Iris thermal/visible face database. 1
- [2] Jean-Baptiste Alayrac, Jeff Donahue, Pauline Luc, Antoine Miech, Iain Barr, Yana Hasson, Karel Lenc, Arthur Mensch, Katherine Millican, Malcolm Reynolds, Roman Ring, Eliza Rutherford, Serkan Cabi, Tengda Han, Zhitao Gong, Sina Samangooei, Marianne Monteiro, Jacob L Menick, Sebastian Borgeaud, Andy Brock, Aida Nematzadeh, Sahand Sharifzadeh, Mikolaj Binkowski, Ricardo Barreira, Oriol Vinyals, Andrew Zisserman, and Karén Simonyan. Flamingo: a visual language model for few-shot learning. In *NeurIPS*, 2022. 2
- [3] Xiang An, Yin Xie, Kaicheng Yang, Wenkang Zhang, Xiuwei Zhao, Zheng Cheng, Yirui Wang, Songcen Xu, Changrui Chen, Chunsheng Wu, Huajie Tan, Chunyuan Li, Jing Yang, Jie Yu, Xiyao Wang, Bin Qin, Yumeng Wang, Zizhen Yan, Ziyong Feng, Ziwei Liu, Bo Li, and Jiankang Deng. Llava-onevision-1.5: Fully open framework for democratized multimodal training. *arXiv preprint arXiv:2509.23661*, 2025. 2, 6, 7
- [4] Chris H. Bahnsen and Thomas B. Moeslund. Rain Removal in Traffic Surveillance: Does it matter? *IEEE TITS*, 2019. 1
- [5] Shuai Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Sibao Song, Kai Dang, Peng Wang, Shijie Wang, Jun Tang, Humen Zhong, Yuanzhi Zhu, Mingkun Yang, Zhao-hai Li, Jianqiang Wan, Pengfei Wang, Wei Ding, Zheren Fu, Yiheng Xu, Jiabo Ye, Xi Zhang, Tianbao Xie, Zesen Cheng, Hang Zhang, Zhibo Yang, Haiyang Xu, and Junyang Lin. Qwen2.5-vl technical report. *arXiv preprint arXiv:2502.13923*, 2025. 2, 4, 5, 6, 7, 8
- [6] ByteDance Seed. Doubao-seed-1.6. <https://console.volcengine.com/ark/region:ark+cn-beijing/model/detail?Id=doubao-seed-1-6>, 2025. 2, 6
- [7] ByteDance Seed. Doubao-seed-vision-1.6. <https://console.volcengine.com/ark/region:ark+cn-beijing/model/detail?Id=doubao-seed-1-6-vision>, 2025. 2, 6, 7
- [8] John Canny. A computational approach to edge detection. *IEEE TPAMI*, 1986. 5
- [9] Zhe Cao, Jin Zhang, and Ruiheng Zhang. Irgpt: Understanding real-world infrared image with bi-cross-modal curriculum on large-scale benchmark. In *ICCV*, 2025. 1, 2, 3
- [10] Gheorghe Comanici, Eric Bieber, Mike Schaekermann, Ice Pasupat, Noveen Sachdeva, Inderjit Dhillon, Marcel Blstein, Ori Ram, Dan Zhang, Evan Rosen, Luke Marris, Sam Petulla, Colin Gaffney, Asaf Aharoni, Nathan Lintz, Tiago Cardal Pais, Henrik Jacobsson, Idan Szepktor, Nanjiang Jiang, Krishna Haridasan, Ahmed Omran, Nikunj Saunshi, Dara Bahri, Gaurav Mishra, Eric Chu, Toby Boyd, Brad Hekman, Aaron Parisi, Chaoyi Zhang, Kornrhapop Kawintiranon, Tania Bedrax-Weiss, Oliver Wang, Ya Xu, Ollie Purkiss, Uri Mendlovic, Ilai Deutel, Nam Nguyen, Adam Langley, Flip Korn, Lucia Rossazza, Alexandre Ramé, Sagar Waghmare, Helen Miller, Nathan Byrd, Ashrith Seshan, Raia Hadsell, Sangnie Bhardwaj, Pawel Janus, Tero Rissa, Dan Horgan, Alvin Abdagic, Lior Belenki, James Allingham, Anima Singh, Theo Guidroz, Srivatsan Srinivasan, Herman Schmit, Kristen Chiafullo, Andre Elisseeff, Nilpa Jha, Prateek Kolhar, Leonard Berrada, Frank Ding, Xiance Si, Shrestha Basu Mallick, Franz Och, Sofia Erell, Eric Ni, Tejasi Latkar, Sherry Yang, Petar Sirkovic, Ziqiang Feng, Robert Leland, Rachel Hornung, Gang Wu, Charles Blundell, Hamidreza Alvari, Po-Sen Huang, Cathy Yip, Sanja Deur, Li Liu, Gabriela Surita, et al. Gemini 2.5: Pushing the frontier with advanced reasoning, multimodality, long context, and next generation agentic capabilities. *arXiv preprint arXiv:2507.06261*, 2025. 1, 2, 5, 6, 7, 8
- [11] Edward J Hu, yelong shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. LoRA: Low-rank adaptation of large language models. In *ICLR*, 2022. 1
- [12] Xinyu Jia, Chuang Zhu, Minzhen Li, Wenqi Tang, and Wenli Zhou. LLVIP: A visible-infrared paired dataset for low-light vision. In *ICCVW*, 2021. 1
- [13] Shixin Jiang, Zerui Chen, Jiafeng Liang, Yanyan Zhao, Ming Liu, and Bing Qin. Infrared-LLaVA: Enhancing understanding of infrared images in multi-modal large language models. In *EMNLP*, 2024. 1, 2, 3
- [14] Kunchang Li, Yali Wang, Yanan He, Yizhuo Li, Yi Wang, Yi Liu, Zun Wang, Jilan Xu, Guo Chen, Ping Luo, Limin Wang, and Yu Qiao. Mvbench: A comprehensive multi-modal video understanding benchmark. In *CVPR*, 2024. 3
- [15] Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. Visual instruction tuning. In *NeurIPS*, 2023. 2
- [16] Yuan Liu, Haodong Duan, Yuanhan Zhang, Bo Li, Songyang Zhang, Wangbo Zhao, Yike Yuan, Jiaqi Wang, Conghui He, Ziwei Liu, Kai Chen, and Dahua Lin. Mmbench: Is your multi-modal model an all-around player? In *ECCV*, 2024. 1, 3, 5
- [17] OpenAI. Thinking with images. <https://openai.com/index/thinking-with-images/>, 2025. 1
- [18] OpenAI, Aaron Hurst, Adam Lerer, Adam P. Goucher, Adam Perelman, Aditya Ramesh, Aidan Clark, AJ Ostrow, Akila Welihinda, Alan Hayes, Alec Radford, Aleksander Mądry, Alex Baker-Whitcomb, Alex Beutel, Alex Borzunov, Alex Carney, Alex Chow, Alex Kirillov, Alex Nichol, Alex Paino, Alex Renzin, Alex Tachard Passos, Alexander Kirillov, Alexi Christakis, Alexis Conneau, Ali Kamali, Allan Jabri, Allison Moyer, Allison Tam, Amadou Crookes, Amin Tootoochian, Amin Tootoonchian, Ananya Kumar, Andrea Vallone, Andrej Karpathy, Andrew Braunstein, Andrew Cann, Andrew Codispoti, Andrew Galu, Andrew Kondrich, Andrew Tulloch, Andrey Mishchenko, Angela Baek, Angela Jiang, Antoine Pelisse, Antonia Woodford, Anuj Gosalia, Arka Dhar, Ashley Pantuliano, Avi Nayak, Avital Oliver, Barret Zoph, Behrooz Ghorbani, Ben Leimberger, Ben Rossen, Ben Sokolowsky, Ben Wang, Benjamin Zweig, Beth Hoover, Blake Samic, Bob McGrew, Bobby Spero, Bogo Gertler,

