

Align²LLaVA: Cascaded Human and Large Language Model Preference Alignment for Multi-modal Instruction Curation

Hongzhe Huang¹, Jiang Liu¹, Zhewen Yu¹, Li Cai¹,
Dian Jiao¹, Wenqiao Zhang^{1†}, Siliang Tang¹, Juncheng Li^{1†},
Hao Jiang², Haoyuan Li², Yueting Zhuang¹

¹Zhejiang University, ²Alibaba, [†]Corresponding Authors

Abstract

Recent advances in Multi-modal Large Language Models (MLLMs), such as LLaVA-series models, are driven by massive machine-generated instruction-following data tuning. Such automatic instruction collection pipelines, however, inadvertently introduce significant variability in data quality. This paper introduces a novel instruction curation algorithm, derived from two unique perspectives, **human and LLM preference alignment**, to compress this vast corpus of machine-generated multimodal instructions to a compact and high-quality form: (i) *For human preference alignment*, we have collected a machine-generated multimodal instruction dataset and established a comprehensive set of both subjective and objective criteria to guide the data quality assessment critically from human experts. By doing so, a reward model was trained on the annotated dataset to internalize the nuanced human understanding of instruction alignment. (ii) *For LLM preference alignment*, given the instruction selected by the reward model, we propose leveraging the inner LLM used in MLLM to align the writing style of visual instructions with that of the inner LLM itself, resulting in LLM-aligned instruction improvement. Extensive experiments demonstrate that we can **maintain or even improve model performance by compressing synthetic multimodal instructions by up to 90%**. Impressively, by aggressively reducing the training instructions from 158k to 14k (9× smaller), our model consistently outperforms its full-size dataset counterpart across various MLLM benchmarks. Our project is available at <https://github.com/DCDmllm/Align2LLaVA>.

1 Introduction

Recent triumphs in artificial intelligence have been driven notably by the development of large language models (LLMs) (Achiam et al., 2023; Team et al., 2023; Dubey et al., 2024). Inspired by the impressive instruction-following capability of LLMs,

visual instruction tuning (Liu et al., 2023b) has been proposed to extend LLMs into Multimodal LLMs (MLLMs) to perceive and understand visual signals (e.g., images) (Zhang et al., 2022, 2024b; Lin et al., 2025; Zhang et al., 2019; Zhu et al., 2023; Yuan et al., 2024). Due to the scarcity of visual-language instruction-following data, contemporary MLLMs prefer to utilize the data reformation approach (Liu et al., 2023b), which leverages text-only LLMs conditioned on image captions and bounding boxes to create synthetic multi-modal instructions involving visual content. Following the evolution, other works such as ALLaVA (Chen et al., 2024) and VIGC (Wang et al., 2024), have extended this automatic paradigm by utilizing visually-capable LLMs, demonstrating the potential value of utilizing synthetic instruction generation.

Despite their efforts and achievements, the simply synthetic instruction generation paradigm may result in sub-optimal multimodal instruction following capability, which can be summarized as two key aspects: (i) **Inherently Noisy Instructions**: The generated questions and corresponding responses of instructions may be incongruent or incomplete with the visual content, as they are produced by text-only large language models. Regarding question generation, common issues include irrelevance to the image, independence from visual cues, and repetitive queries in multi-turn dialogues. On the other hand, the corresponding responses may suffer from textual hallucinations or ignored visual details, as they are produced by text-only large language models. Figure 1(a) shows an example of the synthetic instructions. Of the two LLM-generated questions, the previous one demonstrates superior quality by incorporating queries relevant to image content, and providing subtle prompts to stimulate reasoning without revealing specific visual information. Conversely, the subsequent question lacks pertinence to the image, as it can-

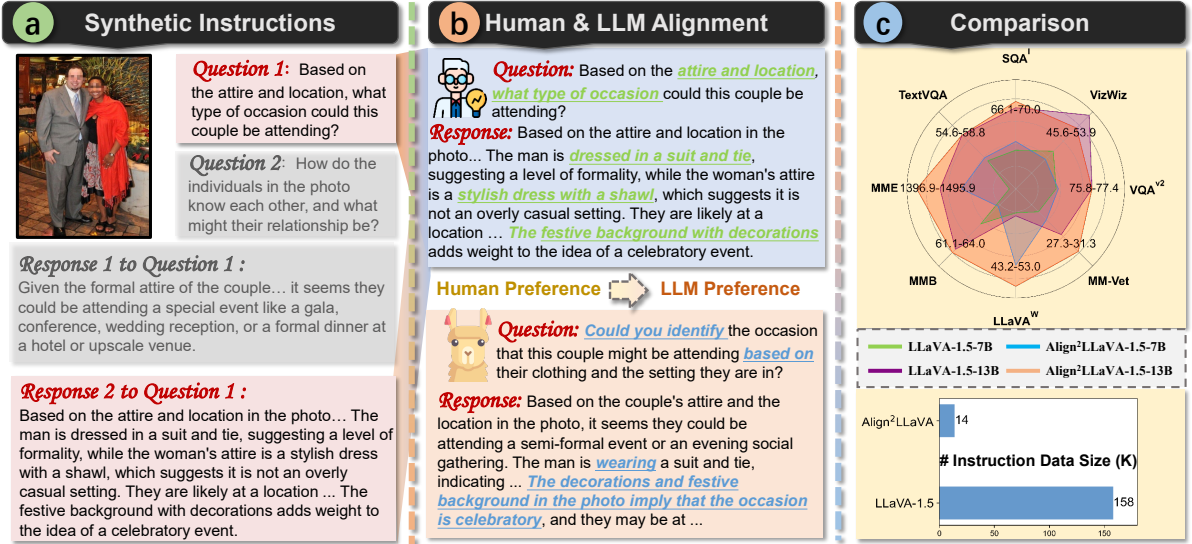


Figure 1: (a) An example of synthetic visual instruction data generated by LLMs. (b) Demonstration of proposed cascaded human and LLM preference alignment. (c) Across 8 benchmarks, our approach achieves comparable or superior performance to LLaVA-1.5 trained on the full dataset (top), using a significantly reduced instruction set (bottom), demonstrating the efficiency of our method.

not be accurately answered based on the visual data provided. Moreover, comparing the two responses to the first question, the first one is notably deficient in visual details, as highlighted in the upper part of Figure 1(b). Such details are crucial for comprehensive and persuasive analysis. These inappropriate questions and responses inevitably compromise the accurately visual-language alignments of the MLLMs. (ii) **Internally Linguistic Gap:** We refer the linguistic characteristic to the specific habits of vocabulary selection, grammar usage and sentence structuring used to express particular semantics. We argue that well-trained LLMs have unique writing style preferences, which are expressed in the output probabilities of candidate tokens when generating new tokens. In Figure 1(b), we present an instance of sentence writing manner gap. A comparison between the original instruction (top) and the LLM-revised instruction (bottom) reveals subtle stylistic variations, including word choice, sentence structure, and the degree of complexity in expressing the same idea. During the visual instruction tuning phase, the writing manner gap forces the LLM to change its original writing style, leading to performance degradation or even catastrophic forgetting. Based on the aforementioned insights, a meaningful optimization goal of synthetic instruction generation is to effectively explore the well-aligned vision-instructions and minimize the linguistic gap, thereby boosting the

multi-modal instruction-following capabilities.

Generally speaking, a prevalent strategy for quality enhancement of textually training data is *data filtering*, e.g., employing pre-trained models to assess answer quality relative to given questions, subsequently retaining only instructions with high-quality answers (Cao et al., 2023; Chen et al., 2023a; Li et al., 2023a). However, this paradigm is not the true savior for multi-modal instructions, mainly due to the current filtering methods being designed for text-only instruction data, i.e., lacking the reflect in depth of vision-language consistency among the image, question, and answer. Notably, these methods are trained on unfiltered data with noise and bias, which may lead to inaccurate assessments. What’s more, current data filtering methods have not considered the linguistic gap between LLMs, handling the potential of synthetic instructions’ potential. Summing up, the premise of better multi-modal instruction filtering is to mitigate the aforementioned issues, pushing us to revisit the filtering solution for synthetic instructions.