- Bowen Cheng, Brad Lightcap, Brandon Walkin, Brendan Quinn, Brian Guarraci, Brian Hsu, Bright Kellogg, Brydon Eastman, Camillo Lugaresi, Carroll Wainwright, Cary Bassin, Cary Hudson, Casey Chu, Chad Nelson, Chak Li, Chan Jun Shern, Channing Conger, Charlotte Barette, Chelsea Voss, Chen Ding, Cheng Lu, Chong Zhang, Chris Beaumont, Chris Hallacy, Chris Koch, Christian Gibson, Christina Kim, Christine Choi, Christine McLeavey, Christopher Hesse, Claudia Fischer, Clemens Winter, Coley Czarnecki, Colin Jarvis, Colin Wei, Constantin Koumouzelis, Dane Sherburn, Daniel Kappler, et al. Gpt-4o system card. *arXiv preprint arXiv:2410.21276*, 2024. 1, 2
- [19] Cristina Palmero, Albert Clapés, Chris Bahnsen, Andreas Møgelmoose, Thomas B Moeslund, and Sergio Escalera. Multi-modal rgb–depth–thermal human body segmentation. *IJCV*, 2016. 1
- [20] Atin Pothiraj, Elias Stengel-Eskin, Jaemin Cho, and Mohit Bansal. Capture: Evaluating spatial reasoning in vision language models via occluded object counting. *arXiv preprint arXiv:2504.15485*, 2025. 3
- [21] Team Seedream, Yunpeng Chen, Yu Gao, Lixue Gong, Meng Guo, Qiushan Guo, Zhiyao Guo, Xiaoxia Hou, Weilin Huang, Yixuan Huang, Xiaowen Jian, Huafeng Kuang, Zhichao Lai, Fanshi Li, Liang Li, Xiaochen Lian, Chao Liao, Liyang Liu, Wei Liu, Yanzuo Lu, Zhengxiong Luo, Tongtong Ou, Guang Shi, Yichun Shi, Shiqi Sun, Yu Tian, Zhi Tian, Peng Wang, Rui Wang, Xun Wang, Ye Wang, Guofeng Wu, Jie Wu, Wenxu Wu, Yonghui Wu, Xin Xia, Xuefeng Xiao, Shuang Xu, Xin Yan, Ceyuan Yang, Jianchao Yang, Zhonghua Zhai, Chenlin Zhang, Heng Zhang, Qi Zhang, Xinyu Zhang, Yuwei Zhang, Shijia Zhao, Wenliang Zhao, and Wenjia Zhu. Seedream 4.0: Toward next-generation multimodal image generation. *arXiv preprint arXiv:2509.20427*, 2025. 2, 5, 8, 1
- [22] Oriane Siméoni, Huy V. Vo, Maximilian Seitzer, Federico Baldassarre, Maxime Oquab, Cijo Jose, Vasil Khalidov, Marc Szafraniec, Seungeun Yi, Michaël Ramamonjisoa, Francisco Massa, Daniel Haziza, Luca Wehrstedt, Jianyuan Wang, Timothée Darcet, Théo Moutakanni, Leonel Sentana, Claire Roberts, Andrea Vedaldi, Jamie Tolan, John Brandt, Camille Couprie, Julien Mairal, Hervé Jégou, Patrick Labatut, and Piotr Bojanowski. Dinov3. *arXiv preprint arXiv:2508.10104*, 2025. 6
- [23] Yiming Sun, Bing Cao, Pengfei Zhu, and Qinghua Hu. Drone-based rgb-infrared cross-modality vehicle detection via uncertainty-aware learning. *IEEE TCSVT*, 2022. 1
- [24] TELEDYNE FLIR Team. FLIR thermal dataset for algorithm training. Online, 2019. 1
- [25] V Team, Wenyi Hong, Wenmeng Yu, Xiaotao Gu, Guo Wang, Guobing Gan, Haomiao Tang, Jiale Cheng, Ji Qi, Junhui Ji, Lihang Pan, Shuaiqi Duan, Weihang Wang, Yan Wang, Yean Cheng, Zehai He, Zhe Su, Zhen Yang, Ziyang Pan, Aohan Zeng, Baoxu Wang, Bin Chen, Boyan Shi, Changyu Pang, Chenhui Zhang, Da Yin, Fan Yang, Guoqing Chen, Jiazheng Xu, Jiale Zhu, Jiali Chen, Jing Chen, Jinhao Chen, Jinghao Lin, Jinjiang Wang, Junjie Chen, Leqi Lei, Letian Gong, Leyi Pan, Mingdao Liu, Mingde Xu, Mingzhi Zhang, Qinkai Zheng, Sheng Yang, Shi Zhong, Shiyu Huang, Shuyuan Zhao, Siyan Xue, Shangqin Tu, Shengbiao Meng, Tianshu Zhang, Tianwei Luo, Tianxiang Hao, Tianyu Tong, Wenkai Li, Wei Jia, Xiao Liu, Xiaohan Zhang, Xin Lyu, Xinyue Fan, Xuancheng Huang, Yanling Wang, Yadong Xue, Yanfeng Wang, Yanzi Wang, Yifan An, Yifan Du, Yiming Shi, Yiheng Huang, Yilin Niu, Yuan Wang, Yuanchang Yue, Yuchen Li, Yutao Zhang, Yuting Wang, Yu Wang, Yuxuan Zhang, Zhao Xue, Zhenyu Hou, Zhengxiao Du, Zihan Wang, Peng Zhang, Debing Liu, Bin Xu, Juanzi Li, Minlie Huang, Yuxiao Dong, and Jie Tang. Glm-4.5v and glm-4.1v-thinking: Towards versatile multimodal reasoning with scalable reinforcement learning. *arXiv preprint arXiv:2507.01006*, 2025. 6, 7, 8
- [26] Huy V. Vo, Vasil Khalidov, Timothée Darcet, Théo Moutakanni, Nikita Smetanin, Marc Szafraniec, Hugo Touvron, camille couprie, Maxime Oquab, Armand Joulin, Herve Jegou, Patrick Labatut, and Piotr Bojanowski. Automatic data curation for self-supervised learning: A clustering-based approach. *TMLR*, 2024. 6
- [27] Weiyun Wang, Zhangwei Gao, Lixin Gu, Hengjun Pu, Long Cui, Xingguang Wei, Zhaoyang Liu, Linglin Jing, Shenglong Ye, Jie Shao, Zhaokai Wang, Zhe Chen, Hongjie Zhang, Ganlin Yang, Haomin Wang, Qi Wei, Jinhui Yin, Wenhao Li, Erfei Cui, Guanzhou Chen, Zichen Ding, Changyao Tian, Zhenyu Wu, Jingjing Xie, Zehao Li, Bowen Yang, Yuchen Duan, Xuehui Wang, Zhi Hou, Haoran Hao, Tianyi Zhang, Songze Li, Xiangyu Zhao, Haodong Duan, Nianchen Deng, Bin Fu, Yinan He, Yi Wang, Conghui He, Botian Shi, Junjun He, Yingtong Xiong, Han Lv, Lijun Wu, Wenqi Shao, Kaipeng Zhang, Huipeng Deng, Biqing Qi, Jiaye Ge, Qipeng Guo, Wenwei Zhang, Songyang Zhang, Maosong Cao, Junyao Lin, Kexian Tang, Jianfei Gao, Haihan Huang, Yuzhe Gu, Chengqi Lyu, Huanze Tang, Rui Wang, Haijun Lv, Wanli Ouyang, Limin Wang, Min Dou, Xizhou Zhu, Tong Lu, Dahua Lin, Jifeng Dai, Weijie Su, Bowen Zhou, Kai Chen, Yu Qiao, Wenhao Wang, and Gen Luo. Internvl3.5: Advancing open-source multimodal models in versatility, reasoning, and efficiency. *arXiv preprint arXiv:2508.18265*, 2025. 2, 6, 7, 8, 1, 4
- [28] Chenfei Wu, Jiahao Li, Jingren Zhou, Junyang Lin, Kaiyuan Gao, Kun Yan, Sheng ming Yin, Shuai Bai, Xiao Xu, Yilei Chen, Yuxiang Chen, Zecheng Tang, Zekai Zhang, Zhengyi Wang, An Yang, Bowen Yu, Chen Cheng, Dayiheng Liu, Deqing Li, Hang Zhang, Hao Meng, Hu Wei, Jingyuan Ni, Kai Chen, Kuan Cao, Liang Peng, Lin Qu, Minggang Wu, Peng Wang, Shuting Yu, Tingkun Wen, Wensen Feng, Xiaoxiao Xu, Yi Wang, Yichang Zhang, Yongqiang Zhu, Yujia Wu, Yuxuan Cai, and Zenan Liu. Qwen-image technical report. *arXiv preprint arXiv:2508.02324*, 2025. 2, 5, 6, 8, 1
- [29] An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, Chujie Zheng, Dayiheng Liu, Fan Zhou, Fei Huang, Feng Hu, Hao Ge, Haoran Wei, Huan Lin, Jialong Tang, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Yang, Jiayi Yang, Jing Zhou, Jingren Zhou, Junyang Lin, Kai Dang, Keqin Bao, Kexin Yang, Le Yu, Lianghao Deng, Mei Li, Mingfeng Xue, Mingze Li, Pei Zhang, Peng Wang, Qin Zhu, Rui Men, Ruize Gao, Shixuan Liu, Shuang