To this end, we propose a novel data curation paradigm for synthetic multi-modal instructions, called **Align²LLaVA**. Our bootstrapping philosophy is to progressively align human expert’s and pre-trained LLM’s preference, thereby respectively addressing the inherently noisy instructions and internally linguistic gap: **Step 1, Human Knowledge Alignment**, this step is built upon the concept of

integrating human feedback signals to comprehensively evaluate and filter both question and answer components of visual instructions. In detail, we collect human preference instructions emphasizing distinct quality criteria at question- and response-aligned level given images, *i.e.*, respectively image- k questions and question- l responses datasets. Human annotators independently rank each part based on metrics including image relevance, question comprehensiveness, answer accuracy, etc. By doing so, human-aligned instructions are obtained and thereafter employed to train two separate reward models for evaluation. Finally, utilizing the trained reward models, a two-stage filtering process is then applied to large-scale synthetic instructions, preserving examples that optimally conform to human-established quality standards within both question and answer domains. **Step 2, LLM Characteristic Alignment**, we leverage the inner LLM to align the writing manner of soft-format visual instructions with that of the inner LLM itself, ensuring the original semantics of these instructions remain unchanged. The answer part of soft-format visual instructions is first rewritten by the inner LLM to match its writing manner, and then reviewed by the inner LLM to ensure the alignment of writing manner while preserving the original meaning. If the revised answer is deemed unqualified during the review, the original answer is retained. By combining this two-step aligning process, the proposed synthetic multi-modal instruction dataset is created.

Our experiments empirically validated the efficacy of the proposed data curation method. We applied the pipeline to a 158K synthetic instruction dataset, resulting in a compressed dataset, **Align²LLaVA-Instruct**. Align²LLaVA-Instruct consists of 30% of the original questions and 30% of the remaining answers, which is 9% of the full size. We subsequently fine-tuned the LLaVA-1.5 (Liu et al., 2023a) model using Align²LLaVA-Instruct combined with the caption and short VQA parts of the LLaVA-1.5-665K (Liu et al., 2023a) dataset, achieving comparable or even superior performance to the same architecture trained on the full instruction data, across a suite of 8 benchmarks. These benchmarks encompassed a diverse range of tasks, spanning traditional academic challenges and comprehensive instruction-following scenarios.

In summary, our contributions are threefold:

- We introduce a novel pipeline, **Align²LLaVA**, to curate high-quality synthetic visual instruc-

tion data by sequentially aligning with human preferences and LLM characteristics.

- Utilizing the curation pipeline, we present an aligned synthetic visual instruction dataset, **Align²LLaVA-Instruct**, distinguished by its high quality and compact size.
- Through extensive experiments, we demonstrate the effectiveness of the proposed curation pipeline and dataset for MLLM training.

2 Related Work

2.1 Learning from Human Knowledge

The integration of human knowledge is becoming increasingly pivotal in aligning model behavior with human intent (Lee et al., 2023; Nakano et al., 2021; Wu et al., 2023; Xu et al., 2024; Zhang et al., 2024a; Yu et al., 2024b; Zhang et al., 2023). A crucial aspect is training a reward model (MacGlashan et al., 2017) to directly capture human preferences for the outputs generated by the model. Recently, researchers have proposed to use reinforcement learning (Schulman et al., 2017) to fine-tune language models (Ouyang et al., 2022) and diffusion models (Lee et al., 2023; Xu et al., 2024; Zhang et al., 2024a) utilizing the signals from the reward model. Additionally, there are works that leverage human knowledge to directly enhance the quality of datasets (Yu et al., 2024b; Zhang et al., 2023). These works integrate human knowledge with image-text data to mitigate visual and textual biases present in large-scale image-text datasets.

2.2 Instruction Data Selection

Recent researches show that instruction tuning could enable large language models (LLMs) to be tailored to specific domains, tasks, or applications by providing explicit instructions or guidelines (Wei et al., 2021). In order to enhance the instruction-following abilities of LLMs, previous works mainly focus on increasing the data sizes through various strategies (Honovich et al., 2022; Wang et al., 2022; Taori et al., 2023; Köpf et al., 2024). However, LIMA (Zhou et al., 2024) illustrates that even a small number of constructed high-quality instructions could empower the model with a powerful instruction-following capability.

Inspired by LIMA (Zhou et al., 2024), Instruction Mining (Cao et al., 2023) adopts a linear quality rule and bag of indicators to evaluate the quality of instruction-following data. AlpaGasus (Chen

et al., 2023c) directly leverages an external LLM (ChatGPT) to score each instruction and then selects 9K Alpaca data with a threshold. While their model surpasses the performance of the official Alpaca model which is trained on the complete dataset, this approach may neglect the intrinsic abilities of the base model, relying excessively on external models. Li et al. (2023a) present a self-guided approach for LLMs to independently identify and choose relevant instruction pairs from extensive open-source datasets and in their approach they introduce an Instruction-Following Difficulty (IFD) metric as a tool to identify gaps in a model’s responses versus its autonomous generation capability. DataTailor (Yu et al., 2024a) focuses on instruction selection through three criteria: informativeness, uniqueness, and representativeness.

3 Method

Our visual instruction curation pipeline comprises three sequential steps, as depicted in Figure 2. Initially, two distinct datasets, enriched with human preferences, are constructed focusing on the question and answer components of instructions, respectively (Section 3.1). Leveraging these human-ranked datasets, two separate reward models are trained to align with human values, and subsequently serve as human-like evaluators to filter large-scale synthetic visual instruction data through a two-stage process (Section 3.2). Finally, an inner LLM is employed to refine the selected instructions in a rewrite-review manner, ensuring alignment with the LLM’s writing style while preserving the original semantic content (Section 3.3).

3.1 Human Preferences Collection

The initial step of our pipeline involves the curation of two distinct visual instruction datasets, each accompanied by corresponding human preference rankings. As depicted in Step 1 of Figure 2, these datasets prioritize either the question or answer component of the instructions. The data collection process unfolds in two sequential phases: *visual instruction generation* and *human preference annotation*. The anticipated outcome is a question dataset ranked by comprehensiveness and visual relevance, as well as an answer dataset evaluated based on accuracy and answer alignment with the provided image and question.

3.1.1 Visual Instruction Generation

To initiate our pipeline, a seed dataset of 960 images is randomly selected from the LLaVA-instruct-158K (Liu et al., 2023b) dataset. Each image is paired with exactly one group of instructions generated by GPT-4 (Achiam et al., 2023). Following LLaVA, these instructions are classified into three categories: *detail description*, *complex reasoning*, and *conversation*, where only the *conversation* samples incorporate multiple question-answer pairs. We also ignore the question part of *detail description* instances, because they are randomly sampled from a fixed, limited set of pre-defined options.

Leveraging the seed dataset, we construct question and answer datasets with diverse instruction variations. To generate the question dataset, LLMs are prompted with images and tasked with producing multiple related queries, and images within the *detail description* part are excluded from this dataset. For the answer dataset, LLMs are provided with both the original image and seed question to produce alternative responses. To enhance data diversity, both GPT-4V and text-only GPT-4 (following LLaVA’s approach of incorporating MSCOCO (Lin et al., 2015) bounding boxes and captions) are employed. Each model generates three responses, with low-quality outputs (e.g., refusals with initial parts like "I’m sorry...") discarded. The seed questions or answers are also included in the datasets.

We finally collected a question dataset of 800 samples and an answer dataset of 960 samples. Formally, each seed sample is represented as (I, Q_0, A_0) , where I is the image, Q_0 is the seed question, and A_0 is the seed answer. Question dataset samples are denoted as $(I, Q_0, Q_1, \dots, Q_k)$, and answer dataset samples as $(I, Q_0, A_0, A_1, \dots, A_l)$, with $k, l \in [3, 6]$ representing the number of generated questions or answers, respectively.

3.1.2 Human Preference Annotation

To evaluate instruction quality, we employ 7 CS students as labelers. They are tasked with ranking each of the question and answer datasets according to distinct criteria, involving correctness, diversity, image-question-answer relevance, etc. Detailed annotation guidelines are provided in the appendix. In general, these criteria establish a multi-faceted evaluation framework that can assess visual instructions holistically. Moreover, to guarantee annotation accuracy and consistency, we have employed

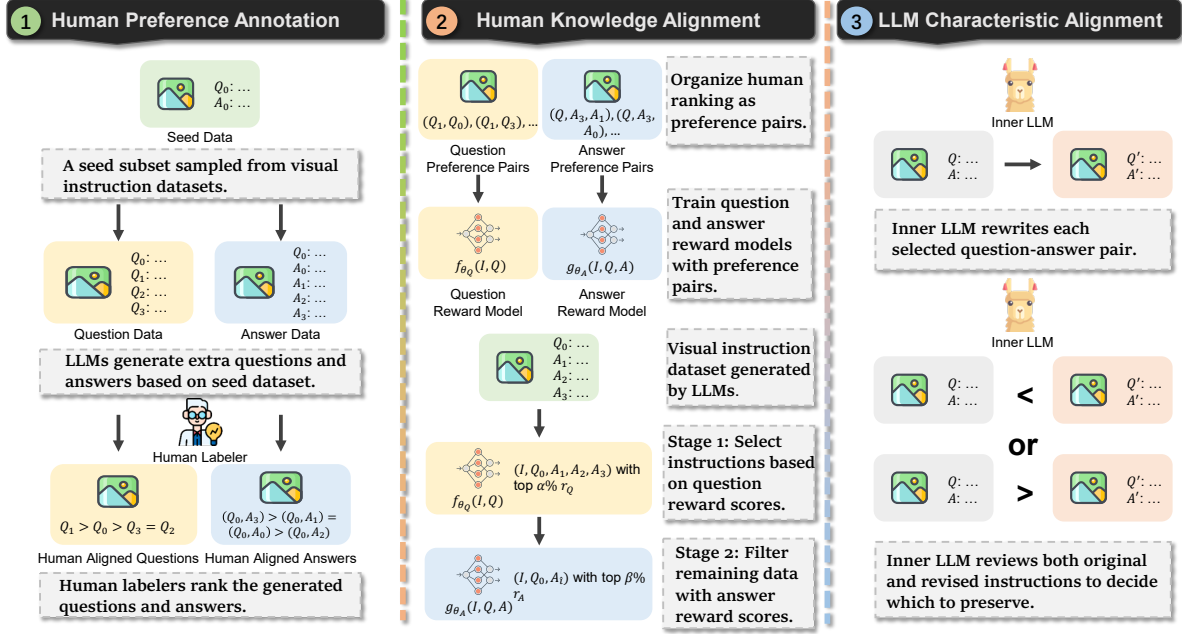


Figure 2: An overview of our data curation pipeline incorporating human knowledge and LLM characteristics. The process comprises three sequential steps: (1) Human preference data is curated through manual annotation on LLM generated questions and answers. (2) Two reward models are trained to align with human values, and subsequently utilized for large-scale data filtration. (3) An inner LLM is employed to rewrite and review the selected instructions.

a multi-annotator approach. Each data sample is independently annotated by one annotator and subsequently reviewed by two or three additional annotators.

3.2 Human Knowledge Alignment

To align large-scale visual instruction data with human knowledge, we employ a human-preference-based filtration process utilizing curated datasets enriched with human feedback. Initially, two distinct reward models are trained separately on question and answer datasets to learn human preferences for these respective components. Subsequently, these trained reward models function as human-like evaluators to filter large-scale synthetic dataset across two sequential stages.

3.2.1 Reward Model Training

To provide a multi-dimensional assessment on visual instructions, we propose to train two distinct reward models on the human-ranked datasets, dedicated to evaluating the question and answer parts, respectively. By aligning these models with human preferences, we aim to cultivate a in-depth understanding of complex semantics and subtle nuances within instructions.

Following InstructGPT (Ouyang et al., 2022),

we structure human preferences as pairwise comparisons, and adopt a pairwise ranking loss for model training. Specifically, within the question dataset, each pair of questions associated with the same image are compared, resulting in a preference pair indicated by (I, Q_i, Q_j) if Q_i is superior to Q_j . Similarly, the answer dataset is organized as preference pairs (I, Q_0, A_i, A_j) , where A_i is preferred over A_j within the context of image I and seed question Q_0 . This yields at most C_k^2 question preference pairs and C_l^2 answer preference pairs for each image.

Subsequently, two reward models are independently trained on question and answer preference pairs, with loss functions defined as:

$$loss_Q(\theta_Q) = -\mathbb{E}_{(I, Q_i, Q_j) \sim \mathcal{D}_Q^{\text{HF}}} [\log(\sigma(f_{\theta_Q}(I, Q_i) - f_{\theta_Q}(I, Q_j)))], \quad (1)$$

$$loss_A(\theta_A) = -\mathbb{E}_{(I, Q_0, A_i, A_j) \sim \mathcal{D}_A^{\text{HF}}} [\log(\sigma(g_{\theta_A}(I, Q_0, A_i) - g_{\theta_A}(I, Q_0, A_j)))], \quad (2)$$

where f is the question reward model parameterized by θ_Q that takes an image I and a question Q as input, and g is the answer reward model parameterized by θ_A , provided with an extra answer A . Both $f_{\theta_Q}(I, Q)$ and $g_{\theta_A}(I, Q, A)$ are scalar scores for given instructions. σ is the sigmoid function,

and $\mathcal{D}_Q^{\text{HF}}$, $\mathcal{D}_A^{\text{HF}}$ are the question and answer dataset, respectively.

3.2.2 Instruction Data Filtration

Leveraging the reward models well-aligned with human preferences, we implement a two-stage filtration on large-scale datasets, acquiring a substantially smaller yet high-quality dataset. Initially, a visual instruction dataset is synthesized, with each image accompanied by one question and three answers, formally represented as $\mathcal{D} = \{(I^n, Q_0^n, A_1^n, A_2^n, A_3^n)\}_{n \in N}$, where $N = \{1, \dots, n_0\}$ is the index set of the full dataset.

In the first filtration stage, we rank all instances in \mathcal{D} based on their question quality using the question reward model f_{θ_Q} , defined as $r_Q^n = f_{\theta_Q}(I^n, Q_0^n)$. Retaining the top $\alpha\%$ of instances yields dataset $\mathcal{D}' = \{(I^{n'}, Q_0^{n'}, A_1^{n'}, A_2^{n'}, A_3^{n'})\}_{n' \in N'}$, where N' is a subset of N , with size $|N'| = \alpha\% \cdot |N|$. The second stage concentrates on evaluating answer quality. For each image in \mathcal{D}' , the answer reward model g_{θ_A} assigns scores to the three answer options, computed as $r_{A,i}^{n'} = g_{\theta_A}(I^{n'}, Q_0^{n'}, A_i^{n'})$ for $i \in \{1, 2, 3\}$. The answer with the highest score, denoted as $A_i^{n'}$, is selected for each image. Subsequently, these selected image-question-answer triplets undergo a ranking process analogous to the first stage, where the top $\beta\%$ instances based on answer scores are preserved, leading to the final dataset $\mathcal{D}'' = \{(I^{n''}, Q_0^{n''}, A_i^{n''})\}_{n'' \in N''}$, where $N'' \subset N'$ and $|N''| = \beta\% \cdot |N'|$. A notable exception is the *detailed description* instructions, which bypass the initial question-based filtration. To maintain consistency in the overall data scale, a direct sampling rate of $\alpha\% \cdot \beta\%$ is applied to this category in the second stage. This two-stage process achieves a sampling rate of $\alpha\% \cdot \beta\%$ relative to the full dataset, effectively retaining instructions best aligned with human preferences in both question and answer domains.