- Luo, Tianhao Li, Tianyi Tang, Wenbiao Yin, Xingzhang Ren, Xinyu Wang, Xinyu Zhang, Xuancheng Ren, Yang Fan, Yang Su, Yichang Zhang, Yinger Zhang, Yu Wan, Yuqiong Liu, Zekun Wang, Zeyu Cui, Zhenru Zhang, Zhipeng Zhou, and Zihan Qiu. Qwen3 technical report. *arXiv preprint arXiv:2505.09388*, 2025. [1](#), [2](#), [5](#), [6](#), [7](#), [8](#), [4](#)
- [30] Biao Yang, Bin Wen, Boyang Ding, Changyi Liu, Chenglong Chu, Chengru Song, Chongling Rao, Chuan Yi, Da Li, Dunju Zang, Fan Yang, Guorui Zhou, Guowang Zhang, Han Shen, Hao Peng, Haojie Ding, Hao Wang, Haonan Fan, Hengrui Ju, Jiaming Huang, Jiangxia Cao, Jiankang Chen, Jingyun Hua, Kaibing Chen, Kaiyu Jiang, Kaiyu Tang, Kun Gai, Muhao Wei, Qiang Wang, Ruitao Wang, Sen Na, Shengnan Zhang, Siyang Mao, Sui Huang, Tianke Zhang, Tingting Gao, Wei Chen, Wei Yuan, Xiangyu Wu, Xiao Hu, Xingyu Lu, Yi-Fan Zhang, Yiping Yang, Yulong Chen, Zeyi Lu, Zhenhua Wu, Zhixin Ling, Zhuoran Yang, Ziming Li, Di Xu, Haixuan Gao, Hang Li, Jing Wang, Lejian Ren, Qigen Hu, Qianqian Wang, Shiyao Wang, Xincheng Luo, Yan Li, Yuhang Hu, and Zixing Zhang. Kwai keye-v1 1.5 technical report. *arXiv preprint arXiv:2509.01563*, 2025. [6](#), [7](#)
- [31] Sihan Yang, Runsen Xu, Yiman Xie, Sizhe Yang, Mo Li, Jingli Lin, Chenming Zhu, Xiaochen Chen, Haodong Duan, Xiangyu Yue, Dahua Lin, Tai Wang, and Jiangmiao Pang. Mmsi-bench: A benchmark for multi-image spatial intelligence. *arXiv preprint arXiv:2505.23764*, 2025. [3](#)
- [32] Weihao Yu, Zhengyuan Yang, Linjie Li, Jianfeng Wang, Kevin Lin, Zicheng Liu, Xinchao Wang, and Lijuan Wang. MM-vet: Evaluating large multimodal models for integrated capabilities. In *ICML*, 2024. [3](#)
- [33] Xiang Yue, Yuansheng Ni, Kai Zhang, Tianyu Zheng, Ruoqi Liu, Ge Zhang, Samuel Stevens, Dongfu Jiang, Weiming Ren, Yuxuan Sun, Cong Wei, Botao Yu, Ruibin Yuan, Renliang Sun, Ming Yin, Boyuan Zheng, Zhenzhu Yang, Yibo Liu, Wenhao Huang, Huan Sun, Yu Su, and Wenhui Chen. Mmmu: A massive multi-discipline multimodal understanding and reasoning benchmark for expert agi. In *CVPR*, 2024. [1](#), [3](#)
- [34] Tao Zhang, Jinyong Wen, Zhen Chen, Kun Ding, Shiming Xiang, and Chunhong Pan. UNIP: Rethinking pre-trained attention patterns for infrared semantic segmentation. In *ICLR*, 2025. [4](#), [1](#), [3](#)
- [35] Jinguo Zhu, Weiyun Wang, Zhe Chen, Zhaoyang Liu, Shenglong Ye, Lixin Gu, Hao Tian, Yuchen Duan, Weijie Su, Jie Shao, Zhangwei Gao, Erfei Cui, Xuehui Wang, Yue Cao, Yangzhou Liu, Xingguang Wei, Hongjie Zhang, Haomin Wang, Weiye Xu, Hao Li, Jiahao Wang, Nianchen Deng, Songze Li, Yinan He, Tan Jiang, Jiapeng Luo, Yi Wang, Conghui He, Botian Shi, Xingcheng Zhang, Wenqi Shao, Junjun He, Yingdong Xiong, Wenwen Qu, Peng Sun, Penglong Jiao, Han Lv, Lijun Wu, Kaipeng Zhang, Huipeng Deng, Jiaye Ge, Kai Chen, Limin Wang, Min Dou, Lewei Lu, Xizhou Zhu, Tong Lu, Dahua Lin, Yu Qiao, Jifeng Dai, and Wenhui Wang. Internvl3: Exploring advanced training and test-time recipes for open-source multimodal models. *arXiv preprint arXiv:2504.10479*, 2025. [6](#), [7](#)

# IF-Bench: Benchmarking and Enhancing MLLMs for Infrared Images with Generative Visual Prompting

## Supplementary Material

We illustrate additional example cases and the image distribution of IF-Bench in Appendix A. All prompts used in the paper are listed in Appendix B. Appendix C details the training and inference settings for editing models in GenViP. In Appendix D, we conduct further analyses, including comparisons of the language preferences of MLLMs on IF-Bench, the stricter correctness judgement strategy, the connection between our work and Thinking-with-Image, and the translated quality across editing models.

### A. Details of IF-Bench

In this section, we present additional details of IF-Bench. Fig. 7 illustrates some cases in IF-Bench, with images, questions, options, and ground-truth answers. Fig. 8 depicts the image distribution in IF-Bench across 23 sub-datasets in InfPre [34]. Notably, after manual filtering, the images in IF-Bench remain relatively evenly distributed across the constituent sub-datasets, ensuring comprehensive coverage of diverse scenarios. The datasets OTCBVS-IRIS-FACE [1], VAP [19], and Rain [4] contain fewer samples primarily because their scenes are comparatively homogeneous.