3.3 LLM Characteristic Alignment

3.3.1 Inner LLM Rewrite

To harmonize instructions with the writing style of the inner LLM, we finally introduce an LLM alignment step. This alignment comprises two sequential phases: rewriting and review. During the rewriting phase, the inner LLM is tasked with modifying both question and answer components to conform to its characteristic writing style while maintaining semantic equivalence. Accompanying each

revised instruction is an explanatory note detailing the modifications, thereby enhancing the process's robustness and interpretability. Crucially, the LLM is instructed to maintain the original phrasing if it already adheres to the LLM's stylistic expectations, mitigating the risk of compromising high-quality instructions.

3.3.2 Inner LLM Review

In the subsequent review phase, both the original and revised instructions are presented to the LLM. The model is prompted to determine whether the revised content more closely aligns with its writing style while satisfying the semantics preserving demand. This serves as an extra stage of quality assurance. To encourage logical decision-making, the LLM is also required to provide a rationale for its selection. By refining instruction formats to better suit the target MLLM, this alignment process ultimately enhances training efficiency.

4 Experiments

4.1 Setting and Benchmark

4.1.1 Experimental Setting

To obtain the synthetic visual instruction set, we first generated instructions for 158K images from the MSCOCO (Lin et al., 2015) dataset, using CogVLM-17B (Wang et al., 2023) with 4-bit quantization. Our curation pipeline was applied to this synthetic data, with 30% sampling rate for both questions and answers, resulting in a reduced dataset, Align²LLaVA-Instruct, constituting 9% of the original samples. Combined with the image-caption and short VQA data from the LLaVA-1.5-665K (Liu et al., 2023a) dataset, Align²LLaVA-Instruct was subsequently used to fine-tune the LLaVA-1.5 (Liu et al., 2023a) model, initialized with weights from the previous stage pre-trained on caption datasets. We compared the resulting model against the baseline of LLaVA-1.5 fine-tuned on the entire synthetic instruction dataset with caption and VQA data, as well as other state-of-the-art methods.

4.1.2 Benchmarks

This evaluation suite comprises academic VQA benchmarks, including VQAv2 (Goyal et al., 2017), VizWiz (Gurari et al., 2018), ScienceQA (IMG) (Lu et al., 2022), and TextVQA (Singh et al., 2019). Additionally, to measure instruction-following abilities, we employ MME-Perception (Fu et al., 2023), MMBench (Liu et al.,

Method	LLM	SMI	VQA Datasets				Instruction-Following Benchmarks			
			VQA ^{v2}	VizWiz	SQA ^I	VQA ^T	MME	MMB	LLaVA ^W	MM-Vet
InstructBLIP (Dai et al., 2023)	Vicuna-7B	158K	76.5	34.5	60.5	50.1	1251.8	36.0	56.2	26.2
IDEFICS-9B (Laurençon et al., 2023)	LLama-7B	-	50.9	35.5	56.7	25.9	764.0	48.2	58.2	27.4
Qwen-VL (Bai et al., 2023)	Qwen-7B	-	78.8	35.2	67.1	63.8	1185.4	38.2	37.8	31.6
LLaVA1.5 [†] (Liu et al., 2023a)	Vicuna-7B	158K	75.8	47.1	66.1	54.6	1396.9	62.2	43.2	27.3
Align²LLaVA[†] (Ours)	Vicuna-7B	14K	75.9	45.6	66.8	55.1	1448.0	61.1	50.2	28.1
Improvement. \uparrow	-	-	+0.1	-	+0.7	+0.5	+51.1	-	+7.0	+0.8

InstructBLIP (Dai et al., 2023)	Vicuna-13B	158K	74.8	33.4	63.1	50.7	1212.8	34.2	50.1	25.6
Shikra (Chen et al., 2023b)	Vicuna-13B	156K	77.4	-	-	-	-	58.8	-	-
LLaVA1.5 [†] (158K) (Liu et al., 2023a)	Vicuna-13B	158K	77.4	53.9	69.5	58.8	1471.1	63.8	43.6	29.6
Align²LLaVA[†] (14K) (Ours)	Vicuna-13B	14K	77.4	51.5	70.0	58.8	1495.9	64.0	53.0	31.3
Improvement. \uparrow	-	-	-	-	+0.5	-	+24.8	+0.2	+9.4	+1.7

Table 1: Comparison with Baselines on 8 benchmarks. For fair comparison, we refer [†] to the model trained on our automatically synthesized instructions. SMI indicates the number of synthetic multimodal instruction samples.

2023c), LLaVA-Bench-in-the-Wild (Liu et al., 2023b), and MM-Vet (Yu et al., 2023). These benchmarks collectively provide a robust assessment of multimodal understanding, open-ended conversation, and other relevant skills.

4.2 Overall Performance

Table 1 presents the main experiment results. Our method outperforms the baselines on the majority of benchmarks while achieving highly competitive results on the remaining ones. A key finding is the significant reduction in synthetic instruction data required for our model to achieve such strong performance. Specifically, our model was fine-tuned on only 9% of the synthetic instruction data compared to the full-data baseline using the same model architecture. This outcome strongly suggests that our proposed alignment pipeline not only effectively compresses data size, but also enhances data quality and training efficiency.

4.3 Ablation Studies

4.3.1 Ablation on Human Alignment

To investigate the impact of human knowledge alignment, we conducted an ablation study comparing our human-preference-based data filtration method to a random sampling baseline. The results, presented in Table 2, demonstrate the consistent superiority of our proposed method. Notably, the performance gains are particularly pronounced on instruction-following benchmarks. For instance, on MME (Fu et al., 2023), we observe a significant improvement from 1405.3 to 1448.0, while on LLaVA-Bench-In-the-Wild (Liu et al., 2023b), the score rises from 41.6 to 50.2. These substantial gains strongly suggest that aligning data with human preferences is a critical factor in enhanc-

ing the model’s capacity to generate responses that closely correspond to queries and visual context. This alignment ultimately contributes to the production of human-aligned outputs that effectively fulfill user intents.

4.3.2 Ablation on LLM Alignment

We additionally performed an ablation study to investigate the individual contributions of the rewriting and review stages in our LLM alignment process. Specifically, we sequentially removed each stage to assess its impact on model performance. As detailed in Table 3, the combined application of LLM rewriting and review consistently outperforms the baseline method with solely human-aligned filtration. However, omitting the review stage and directly adopting refined instructions yields unstable results. While achieving a peak accuracy of 67.4 on ScienceQA (IMG) (Lu et al., 2022), a substantial performance degradation from 1431.3 to 1408.8 is also observed on the MME (Fu et al., 2023) benchmark. These results indicate that LLM-rewritten instructions are not reliably improved and can potentially compromise the quality of originally superior data. Consequently, the subsequent review stage is essential for identifying and rectifying such issues, thereby enhancing the overall stability and effectiveness of our LLM alignment approach.

4.4 In-Depth Analysis

We also conducted a series of experiments to further analyze the effectiveness of Align²LLaVA.

4.4.1 Object Hallucination Evaluation

We adopted the POPE (Li et al., 2023b) benchmark to evaluate the models’ propensity for hallucination, which utilizes a polling-based query method

Data Selection	Benchmarks			
	VizWiz	SQA ¹	MME	LLaVA ^W
Random (14K)	44.5	66.5	1405.3	41.6
Ours (14K)	45.6	66.8	1448.0	50.2

Table 2: Comparison between our human-preference-aligned filtration and random compression strategy.

Method	Stages			Benchmarks			
	Human	Rewrite	Review	SQA ¹	VQA ^T	MME	LLaVA ^W
Non-LLM-Aligned	✓			66.3	54.8	1431.3	42.9
LLM Rewrite	✓	✓		67.4	54.9	1408.8	49.0
Ours	✓	✓	✓	66.8	55.1	1448.0	50.2

Table 3: Abalation results of involving different stages in LLM alignment.

involving yes/no questions. We report accuracy and F1-scores as primary performance metrics. The proportion of positive responses, denoted as "Yes" ratio, is also included as a reference for analyzing model behaviors. The results, presented in Table 4, indicate a consistent superiority of our method over the baseline across all three experimental settings. Notably, despite being trained on only 9% of the original synthetic instructions, our model demonstrates a slight reduction in hallucination. The results underscores the effectiveness of our pipeline in identifying and eliminating factual inaccuracies through data alignment, which is a critical aspect of reliable LLM performance.