### B. Details of Prompts

In this section, we provide the specific prompts used in this work. Fig. 10 illustrates the unified system prompt when evaluating MLLMs on IF-Bench. Fig. 11 depicts LLM-based parsing prompt, as described in Sec. 3.3. Tab. 12 and Tab. 13 provide the evaluation prompt for the dual-image input in GenViP and the textual prior, respectively. For simplicity, only the English versions are presented here, with the Chinese versions being semantically equivalent to their English counterparts. Fig. 14 lists the system prompt used for automatic VQA generation in Sec. 3.2.

### C. Details of Editing Models

In this section, we clarify the training and inference details involved in the editing models of GenViP.

**Inference of Editing Models.** For Seedream 4.0 [21] and Gemini-2.5-Flash [10], we use their official APIs for infrared-to-RGB translation. For Qwen-Edit-2509 [28], inference is performed via local deployment. The number of inference steps is set to 40 for both Qwen-Edit-2509 and Qwen-Edit-2509-FT.

**Training of Qwen-Edit-2509.** We fine-tune Qwen-Edit-2509 [28] using LoRA [11] with a rank of 32. Qwen-Edit-2509-FT is trained for two epochs on the RGB-T dataset

consisting of 50,000 pairs, with a batch size of 8 and a fixed resolution of 1024.

**Editing Prompt.** For both the training of Qwen-Edit-2509-FT and inference of all editing models, we use the same editing prompt: *Translate the infrared image into the corresponding visible light (RGB) image.* We also experimented with more sophisticated editing prompts, but found that they led to similar results. Therefore, we adopt this simple prompt throughout.

### D. More Analyses

**Language Preference.** Fig. 9 presents the Chinese and English scores of various MLLMs on IF-Bench. Most models exhibit no notable language preference, with score differences between the Chinese and English settings remaining within one point. Only a few smaller models, such as InternVL3.5-1B [27], InternVL3.5-2B [27], and Qwen3-VL-4B-Instruct [29], show slightly larger discrepancies of approximately one to two points.

**Stricter Correctness Judgement.** In the main paper, we adopt a circular evaluation strategy and report the average accuracy across different option orders. Here, we further apply a more stringent correctness criterion: a question is considered correct only if the model answers it correctly under all option orders. Tab. 7 compares the score on IF-Bench of the two criteria. Notably, the stricter criterion substantially reduces accuracy, with the average score dropping by more than 10 points for several models. This indicates that current MLLMs exhibit limited robustness to option-order variations and still have considerable room for improvement in infrared image understanding.

**Relationship with Thinking-with-Image.** Thinking-with-Image was introduced in OpenAI o3 [17], where the model performs a variety of operations on the input image to actively generate intermediate visual representations, thereby enhancing its capability in multimodal reasoning. Our proposed GenViP shares a similar intuition: processing the image to improve the model’s understanding. But the key difference is that in our work, the model is explicitly required to invoke an editing model at the initial stage. The idea behind GenViP can be naturally extended to the Thinking-with-Image paradigm, where the model could be trained to autonomously decide whether to call an editing model. We leave this direction for future work.

**Visualization of Translated Quality.** Fig. 15 illustrates the translated RGB images produced by different editing models. We observe that for some infrared inputs, closed-

### Scene Understanding



What type of environment does this image depict?

- A: Rural road.
- B: Urban highway.
- C: Suburban neighborhood.
- D: Industrial area.

GT: B

### Viewpoint of Capture



Compared to people in the elevator, from which viewpoint was this image captured?

- A: Top-down.
- B: Frontal.
- C: Side-view.
- D: Bottom-up.

GT: A

### Spatial Relationship Understanding



What is the spatial relationship between the two cars in the image?

- A: The car on the right is behind the car on the left.
- B: The cars are side by side.
- C: The car on the right is in front of the car on the left.
- D: The cars are not visible.

GT: C

### Thermal Feature Understanding



Which area in the image appears to have the highest thermal activity?

- A: The crowd of people.
- B: The trees on the right.
- C: The ground in the foreground.
- D: The sky above.

GT: A

### Thermal Feature Reasoning



What could be the reason for the uniform heat signatures among the individuals?

- A: They are all performing different actions.
- B: They are standing in a cold environment.
- C: The ambient temperature is very high.
- D: They are wearing identical clothing materials.

GT: D

### Image Theme



What is the most likely theme of this infrared image?

- A: Traffic surveillance.
- B: Wildlife monitoring.
- C: Security surveillance.
- D: Agricultural monitoring.

GT: A

### Target Localization



What are the coordinates of the person closest to the straight pole in the image?

- A: (0.4, 0.8, 0.1, 0.3).
- B: (0.54, 0.9, 0.06, 0.2).
- C: (0.6, 0.7, 0.05, 0.15).
- D: (0.5, 0.95, 0.08, 0.25).

GT: B

### Object Counting



How many people are visible in the image?

- A: Between 20 and 30.
- B: Between 10 and 20.
- C: Less than 10.
- D: More than 30.

GT: A

### Action Recognition



What action is the animal in the image performing?

- A: Running.
- B: Walking.
- C: Jumping.
- D: Standing still.

GT: D

### Commonsense Reasoning



Based on the image, what is the most likely purpose of the barrier in the foreground?

- A: To guide traffic flow.
- B: To block unauthorized access.
- C: To provide shade.
- D: To serve as a decorative element.

GT: B

Figure 7. Example cases for each dimension in IF-Bench, aligned with the questions and options in Tab. 2.

source models such as Seedream 4.0 [21] and Gemini-2.5-Flash [10] show strong generalization ability (e.g., rows 1, 3, and 5). However, their translation quality remains suboptimal for certain other cases, exhibiting issues such as unnatural color rendering and imperfect spatial correspondence. Qwen-Edit-2509 [28] performs worse overall

compared to the closed-source models. After fine-tuning on 50,000 infrared- $\text{RGB}$  image pairs, Qwen-Edit-2509-FT significantly improves the naturalness of the translated images as well as their spatial and semantic consistency, thereby leading to a substantial performance gain for Gen-ViP.

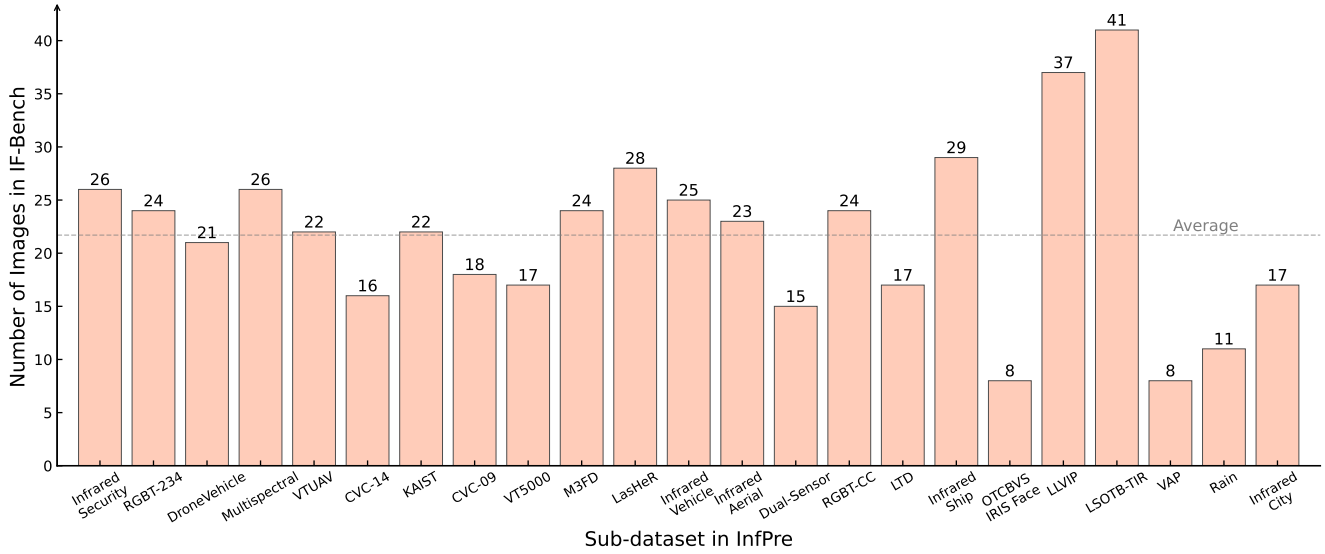


Figure 8. The distribution images in IF-Bench across various sub-datasets in InfPre [34].