4.4.2 Human Evaluation

To further assess model performance, a human evaluation was conducted on two sets of 50 images randomly sampled from each of the MSCOCO (Lin et al., 2015) and Flickr30K (Young et al., 2014) datasets. For each image, GPT-4 was prompted to generate questions, which were subsequently refined manually. Align²LLaVA model and the baseline LLaVA-1.5 generated responses to these questions, and 7 human evaluators independently compare the quality of the generated responses to each question. As illustrated in Figure 3, our method consistently surpasses the baseline model in terms of human preference. These results provide empirical evidence that our approach effectively aligns the model with human expectations, enhancing the model’s capacity of producing human-satisfying outputs.

4.4.3 Transfer to Different LLMs

We have also explored the transferability of our curated instructions to different LLMs. Beyond the LLaVA-1.5 based on Vicuna-7B, we fine-tuned Qwen-VL-7B (Bai et al., 2023) under identical

Method	Adversarial			Popular			Random		
	Acc	F1	Yes	Acc	F1	Yes	Acc	F1	Yes
LLaVA1.5 [†] (Liu et al., 2023a)	84.1	82.7	42.2	86.1	84.6	42.2	86.7	85.5	40.4
Align ² LLaVA [†] (Ours)	84.5	83.4	43.4	86.6	85.3	41.3	87.1	86.1	41.6

Table 4: Zero-shot object hallucination evaluation results on POPE benchmark.

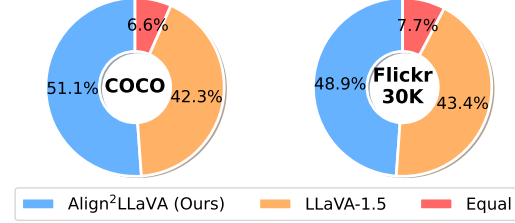


Figure 3: Human evaluation of LLaVA and Align²LLaVA.

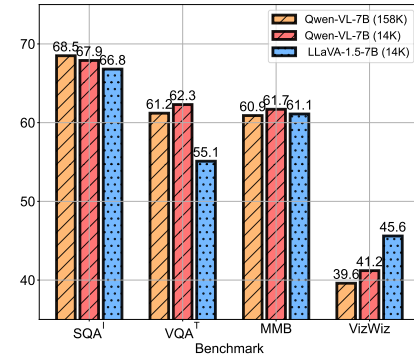


Figure 4: Performance transferring to different LLMs.

settings of Section 4.1. As detailed in Figure 4, on Qwen-7B models, our method achieves comparable performance compared to the counterpart tuned on full 158K synthetic instructions, similar to the findings of fine-tuning Vicuna-7B models. This outcome suggests the generalizability of our Align²LLaVA in across various LLM architectures.

5 Conclusion

This paper contributes to the advancement of high-quality visual instruction synthesis by introducing a novel filtration pipeline that significantly reduces dataset size. Our pipeline sequentially refines instruction quality through alignment with human preferences and LLM writing style. By applying this pipeline to synthetic data, we demonstrate that the resulting substantially smaller dataset yields performance on par with the full instruction dataset when fine-tuning MLLMs. Our results highlight the potential for improving instruction quality via human-in-the-loop refinement and offer a path towards more efficient MLLM training.

Acknowledgements

This work has been supported in part by the NSFC (No. 62402426, 62272411), the Key Research and Development Projects in Zhejiang Province (No. 2025C01128, 2024C01106, 2025C01030, 2025C02156), Ningbo Yongjiang Talent Introduction Programme (2023A-400-G), Zhejiang University Education Foundation Qizhen Scholar Foundation.

Limitations

Our approach, while effective, is subject to certain limitations. Firstly, despite the implementation of unified criteria and rigorous reviews to maintain annotation consistency, inherent subjectivity in human judgment can still introduce discrepancies in human preference data, particularly when multiple annotators are involved. Secondly, while the required amount of human preference data is relatively modest, human annotation remains an essential step in our pipeline. While scaling the synthetic dataset within the current domain does not incur additional annotation costs, extending this method to new domains necessitates the annotation of human preference data in the target domain to enhance the accuracy of the reward models. This potential increase in labor costs associated with adapting our method to novel scenarios should be carefully considered. Furthermore, the acquisition of human preference data in real-world applications may raise privacy concerns. It is imperative to identify and address these privacy issues to ensure the ethical and responsible application of our method.

References

- Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.
- Jinze Bai, Shuai Bai, Shusheng Yang, Shijie Wang, Sinan Tan, Peng Wang, Junyang Lin, Chang Zhou, and Jingren Zhou. 2023. Qwen-vl: A frontier large vision-language model with versatile abilities. *arXiv preprint arXiv:2308.12966*.
- Yihan Cao, Yanbin Kang, and Lichao Sun. 2023. Instruction mining: High-quality instruction data selection for large language models. *arXiv preprint arXiv:2307.06290*.
- Guiming Hardy Chen, Shunian Chen, Ruifei Zhang, Junying Chen, Xiangbo Wu, Zhiyi Zhang, Zhihong Chen, Jianquan Li, Xiang Wan, and Benyou Wang. 2024. Allava: Harnessing gpt4v-synthesized data for a lite vision-language model. *arXiv preprint arXiv:2402.11684*.
- Hao Chen, Yiming Zhang, Qi Zhang, Hantao Yang, Xiaomeng Hu, Xuetao Ma, Yifan Yanggong, and Junbo Zhao. 2023a. [Maybe only 0.5% data is needed: A preliminary exploration of low training data instruction tuning](#). *Preprint*, arXiv:2305.09246.
- Keqin Chen, Zhao Zhang, Weili Zeng, Richong Zhang, Feng Zhu, and Rui Zhao. 2023b. Shikra: Unleashing multimodal llm’s referential dialogue magic. *arXiv preprint arXiv:2306.15195*.
- Lichang Chen, Shiyang Li, Jun Yan, Hai Wang, Kalpa Gunaratna, Vikas Yadav, Zheng Tang, Vijay Srinivasan, Tianyi Zhou, Heng Huang, et al. 2023c. Alpaga: Training a better alpaca with fewer data. *arXiv preprint arXiv:2307.08701*.
- W Dai, J Li, D Li, AMH Tiong, J Zhao, W Wang, B Li, P Fung, and S Hoi. 2023. Instructblip: towards general-purpose vision-language models with instruction tuning. *arxiv. Preprint posted online on June, 15:2023*.
- Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, et al. 2024. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*.
- Chaoyou Fu, Peixian Chen, Yunhang Shen, Yulei Qin, Mengdan Zhang, Xu Lin, Jinrui Yang, Xiawu Zheng, Ke Li, Xing Sun, et al. 2023. Mme: A comprehensive evaluation benchmark for multimodal large language models. *arXiv preprint arXiv:2306.13394*.
- Yash Goyal, Tejas Khot, Douglas Summers-Stay, Dhruv Batra, and Devi Parikh. 2017. Making the v in vqa matter: Elevating the role of image understanding in visual question answering. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 6904–6913.
- Danna Gurari, Qing Li, Abigale J Stangl, Anhong Guo, Chi Lin, Kristen Grauman, Jiebo Luo, and Jeffrey P Bigham. 2018. Vizwiz grand challenge: Answering visual questions from blind people. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3608–3617.
- Or Honovich, Thomas Scialom, Omer Levy, and Timo Schick. 2022. Unnatural instructions: Tuning language models with (almost) no human labor. *arXiv preprint arXiv:2212.09689*.
- Andreas Köpf, Yannic Kilcher, Dimitri von Rütte, Sotiris Anagnostidis, Zhi Rui Tam, Keith Stevens, Abdullah Barhoum, Duc Nguyen, Oliver Stanley, Richárd Nagyfi, et al. 2024. Openassistant conversations-democratizing large language model alignment. *Advances in Neural Information Processing Systems*, 36.

- Hugo Laurençon, Lucile Saulnier, Léo Tronchon, Stas Bekman, Amanpreet Singh, Anton Lozhkov, Thomas Wang, Siddharth Karamcheti, Alexander M. Rush, Douwe Kiela, Matthieu Cord, and Victor Sanh. 2023. [Obelics: An open web-scale filtered dataset of interleaved image-text documents](#). *Preprint*, arXiv:2306.16527.
- Kimin Lee, Hao Liu, Moonkyung Ryu, Olivia Watkins, Yuqing Du, Craig Boutilier, Pieter Abbeel, Mohammad Ghavamzadeh, and Shixiang Shane Gu. 2023. Aligning text-to-image models using human feedback. *arXiv preprint arXiv:2302.12192*.
- Ming Li, Yong Zhang, Zhitao Li, Jiu Hai Chen, Lichang Chen, Ning Cheng, Jianzong Wang, Tianyi Zhou, and Jing Xiao. 2023a. From quantity to quality: Boosting llm performance with self-guided data selection for instruction tuning. *arXiv preprint arXiv:2308.12032*.
- Yifan Li, Yifan Du, Kun Zhou, Jinpeng Wang, Wayne Xin Zhao, and Ji-Rong Wen. 2023b. Evaluating object hallucination in large vision-language models. *arXiv preprint arXiv:2305.10355*.
- Tianwei Lin, Wenqiao Zhang, Sijing Li, Yuqian Yuan, Binhe Yu, Haoyuan Li, Wangui He, Hao Jiang, Mengze Li, Xiaohui Song, et al. 2025. Healthgpt: A medical large vision-language model for unifying comprehension and generation via heterogeneous knowledge adaptation. *arXiv preprint arXiv:2502.09838*.
- Tsung-Yi Lin, Michael Maire, Serge Belongie, Lubomir Bourdev, Ross Girshick, James Hays, Pietro Perona, Deva Ramanan, C. Lawrence Zitnick, and Piotr Dollár. 2015. [Microsoft coco: Common objects in context](#). *Preprint*, arXiv:1405.0312.
- Haotian Liu, Chunyuan Li, Yuheng Li, and Yong Jae Lee. 2023a. Improved baselines with visual instruction tuning. *arXiv preprint arXiv:2310.03744*.
- Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. 2023b. Visual instruction tuning. *arXiv preprint arXiv:2304.08485*.
- Yuan Liu, Haodong Duan, Yuanhan Zhang, Bo Li, Songyang Zhang, Wangbo Zhao, Yike Yuan, Jiaqi Wang, Conghui He, Ziwei Liu, et al. 2023c. Mmbench: Is your multi-modal model an all-around player? *arXiv preprint arXiv:2307.06281*.
- Pan Lu, Swaroop Mishra, Tanglin Xia, Liang Qiu, Kai-Wei Chang, Song-Chun Zhu, Oyvind Tafjord, Peter Clark, and Ashwin Kalyan. 2022. Learn to explain: Multimodal reasoning via thought chains for science question answering. *Advances in Neural Information Processing Systems*, 35:2507–2521.
- James MacGlashan, Mark K Ho, Robert Loftin, Bei Peng, Guan Wang, David L Roberts, Matthew E Taylor, and Michael L Littman. 2017. Interactive learning from policy-dependent human feedback. In *International conference on machine learning*, pages 2285–2294. PMLR.
- Reiichiro Nakano, Jacob Hilton, Suchir Balaji, Jeff Wu, Long Ouyang, Christina Kim, Christopher Hesse, Shantanu Jain, Vineet Kosaraju, William Saunders, et al. 2021. Webgpt: Browser-assisted question-answering with human feedback, 2021. URL <https://arxiv.org/abs/2112.09332>.
- Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, John Schulman, Jacob Hilton, Fraser Kelton, Luke Miller, Maddie Simens, Amanda Askell, Peter Welinder, Paul F Christiano, Jan Leike, and Ryan Lowe. 2022. [Training language models to follow instructions with human feedback](#). In *Advances in Neural Information Processing Systems*, volume 35, pages 27730–27744. Curran Associates, Inc.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. 2021. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR.
- John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. 2017. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*.
- Amanpreet Singh, Vivek Natarajan, Meet Shah, Yu Jiang, Xinlei Chen, Dhruv Batra, Devi Parikh, and Marcus Rohrbach. 2019. Towards vqa models that can read. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 8317–8326.
- Rohan Taori, Ishaan Gulrajani, Tianyi Zhang, Yann Dubois, Xuechen Li, Carlos Guestrin, Percy Liang, and Tatsunori B Hashimoto. 2023. Stanford alpaca: An instruction-following llama model.
- Gemini Team, Rohan Anil, Sebastian Borgeaud, Yonghui Wu, Jean-Baptiste Alayrac, Jiahui Yu, Radu Soricut, Johan Schalkwyk, Andrew M Dai, Anja Hauth, et al. 2023. Gemini: a family of highly capable multimodal models. *arXiv preprint arXiv:2312.11805*.
- Bin Wang, Fan Wu, Xiao Han, Jiahui Peng, Huaping Zhong, Pan Zhang, Xiaoyi Dong, Weijia Li, Wei Li, Jiaqi Wang, and Conghui He. 2024. [VIGC: Visual Instruction Generation and Correction](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 38(6):5309–5317.
- Wei Han Wang, Qingsong Lv, Wenmeng Yu, Wenyi Hong, Ji Qi, Yan Wang, Junhui Ji, Zhuoyi Yang, Lei Zhao, Xixuan Song, et al. 2023. Cogvlm: Visual expert for pretrained language models. *arXiv preprint arXiv:2311.03079*.
- Yizhong Wang, Yeganeh Kordi, Swaroop Mishra, Alisa Liu, Noah A Smith, Daniel Khashabi, and Hananeh Hajishirzi. 2022. Self-instruct: Aligning language model with self generated instructions. *arXiv preprint arXiv:2212.10560*.