Figure 9. The language-specific score on IF-Bench of various MLLMs.

## E. Limitations and Future Work

The current IF-Bench contains a relatively limited number of images and questions, and does not cover some more challenging task types. In future work, we will further expand and diversify the benchmark to keep pace with the rapid progress of MLLMs and explore more methods to improve infrared understanding.

Table 7. Comparison of different correctness judgement strategies.

Model	Stricter Strategy	Avg	SU	IT	VC	TL	SRU	OC	TFU	AR	TFR	CR
Qwen2.5-VL-7B [5]	✗	<b>71.1</b>	93.3	86.5	74.4	60.3	49.5	42.7	58.9	74.0	88.6	83.1
	✓	53.3 (-17.8)	89.7	71.6	55.2	10.3	31.5	17.8	43.6	58.5	77.1	77.4
Qwen3-VL-8B-Instruct [29]	✗	<b>78.8</b>	93.9	93.1	79.7	74.1	65.5	60.2	68.6	75.8	88.3	88.9
	✓	69.1 (-9.7)	91.7	89.2	69.0	54.8	51.9	40.8	61.4	66.2	81.4	84.7
InternVL3.5-30B-A3B [27]	✗	<b>74.4</b>	94.9	88.2	73.7	75.0	59.0	45.0	69.0	68.1	85.8	85.7
	✓	58.8 (-15.6)	88.5	79.7	59.5	50.0	32.4	26.4	52.9	58.5	65.3	75.0
InternVL3.5-38B [27]	✗	<b>79.0</b>	93.4	93.2	77.8	80.5	57.6	59.6	70.0	77.7	91.5	88.7
	✓	65.2 (-13.8)	89.1	80.4	66.4	62.3	43.5	29.3	55.7	62.3	80.5	82.3
Qwen3-VL-235B-A22B-Instruct [29]	✗	<b>83.7</b>	95.8	94.9	84.5	83.2	65.3	68.4	71.6	84.0	97.5	92.1
	✓	78.5 (-5.2)	94.9	90.5	75.9	74.0	57.4	59.2	67.1	80.0	96.6	89.5

You are a professional multimodal large language model assistant. You will be given a single-choice question that includes:

1. One infrared image.
2. One question related to the image.
3. Four answer options (A, B, C, D).

Your task:  
Carefully analyze the image and the question, evaluate all answer choices, and select the most appropriate one. Please output only a single uppercase letter (A, B, C, or D) as your final answer. Do not include any explanations, reasoning, or additional text.

Evaluation Guidelines:

1. Each question has only one correct answer.
2. Random guessing is not allowed; answers must be based on accurate analysis of the image and the question.
3. The output format must be a single uppercase letter: A, B, C, or D.

The input question is: <input\_question>.  
The input options are: <input\_options>.

Figure 10. The unified system prompt when evaluating MLLMs on IF-Bench (English Version).

You are a professional large model. You will receive responses from a model to single-choice questions. The options for each question are A, B, C, and D. Please determine which option (A, B, C, or D) the response considers to be the correct answer. If you believe that none of the options A, B, C, or D are present in the response, please reply with "F".

Remember, you can only reply with one of the characters A, B, C, D, or F. Do not output anything else.

Input: <input\_response>

Figure 11. The system prompt for LLM parsing in Sec. 3.3.

You are a professional multimodal large language model assistant. You will be given a single-choice question that includes:

1. One infrared image and one corresponding RGB image. The RGB image is translated from the corresponding infrared image by an image translation model.
2. One question related to the image.
3. Four answer options (A, B, C, D).

Your task:  
Carefully analyze the image and the question, evaluate all answer choices, and select the most appropriate one. Please output only a single uppercase letter (A, B, C, or D) as your final answer. Do not include any explanations, reasoning, or additional text.

Evaluation Guidelines:

1. Each question has only one correct answer.
2. Random guessing is not allowed; answers must be based on accurate analysis of the image and the question.
3. The output format must be a single uppercase letter: A, B, C, or D.

The input question is: <input\_question>.  
The input options are: <input\_options>.

Figure 12. The evaluation prompt when using the dual-image input in GenViP (English Version).

When completing the above tasks, please refer to the following prior knowledge about infrared images:

1. Imaging Mechanism: Infrared imaging does not rely on visible light reflected from objects but instead captures the infrared radiation (thermal radiation) emitted by the objects themselves or their environment. The higher the temperature of an object, the stronger its infrared radiation; therefore, infrared images usually reflect temperature distribution rather than surface color. Infrared imaging is insensitive to lighting conditions and can function even in complete darkness.

2. Image Characteristics: Infrared images are usually presented in grayscale, where brightness corresponds to temperature. The resolution of infrared images is generally lower, resulting in less detail and poorer edge sharpness. Due to environmental interference (e.g., atmospheric absorption, sensor noise), infrared images often contain more noise. Moreover, different materials have different infrared emissivities at the same temperature, which may cause brightness differences in the resulting images.

Figure 13. The textual prior for infrared images (English Version).

### ## Task

You are an expert in designing infrared image understanding questions. You will be provided with an infrared image. Your task is to design 4 single-choice questions for each image. Each question should include the following components: Question, Options (A/B/C/D), Answer (Correct Option). The questions should be drawn from the 4 evaluation dimensions: <sampled\_dimensin\_list>. You should provide one question for each dimension.

(Optional) The provided image include target position information in the following format: [(Target Description, Target Center X, Target Center Y, Target Width, Target Height)]. All position values are normalized relative to the full image, ranging from 0 to 1. The coordinate origin (0, 0) is located at the top-left corner of the image. The target position information is <target\_position>. You must include one precise localization question with the definition of format, as described in the "2.1 Target Localization" dimension.

### ## Evaluation Dimensions

#### ### 1. Coarse-Grained Image Perception

##### 1.1 Scene Understanding

Definition: Identify the general environment or scene of the image, such as indoor, outdoor, forest, highway, etc.

##### 1.2 Image Theme

Definition: Identify the primary content or application context of the image, such as traffic surveillance, wildlife monitoring, security surveillance, etc.

##### 1.3 Viewpoint of Capture

Definition: Determine the viewpoint from which the image was taken, such as top-down, frontal, side-view, etc.

#### ### 2. Fine-Grained Image Perception

##### 2.1 Target Localization

Definition: Identify the spatial location of a target in the image. If bounding box data is provided, a precise localization question must be created; otherwise, a rough location question can be designed.

Precise localization (when coordinates are given, e.g., "Leftmost pedestrian, 0.2, 0.5, 0.05, 0.18". Format: (x, y, w, h), all values are relative to the full image, with x and y being the center coordinates):

Example Question: What are the coordinates of the leftmost pedestrian in the image?

- A. (0.3, 0.4, 0.1, 0.4)
- B. (0.08, 0.67, 0.2, 0.5)
- C. (0.15, 0.7, 0.08, 0.23)
- D. (0.2, 0.5, 0.05, 0.18)

Rough localization (when no coordinates are provided):

Example Question: Where is the target located in the image?

- A. Top-left corner
- B. Center-right
- C. Bottom-middle
- D. Top-right corner

##### 2.2 Spatial Relationship Understanding

Definition: Recognize the spatial relationship between multiple objects in the image, such as front/back, left/right, near/far, etc.

##### 2.3 Object Counting

Definition: Count the number of similar objects in the image.

##### 2.4 Thermal Feature Understanding

Definition: Determine thermal variations across image regions, such as identifying heat sources or thermal leaks.

##### 2.5 Action Recognition

Definition: Determine whether a target is active and identify its action based on infrared contours.

#### 3. Image Reasoning Ability

##### 3.1 Thermal Feature Reasoning

Definition: Determine the potential cause of temperature changes using thermal patterns.

##### 3.2 Commonsense Reasoning

Definition: Infer the function or use of an object based on image content, infer the action the object in the image is about to do.

### ## Output Format Example

Please generate 4 questions per image using the following JSON structure:

```
[
  {
    "dimension": "Image Theme",
    "en_question": "What is the most likely theme of this infrared image?",
    "en_options": {
      "A": "Industrial equipment monitoring",
      "B": "Wildlife observation",
      "C": "Medical diagnosis",
      "D": "Indoor entertainment"
    },
    "answer": "B",
    "cn_question": "这张红外图像最可能属于哪个主题?",
    "cn_options": {
      "A": "工业设备监控",
      "B": "野生动物观察",
      "C": "医疗诊断",
      "D": "室内娱乐"
    }
  },
]
```

Figure 14. The system prompt used for VQA generation in Sec. 3.2.




Figure 15. The visualization of translated RGB images from infrared ones across different editing models.

## References

- [1] Besma Abidi. Iris thermal/visible face database. 1
- [2] Jean-Baptiste Alayrac, Jeff Donahue, Pauline Luc, Antoine Miech, Iain Barr, Yana Hasson, Karel Lenc, Arthur Mensch, Katherine Millican, Malcolm Reynolds, Roman Ring, Eliza Rutherford, Serkan Cabi, Tengda Han, Zhitao Gong, Sina Samangooei, Marianne Monteiro, Jacob L Menick, Sebastian Borgeaud, Andy Brock, Aida Nematzadeh, Sahand Sharifzadeh, Mikolaj Binkowski, Ricardo Barreira, Oriol Vinyals, Andrew Zisserman, and Karén Simonyan. Flamingo: a visual language model for few-shot learning. In *NeurIPS*, 2022. 2
- [3] Xiang An, Yin Xie, Kaicheng Yang, Wenkang Zhang, Xiuwei Zhao, Zheng Cheng, Yirui Wang, Songcen Xu, Changrui Chen, Chunsheng Wu, Huajie Tan, Chunyuan Li, Jing Yang, Jie Yu, Xiyao Wang, Bin Qin, Yumeng Wang, Zizhen Yan, Ziyong Feng, Ziwei Liu, Bo Li, and Jiankang Deng. Llava-onevision-1.5: Fully open framework for democratized multimodal training. *arXiv preprint arXiv:2509.23661*, 2025. 2, 6, 7
- [4] Chris H. Bahnsen and Thomas B. Moeslund. Rain Removal in Traffic Surveillance: Does it matter? *IEEE TITS*, 2019. 1
- [5] Shuai Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Sibao Song, Kai Dang, Peng Wang, Shijie Wang, Jun Tang, Humen Zhong, Yuanzhi Zhu, Mingkun Yang, Zhaohai Li, Jianqiang Wan, Pengfei Wang, Wei Ding, Zheren Fu, Yiheng Xu, Jiabo Ye, Xi Zhang, Tianbao Xie, Zesen Cheng, Hang Zhang, Zhibo Yang, Haiyang Xu, and Junyang Lin. Qwen2.5-vl technical report. *arXiv preprint arXiv:2502.13923*, 2025. 2, 4, 5, 6, 7, 8
- [6] ByteDance Seed. Doubao-seed-1.6. <https://console.volcengine.com/ark/region:ark-cn-beijing/model/detail?Id=doubao-seed-1-6>, 2025. 2, 6
- [7] ByteDance Seed. Doubao-seed-vision-1.6. <https://console.volcengine.com/ark/region:ark-cn-beijing/model/detail?Id=doubao-seed-1-6-vision>, 2025. 2, 6, 7
- [8] John Canny. A computational approach to edge detection. *IEEE TPAMI*, 1986. 5
- [9] Zhe Cao, Jin Zhang, and Ruiheng Zhang. Irgpt: Understanding real-world infrared image with bi-cross-modal curriculum on large-scale benchmark. In *ICCV*, 2025. 1, 2, 3
- [10] Gheorghe Comanici, Eric Bieber, Mike Schaekermann, Ice Pasupat, Noveen Sachdeva, Inderjit Dhillon, Marcel Blisstein, Ori Ram, Dan Zhang, Evan Rosen, Luke Marris, Sam Petulla, Colin Gaffney, Asaf Aharoni, Nathan Lintz, Tiago Cardal Pais, Henrik Jacobsson, Idan Szepektor, Nanjiang Jiang, Krishna Haridasan, Ahmed Omran, Nikunj Saunshi, Dara Bahri, Gaurav Mishra, Eric Chu, Toby Boyd, Brad Hekman, Aaron Parisi, Chaoyi Zhang, Kornrapph Kawintiranon, Tania Bedrax-Weiss, Oliver Wang, Ya Xu, Ollie Purkiss, Uri Mendlovic, Ilai Deutel, Nam Nguyen, Adam Langley, Flip Korn, Lucia Rossazza, Alexandre Ramé, Sagar Waghmare, Helen Miller, Nathan Byrd, Ashrith Sheshan, Raia Hadsell, Sangnie Bhardwaj, Pawel Janus, Tero Rissa, Dan Horgan, Alvin Abdagic, Lior Belenki, James Allingham, Anima Singh, Theo Guidroz, Srivatsan Srinivasan, Herman Schmit, Kristen Chiafullo, Andre Elisseeff, Nilpa Jha, Prateek Kolhar, Leonard Berrada, Frank Ding, Xiance Si, Shrestha Basu Mallick, Franz Och, Sofia Erell, Eric Ni, Tejasi Latkar, Sherry Yang, Petar Sirkovic, Ziqiang Feng, Robert Leland, Rachel Hornung, Gang Wu, Charles Blundell, Hamidreza Alvari, Po-Sen Huang, Cathy Yip, Sanja Deur, Li Liu, Gabriela Surita, et al. Gemini 2.5: Pushing the frontier with advanced reasoning, multimodality, long context, and next generation agentic capabilities. *arXiv preprint arXiv:2507.06261*, 2025. 1, 2, 5, 6, 7, 8
- [11] Edward J Hu, yelong shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. LoRA: Low-rank adaptation of large language models. In *ICLR*, 2022. 1
- [12] Xinyu Jia, Chuang Zhu, Minzhen Li, Wenqi Tang, and Wenli Zhou. LLVIP: A visible-infrared paired dataset for low-light vision. In *ICCVW*, 2021. 1
- [13] Shixin Jiang, Zerui Chen, Jiafeng Liang, Yanyan Zhao, Ming Liu, and Bing Qin. Infrared-LLaVA: Enhancing understanding of infrared images in multi-modal large language models. In *EMNLP*, 2024. 1, 2, 3
- [14] Kunchang Li, Yali Wang, Yanan He, Yizhuo Li, Yi Wang, Yi Liu, Zun Wang, Jilan Xu, Guo Chen, Ping Luo, Limin Wang, and Yu Qiao. Mvbench: A comprehensive multi-modal video understanding benchmark. In *CVPR*, 2024. 3
- [15] Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. Visual instruction tuning. In *NeurIPS*, 2023. 2
- [16] Yuan Liu, Haodong Duan, Yuanhan Zhang, Bo Li, Songyang Zhang, Wangbo Zhao, Yike Yuan, Jiaqi Wang, Conghui He, Ziwei Liu, Kai Chen, and Dahua Lin. Mmbench: Is your multi-modal model an all-around player? In *ECCV*, 2024. 1, 3, 5
- [17] OpenAI. Thinking with images. <https://openai.com/index/thinking-with-images/>, 2025. 1
- [18] OpenAI, Aaron Hurst, Adam Lerer, Adam P. Goucher, Adam Perelman, Aditya Ramesh, Aidan Clark, AJ Ostrow, Akila Welihinda, Alan Hayes, Alec Radford, Aleksander Mądry, Alex Baker-Whitcomb, Alex Beutel, Alex Borzunov, Alex Carney, Alex Chow, Alex Kirillov, Alex Nichol, Alex Paino, Alex Renzin, Alex Tachard Passos, Alexander Kirillov, Alexi Christakis, Alexis Conneau, Ali Kamali, Allan Jabri, Allison Moyer, Allison Tam, Amadou Crookes, Amin Tootoochian, Amin Tootoonchian, Ananya Kumar, Andrea Vallone, Andrej Karpathy, Andrew Braunstein, Andrew Cann, Andrew Codispoti, Andrew Galu, Andrew Kondrich, Andrew Tulloch, Andrew Mishchenko, Angela Baek, Angela Jiang, Antoine Pelisse, Antonia Woodford, Anuj Gosalia, Arka Dhar, Ashley Pantuliano, Avi Nayak, Avital Oliver, Barret Zoph, Behrooz Ghorbani, Ben Leimberger, Ben Rossen, Ben Sokolowsky, Ben Wang, Benjamin Zweig, Beth Hoover, Blake Samic, Bob McGrew, Bobby Spero, Bogo Gertler, Bowen Cheng, Brad Lightcap, Brandon Walkin, Brendan Quinn, Brian Guarraci, Brian Hsu, Bright Kellogg, Brydon Eastman, Camillo Lugaresi, Carroll Wainwright, Cary Bassin, Cary Hudson, Casey Chu, Chad Nelson, Chak Li, Chan Jun Shern, Channing Conger, Charlotte Barette, Chelsea Voss, Chen Ding, Cheng Lu, Chong Zhang, Chris