- Jason Wei, Maarten Bosma, Vincent Y Zhao, Kelvin Guu, Adams Wei Yu, Brian Lester, Nan Du, Andrew M Dai, and Quoc V Le. 2021. Finetuned language models are zero-shot learners. *arXiv preprint arXiv:2109.01652*.
- Xiaoshi Wu, Keqiang Sun, Feng Zhu, Rui Zhao, and Hongsheng Li. 2023. Better aligning text-to-image models with human preference. *arXiv preprint arXiv:2303.14420*, 1(3).
- Jiazheng Xu, Xiao Liu, Yuchen Wu, Yuxuan Tong, Qinkai Li, Ming Ding, Jie Tang, and Yuxiao Dong. 2024. Imagereward: Learning and evaluating human preferences for text-to-image generation. *Advances in Neural Information Processing Systems*, 36.
- Peter Young, Alice Lai, Micah Hodosh, and Julia Hockenmaier. 2014. From image descriptions to visual denotations: New similarity metrics for semantic inference over event descriptions. *Transactions of the Association for Computational Linguistics*, 2:67–78.
- Qifan Yu, Zhebei Shen, Zhongqi Yue, Yang Wu, Wenqiao Zhang, Yunfei Li, Juncheng Li, Siliang Tang, and Yueting Zhuang. 2024a. Mastering collaborative multi-modal data selection: A focus on informativeness, uniqueness, and representativeness. *Preprint*, arXiv:2412.06293.
- Tianyu Yu, Yuan Yao, Haoye Zhang, Taiwan He, Yifeng Han, Ganqu Cui, Jinyi Hu, Zhiyuan Liu, Hai-Tao Zheng, Maosong Sun, et al. 2024b. Rlhf-v: Towards trustworthy mllms via behavior alignment from fine-grained correctional human feedback. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13807–13816.
- Weihao Yu, Zhengyuan Yang, Linjie Li, Jianfeng Wang, Kevin Lin, Zicheng Liu, Xinchao Wang, and Lijuan Wang. 2023. Mm-vet: Evaluating large multimodal models for integrated capabilities. *arXiv preprint arXiv:2308.02490*.
- Yuqian Yuan, Hang Zhang, Wentong Li, Zesen Cheng, Boqiang Zhang, Long Li, Xin Li, Deli Zhao, Wenqiao Zhang, Yueting Zhuang, et al. 2024. Videorefer suite: Advancing spatial-temporal object understanding with video llm. *arXiv preprint arXiv:2501.00599*.
- Lei Zhang, Fangxun Shu, Sucheng Ren, Bingchen Zhao, Hao Jiang, and Cihang Xie. 2023. Compress & align: Curating image-text data with human knowledge. *arXiv preprint arXiv:2312.06726*.
- Shu Zhang, Xinyi Yang, Yihao Feng, Can Qin, Chia-Chih Chen, Ning Yu, Zeyuan Chen, Huan Wang, Silvio Savarese, Stefano Ermon, et al. 2024a. Hive: Harnessing human feedback for instructional visual editing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9026–9036.
- Wenqiao Zhang, Tianwei Lin, Jiang Liu, Fangxun Shu, Haoyuan Li, Lei Zhang, He Wanggui, Hao Zhou, Zheqi Lv, Hao Jiang, et al. 2024b. Hyperllava: Dynamic visual and language expert tuning for multimodal large language models. *arXiv preprint arXiv:2403.13447*.
- Wenqiao Zhang, Siliang Tang, Yanpeng Cao, Shiliang Pu, Fei Wu, and Yueting Zhuang. 2019. Frame augmented alternating attention network for video question answering. *IEEE Transactions on Multimedia*, 22(4):1032–1041.
- Wenqiao Zhang, Lei Zhu, James Hallinan, Shengyu Zhang, Andrew Makmur, Qingpeng Cai, and Beng Chin Ooi. 2022. Boostmis: Boosting medical image semi-supervised learning with adaptive pseudo labeling and informative active annotation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 20666–20676.
- Chunting Zhou, Pengfei Liu, Puxin Xu, Srinivasan Iyer, Jiao Sun, Yuning Mao, Xuezhe Ma, Avia Efrat, Ping Yu, Lili Yu, et al. 2024. Lima: Less is more for alignment. *Advances in Neural Information Processing Systems*, 36.
- Jiaqi Zhu, Shaofeng Cai, Fang Deng, Beng Chin Ooi, and Wenqiao Zhang. 2023. Meter: A dynamic concept adaptation framework for online anomaly detection. *arXiv preprint arXiv:2312.16831*.

A Training Details

For the balance between speed and precision, all models are trained with BF16 and TF32 enabled, using LoRA with $r = 128$ and $\alpha = 256$. Model parameters are updated using AdamW optimizer with $\beta_1 = 0.9$, $\beta_2 = 0.999$, $\epsilon = 1e - 8$ and weight decay 0. For all models, learning rate is set to $2e - 5$, with warmup ratio 0.03 and cosine decay schedule. Reward models are trained on $4 \times$ RTX 3090s for 5 epochs, using global batch size 64, spending 8 hours for the question model and 9 hours for the answer model. LLaVA-1.5 instruction fine-tuning is completed on $8 \times$ A100s for 1 epoch with global batch size 128. The training takes 5.5/9.5 hours (7B/13B models) on our aligned dataset and 7/12 hours (7B/13B models) on the full dataset.

B Detailed Guidelines for Human Annotation

Question Annotation Guidelines:

- **Correctness.** Questions are evaluated for their alignment with both image content and general world knowledge. Questions containing inaccuracies, contradictions, or illogical statements are penalized.

- **Fluency.** Questions are assessed for grammatical correctness, clarity, and coherence. Ambiguous or poorly structured questions are unacceptable.
- **Relevance.** Questions are expected to be highly dependent on the visual information provided in the image. Questions easily answered without image context or those that could be directly inferred from general knowledge receive lower scores. Moreover, questions that inadvertently revealed image details were penalized. To ensure question solvability, answers are required to be definitively obtainable from the image, in conjunction with real-world knowledge.
- **Question Distribution (for multi-turn conversation).** In multi-turn dialogues, questions are additionally evaluated based on their diversity in difficulty and subject matter. Ideally, questions should progress from simpler to more complex, covering a broad range of image elements. Repetitive or redundant questions are considered suboptimal.

Answer Annotation Guidelines:

- **Accuracy.** Answers are evaluated for their correctness and coherence, ensuring alignment with both image content and general knowledge. Responses containing factual errors or contradictions are penalized.
- **Completeness.** Answers are expected to provide exhaustive responses to the posed questions, incorporating as much relevant image-derived detail as possible.
- **Reasoning.** For questions requiring logical deduction, answers are assessed based on the clarity and correctness of the reasoning process presented. Step-by-step explanations are preferred.
- **Relevance.** Answers are expected to be focused on the question at hand, avoiding extraneous information unrelated to the image or query.

C Adaptation to Multi-Turn Conversations

To accommodate the *conversation* type instructions with multi-turn question-answer pairs, adaptations

Method	Direct Prediction		Benchmarks	
	Question	Answer	MME	LLaVA ^w
CLIP Score	44.6	51.5	1389.4	40.4
Ours	60.6	65.1	1448.0	50.2

Table 5: Comparison between CLIP and our reward model.

are necessary. Within each example (I, Q_0, A_0) from the seed dataset, the question and answer components can be decomposed into sequences: $Q_0 = \{q_{0,1}, \dots, q_{0,t}\}$ and $A_0 = \{a_{0,1}, \dots, a_{0,t}\}$, where $(q_{0,i}, a_{0,i}), i \in \{1, \dots, t\}$ represents a single-turn question-answer pair, and the number of turns is determined by the question sequence.

Our approach involves a hybrid strategy: evaluating multi-turn questions holistically, while assessing answers on a per-turn basis. This distinction is grounded in the disparate criteria for these two components. Given the inherent interconnectedness and potential redundancies within multi-turn question sequences, a comprehensive evaluation is essential to capture their overall quality. Conversely, the standards for answer quality are predominantly determined on the current question and image, rendering contextual factors less critical. Consequently, isolating answer evaluation to the current turn optimizes computational efficiency and prevents the introduction of irrelevant information.

Specifically, for the question component, human evaluators assess question groups holistically within the same image, and the question reward model also processes entire question groups as input, with individual questions concatenated. For multi-turn answers, each t -turn example $(I, Q_0, A_0, A_1, \dots, A_l)$ is decomposed into t single-turn examples $(I, q_{0,1}, a_{0,1}, \dots, a_{l,1}), \dots, (I, q_{0,t}, a_{0,t}, \dots, a_{l,t})$ for subsequent processing, which mirrors the single-turn pipeline. The sole deviation occurs in the second filtration stage of human alignment, where the optimal answer is independently selected for each turn, and their average score serve as the overall answer quality metric for the multi-turn instance.

D Reward Models Evaluation

To construct question and answer reward models, we adapted the LLaVA-1.5-7B architecture by substituting its language modeling head with a sequence classification head, whose output dimension was set as 1. These models were trained on

human preference datasets derived from the first step of our pipeline, comprising 800 question instances and 960 answer instances, and further partitioned into 700/100 and 800/160 training/testing sets. Model training was conducted on the training split, followed by evaluation of prediction accuracy on the test split. For comparison, we additionally incorporated a baseline method of CLIP (Radford et al., 2021) Score between the image and instruction, which is typical of non-human-feedback evaluation. We also integrated both methods into the data curation pipeline, and subsequently compared the trained LLaVA-1.5 performance. As detailed in Table 5, our proposed reward models consistently outperform the CLIP Score baseline in both direct assessment and trained LLaVA-1.5 performance. These results underscore the effectiveness of our reward models in capturing human-centric evaluation criteria for both questions and answers.

We have further analyzed some failure cases of our reward model in Table 6 and Table 7. It can be found that in reward models’ misprediction cases, the two compared samples are often similar in quality, and the model gives close scores to both samples. Although recorded as a prediction error, the choice of either sample have no significant impact on data quality.