- Beaumont, Chris Hallacy, Chris Koch, Christian Gibson, Christina Kim, Christine Choi, Christine McLeavey, Christopher Hesse, Claudia Fischer, Clemens Winter, Coley Czarnecki, Colin Jarvis, Colin Wei, Constantin Koumouzelis, Dane Sherburn, Daniel Kappler, et al. Gpt-4o system card. *arXiv preprint arXiv:2410.21276*, 2024. 1, 2
- [19] Cristina Palmero, Albert Clapés, Chris Bahnsen, Andreas Møgelmoose, Thomas B Moeslund, and Sergio Escalera. Multi-modal rgb–depth–thermal human body segmentation. *IJCV*, 2016. 1
- [20] Atin Pothiraj, Elias Stengel-Eskin, Jaemin Cho, and Mohit Bansal. Capture: Evaluating spatial reasoning in vision language models via occluded object counting. *arXiv preprint arXiv:2504.15485*, 2025. 3
- [21] Team Seedream, Yunpeng Chen, Yu Gao, Lixue Gong, Meng Guo, Qiushan Guo, Zhiyao Guo, Xiaoxia Hou, Weilin Huang, Yixuan Huang, Xiaowen Jian, Huaifeng Kuang, Zhichao Lai, Fanshi Li, Liang Li, Xiaochen Lian, Chao Liao, Liyang Liu, Wei Liu, Yanzuo Lu, Zhengxiong Luo, Tongtong Ou, Guang Shi, Yichun Shi, Shiqi Sun, Yu Tian, Zhi Tian, Peng Wang, Rui Wang, Xun Wang, Ye Wang, Guofeng Wu, Jie Wu, Wenxu Wu, Yonghui Wu, Xin Xia, Xuefeng Xiao, Shuang Xu, Xin Yan, Ceyuan Yang, Jianchao Yang, Zhonghua Zhai, Chenlin Zhang, Heng Zhang, Qi Zhang, Xinyu Zhang, Yuwei Zhang, Shijia Zhao, Wenliang Zhao, and Wenjia Zhu. Seedream 4.0: Toward next-generation multimodal image generation. *arXiv preprint arXiv:2509.20427*, 2025. 2, 5, 8, 1
- [22] Oriane Siméoni, Huy V. Vo, Maximilian Seitzer, Federico Baldassarre, Maxime Oquab, Cijo Jose, Vasil Khalidov, Marc Szafraniec, Seungeun Yi, Michaël Ramamonjisoa, Francisco Massa, Daniel Haziza, Luca Wehrstedt, Jianyuan Wang, Timothée Darcet, Théo Moutakanni, Leonel Sentana, Claire Roberts, Andrea Vedaldi, Jamie Tolan, John Brandt, Camille Couprie, Julien Mairal, Hervé Jégou, Patrick Labatut, and Piotr Bojanowski. Dinov3. *arXiv preprint arXiv:2508.10104*, 2025. 6
- [23] Yiming Sun, Bing Cao, Pengfei Zhu, and Qinghua Hu. Drone-based rgb-infrared cross-modality vehicle detection via uncertainty-aware learning. *IEEE TCSVT*, 2022. 1
- [24] TELEDYNE FLIR Team. FLIR thermal dataset for algorithm training. Online, 2019. 1
- [25] V Team, Wenyi Hong, Wenmeng Yu, Xiaotao Gu, Guo Wang, Guobing Gan, Haomiao Tang, Jiale Cheng, Ji Qi, Junhui Ji, Lihang Pan, Shuaiqi Duan, Weihang Wang, Yan Wang, Yean Cheng, Zehai He, Zhe Su, Zhen Yang, Ziyang Pan, Aohan Zeng, Baoxu Wang, Bin Chen, Boyan Shi, Changyu Pang, Chenhui Zhang, Da Yin, Fan Yang, Guoqing Chen, Jiazheng Xu, Jiale Zhu, Jiali Chen, Jing Chen, Jinhao Chen, Jinghao Lin, Jinjiang Wang, Junjie Chen, Leqi Lei, Letian Gong, Leyi Pan, Mingdao Liu, Mingde Xu, Mingzhi Zhang, Qinkai Zheng, Sheng Yang, Shi Zhong, Shiyu Huang, Shuyuan Zhao, Siyan Xue, Shangqin Tu, Shengbiao Meng, Tianshu Zhang, Tianwei Luo, Tianxiang Hao, Tianyu Tong, Wenkai Li, Wei Jia, Xiao Liu, Xiaohan Zhang, Xin Lyu, Xinyue Fan, Xuancheng Huang, Yanling Wang, Yadong Xue, Yanfeng Wang, Yanzi Wang, Yifan An, Yifan Du, Yiming Shi, Yiheng Huang, Yilin Niu, Yuan Wang, Yuanchang Yue, Yuchen Li, Yutao Zhang, Yuting Wang, Yu Wang, Yuxuan Zhang, Zhao Xue, Zhenyu Hou, Zhengxiao Du, Zihan Wang, Peng Zhang, Debing Liu, Bin Xu, Juanzi Li, Minlie Huang, Yuxiao Dong, and Jie Tang. Glm-4.5v and glm-4.1v-thinking: Towards versatile multimodal reasoning with scalable reinforcement learning. *arXiv preprint arXiv:2507.01006*, 2025. 6, 7, 8
- [26] Huy V. Vo, Vasil Khalidov, Timothée Darcet, Théo Moutakanni, Nikita Smetanin, Marc Szafraniec, Hugo Touvron, camille couprie, Maxime Oquab, Armand Joulin, Herve Jegou, Patrick Labatut, and Piotr Bojanowski. Automatic data curation for self-supervised learning: A clustering-based approach. *TMLR*, 2024. 6
- [27] Weiyun Wang, Zhangwei Gao, Lixin Gu, Hengjun Pu, Long Cui, Xingguang Wei, Zhaoyang Liu, Linglin Jing, Shenglong Ye, Jie Shao, Zhaokai Wang, Zhe Chen, Hongjie Zhang, Ganlin Yang, Haomin Wang, Qi Wei, Jinhui Yin, Wenhao Li, Erfei Cui, Guanzhou Chen, Zichen Ding, Changyao Tian, Zhenyu Wu, Jingjing Xie, Zehao Li, Bowen Yang, Yuchen Duan, Xuehui Wang, Zhi Hou, Haoran Hao, Tianyi Zhang, Songze Li, Xiangyu Zhao, Haodong Duan, Nianchen Deng, Bin Fu, Yinan He, Yi Wang, Conghui He, Botian Shi, Junjun He, Yingdong Xiong, Han Lv, Lijun Wu, Wenqi Shao, Kaipeng Zhang, Huipeng Deng, Biqing Qi, Jiaye Ge, Qipeng Guo, Wenwei Zhang, Songyang Zhang, Maosong Cao, Junyao Lin, Kexian Tang, Jianfei Gao, Haihan Huang, Yuzhe Gu, Chengqi Lyu, Huanze Tang, Rui Wang, Haijun Lv, Wanli Ouyang, Limin Wang, Min Dou, Xizhou Zhu, Tong Lu, Dahua Lin, Jifeng Dai, Weijie Su, Bowen Zhou, Kai Chen, Yu Qiao, Wenhao Wang, and Gen Luo. Internvl3.5: Advancing open-source multimodal models in versatility, reasoning, and efficiency. *arXiv preprint arXiv:2508.18265*, 2025. 2, 6, 7, 8, 1, 4
- [28] Chenfei Wu, Jiahao Li, Jingren Zhou, Junyang Lin, Kaiyuan Gao, Xun Yan, Sheng ming Yin, Shuai Bai, Xiao Xu, Yilei Chen, Yuxiang Chen, Zecheng Tang, Zekai Zhang, Zhengyi Wang, An Yang, Bowen Yu, Chen Cheng, Dayiheng Liu, Deqing Li, Hang Zhang, Hao Meng, Hu Wei, Jingyuan Ni, Kai Chen, Kuan Cao, Liang Peng, Lin Qu, Minggang Wu, Peng Wang, Shuting Yu, Tingkun Wen, Wensen Feng, Xiaoxiao Xu, Yi Wang, Yichang Zhang, Yongqiang Zhu, Yujia Wu, Yuxuan Cai, and Zenan Liu. Qwen-image technical report. *arXiv preprint arXiv:2508.02324*, 2025. 2, 5, 6, 8, 1
- [29] An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, Chujie Zheng, Dayiheng Liu, Fan Zhou, Fei Huang, Feng Hu, Hao Ge, Haoran Wei, Huan Lin, Jialong Tang, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Yang, Jiayi Yang, Jing Zhou, Jingren Zhou, Junyang Lin, Kai Dang, Keqin Bao, Kexin Yang, Le Yu, Lianghao Deng, Mei Li, Mingfeng Xue, Mingze Li, Pei Zhang, Peng Wang, Qin Zhu, Rui Men, Ruize Gao, Shixuan Liu, Shuang Luo, Tianhao Li, Tianyi Tang, Wenbiao Yin, Xingzhang Ren, Xinyu Wang, Xinyu Zhang, Xuancheng Ren, Yang Fan, Yang Su, Yichang Zhang, Yinger Zhang, Yu Wan, Yuqiong Liu, Zekun Wang, Zeyu Cui, Zhenru Zhang, Zhipeng Zhou, and Zihan Qiu. Qwen3 technical report. *arXiv preprint arXiv:2505.09388*, 2025. 1, 2, 5, 6, 7, 8, 4

- [30] Biao Yang, Bin Wen, Boyang Ding, Changyi Liu, Chenglong Chu, Chengru Song, Chongling Rao, Chuan Yi, Da Li, Dunju Zang, Fan Yang, Guorui Zhou, Guowang Zhang, Han Shen, Hao Peng, Haojie Ding, Hao Wang, Haonan Fan, Hengrui Ju, Jiaming Huang, Jiangxia Cao, Jiankang Chen, Jingyun Hua, Kaibing Chen, Kaiyu Jiang, Kaiyu Tang, Kun Gai, Muhao Wei, Qiang Wang, Ruitao Wang, Sen Na, Shengnan Zhang, Siyang Mao, Sui Huang, Tianke Zhang, Tingting Gao, Wei Chen, Wei Yuan, Xiangyu Wu, Xiao Hu, Xingyu Lu, Yi-Fan Zhang, Yiping Yang, Yulong Chen, Zeyi Lu, Zhenhua Wu, Zhixin Ling, Zhuoran Yang, Ziming Li, Di Xu, Haixuan Gao, Hang Li, Jing Wang, Lejian Ren, Qigen Hu, Qianqian Wang, Shiyao Wang, Xinchun Luo, Yan Li, Yuhang Hu, and Zixing Zhang. Kwai keye-v1 1.5 technical report. *arXiv preprint arXiv:2509.01563*, 2025. 6, 7
- [31] Sihan Yang, Runsen Xu, Yiman Xie, Sizhe Yang, Mo Li, Jingli Lin, Chenming Zhu, Xiaochen Chen, Haodong Duan, Xiangyu Yue, Dahua Lin, Tai Wang, and Jiangmiao Pang. Mmsi-bench: A benchmark for multi-image spatial intelligence. *arXiv preprint arXiv:2505.23764*, 2025. 3
- [32] Weihao Yu, Zhengyuan Yang, Linjie Li, Jianfeng Wang, Kevin Lin, Zicheng Liu, Xinchao Wang, and Lijuan Wang. MM-vet: Evaluating large multimodal models for integrated capabilities. In *ICML*, 2024. 3
- [33] Xiang Yue, Yuansheng Ni, Kai Zhang, Tianyu Zheng, Ruoqi Liu, Ge Zhang, Samuel Stevens, Dongfu Jiang, Weiming Ren, Yuxuan Sun, Cong Wei, Botao Yu, Ruibin Yuan, Renliang Sun, Ming Yin, Boyuan Zheng, Zhenzhu Yang, Yibo Liu, Wenhao Huang, Huan Sun, Yu Su, and Wenhua Chen. Mmmu: A massive multi-discipline multimodal understanding and reasoning benchmark for expert agi. In *CVPR*, 2024. 1, 3
- [34] Tao Zhang, Jinyong Wen, Zhen Chen, Kun Ding, Shiming Xiang, and Chunhong Pan. UNIP: Rethinking pre-trained attention patterns for infrared semantic segmentation. In *ICLR*, 2025. 4, 1, 3
- [35] Jinguo Zhu, Weiyun Wang, Zhe Chen, Zhaoyang Liu, Shenglong Ye, Lixin Gu, Hao Tian, Yuchen Duan, Weijie Su, Jie Shao, Zhangwei Gao, Erfei Cui, Xuehui Wang, Yue Cao, Yangzhou Liu, Xingguang Wei, Hongjie Zhang, Haomin Wang, Weiye Xu, Hao Li, Jiahao Wang, Nianchen Deng, Songze Li, Yinan He, Tan Jiang, Jiapeng Luo, Yi Wang, Conghui He, Botian Shi, Xingcheng Zhang, Wenqi Shao, Junjun He, Yingtong Xiong, Wenwen Qu, Peng Sun, Penglong Jiao, Han Lv, Lijun Wu, Kaipeng Zhang, Huipeng Deng, Jiaye Ge, Kai Chen, Limin Wang, Min Dou, Lewei Lu, Xizhou Zhu, Tong Lu, Dahua Lin, Yu Qiao, Jifeng Dai, and Wenhua Wang. Internvl3: Exploring advanced training and test-time recipes for open-source multimodal models. *arXiv preprint arXiv:2504.10479*, 2025. 6, 7