E Prompts

For question generation using text-only GPT4, we directly use the same prompts as LLaVA, and discard the generated answer parts. For answer generation, we adjust the prompts and provide the questions to ensure the model only produces answers. Moreover, since visual-capable GPT-4V and CogVLM directly see the image, we further modify the prompts to accommodate changes in input format. Table 8, Table 9, Table 10, Table 11, Table 12 show detailed prompts used for generation. Additionally, in the LLM alignment step, our prompts for rewriting and review are detailed in Table 13.

F Additional Examples

We exhibit more examples of different categories to demonstrate our alignment pipeline, as shown in Table 14, Table 15 and Table 16.



Question Failure Case 1	
	
Question 1 (Human Preferred): What tradition or practice could be observed in this city street? RM Score: 4.3	Question 2: Given the presence of both a horse and a car, what kind of transportation conflict might exist? RM Score: 4.8
Question Failure Case 2	
	
Question Group 1 (Human Preferred): 1. How many people are visible in this restaurant scene? 2. Where are the people located? Are they standing or sitting? 3. What types of objects can you identify in this image besides people? 4. What kind of venue does it seem to be, based on the setting and objects present? 5. Based on the image, describe the ambiance or atmosphere of the venue. RM Score: 7.5	Question Group 2: 1. How many people can you see in the restaurant image? 2. What are the people doing in the restaurant? 3. What types of objects are visible in the restaurant? 4. How many televisions are there on the wall? 5. What kind of atmosphere does the image suggest about the restaurant? RM Score: 8.6

Table 6: Question reward model failure cases.

Answer Failure Case 1



Question: Can you see people walking on the city street after the rain in the image?

Answer 1 (Human Preferred):

Yes, there are people walking down the city street after it has been raining. They are likely residents or visitors navigating their way through the city, possibly visiting the restaurants and shops along the street.

RM Score: 13.3

Answer 2:

Yes, people can be seen walking down the city street following the rain, likely resuming their activities now that the weather has improved.

RM Score: 13.6

Answer Failure Case 2



Question: What can be inferred about the airplane's current activity?

Answer 1 (Human Preferred):

Based on the details available in the image, it can be inferred that the airplane is likely either loading or unloading passengers at the airport terminal. The presence of a passenger ramp attached to the airplane typically indicates this kind of activity. The process might involve passengers boarding the plane for an upcoming flight or disembarking after a concluded journey. Additionally, while this process takes place, the airplane's cargo could also be being loaded or unloaded, and it might be undergoing routine checks and refueling in preparation for its next journey.

RM Score: 12.1

Answer 2:

It can be inferred that the airplane is currently engaged in the process of boarding or deboarding passengers. The connected jetway suggests that there is active interaction between the airplane and the airport terminal, indicating the beginning or conclusion of a flight. The positioning at the gate also implies preparations or finalizations are being made for a flight, which could include loading or unloading of luggage, refueling, and necessary checks by the ground crew before departure or after arrival.

RM Score: 12.7

Table 7: Answer reward model failure cases.

Complex Reasoning

GPT-4

You are an AI visual assistant that can analyze a single image. You receive five sentences, each describing the same image you are observing. In addition, specific object locations within the image are given, along with detailed coordinates. These coordinates are in the form of bounding boxes, represented as (x1, y1, x2, y2) with floating numbers ranging from 0 to 1. These values correspond to the top left x, top left y, bottom right x, and bottom right y.

The task is to use the provided caption and bounding box information, create a plausible question about the image, and provide the answer in detail.

Create complex questions beyond describing the scene. To answer such questions, one should require first understanding the visual content, then based on the background knowledge or reasoning, either explain why the things are happening that way, or provide guides and help to user's request. Make the question challenging by not including the visual content details in the question so that the user needs to reason about that first.

Instead of directly mentioning the bounding box coordinates, utilize this data to explain the scene using natural language. Include details like object counts, position of the objects, relative position between the objects.

When using the information from the caption and coordinates, directly explain the scene, and do not mention that the information source is the caption or the bounding box. Always answer as if you are directly looking at the image.

<captions><bounding boxes>

GPT-4V/CogVLM

You are an AI visual assistant that can analyze a single image.

The task is to create a plausible question about the image, and provide the answer in detail.

Create complex questions beyond describing the scene. To answer such questions, one should require first understanding the visual content, then based on the background knowledge or reasoning, either explain why the things are happening that way, or provide guides and help to user's request. Make the question challenging by not including the visual content details in the question so that the user needs to reason about that first.

Include details like object counts, position of the objects, relative position between the objects.

Do not mention anything from the prompt in your response.

<image>

Table 8: Prompts for *complex reasoning* question generation.

Conversation

GPT-4

You are an AI visual assistant, and you are seeing a single image. What you see are provided with five sentences, describing the same image you are looking at. Answer all questions as you are seeing the image.

Design a conversation between you and a person asking about this photo. The answers should be in a tone that a visual AI assistant is seeing the image and answering the question. Ask diverse questions and give corresponding answers.

Include questions asking about the visual content of the image, including the object types, counting the objects, object actions, object locations, relative positions between objects, etc. Only include questions that have definite answers: (1) one can see the content in the image that the question asks about and can answer confidently; (2) one can determine confidently from the image that it is not in the image. Do not ask any question that cannot be answered confidently.

Also include complex questions that are relevant to the content in the image, for example, asking about background knowledge of the objects in the image, asking to discuss about events happening in the image, etc. Again, do not ask about uncertain details. Provide detailed answers when answering complex questions. For example, give detailed examples or reasoning steps to make the content more convincing and well-organized. You can include multiple paragraphs if necessary.

<captions><bounding boxes>

GPT-4V/CogVLM

You are an AI visual assistant, and you are seeing a single image. Answer all questions according to the image.

Design a conversation between you and a person asking about this photo. The answers should be in a tone that a visual AI assistant is seeing the image and answering the question. Ask diverse questions and give corresponding answers.

Include questions asking about the visual content of the image, including the object types, counting the objects, object actions, object locations, relative positions between objects, etc. Only include questions that have definite answers: (1) one can see the content in the image that the question asks about and can answer confidently; (2) one can determine confidently from the image that it is not in the image. Do not ask any question that cannot be answered confidently.

Also include complex questions that are relevant to the content in the image, for example, asking about background knowledge of the objects in the image, asking to discuss about events happening in the image, etc. Again, do not ask about uncertain details. Provide detailed answers when answering complex questions. For example, give detailed examples or reasoning steps to make the content more convincing and well-organized. You can include multiple paragraphs if necessary.

<image>

Table 9: Prompts for *conversation* question generation.

Complex Reasoning

GPT-4

You are an AI visual assistant that can analyze a single image. You receive five sentences, each describing the same image you are observing. In addition, specific object locations within the image are given, along with detailed coordinates. These coordinates are in the form of bounding boxes, represented as (x1, y1, x2, y2) with floating numbers ranging from 0 to 1. These values correspond to the top left x, top left y, bottom right x, and bottom right y.

Then, you are given a plausible question about the image. The task is to use the provided caption and bounding box information to provide the answer in detail.

To answer such questions, you should first understanding the visual content, then based on the background knowledge or reasoning, either explain why the things are happening that way, or provide guides and help to user's request.

Instead of directly mentioning the bounding box coordinates, utilize this data to explain the scene using natural language. Include details like object counts, position of the objects, relative position between the objects.

When using the information from the caption and coordinates, directly explain the scene, and do not mention that the information source is the caption or the bounding box. Always answer as if you are directly looking at the image.

<captions><bounding boxes>

Question: <question>

===

Answer:

GPT-4V/CogVLM

You are an AI visual assistant that can analyze a single image.

The task is to create a plausible question about the image, and provide the answer in detail.

Create complex questions beyond describing the scene. To answer such questions, one should require first understanding the visual content, then based on the background knowledge or reasoning, either explain why the things are happening that way, or provide guides and help to user's request. Make the question challenging by not including the visual content details in the question so that the user needs to reason about that first.

Include details like object counts, position of the objects, relative position between the objects.

Do not mention anything from the prompt in your response.

<image>

Question: <question>

===

Answer:

Table 10: Prompts for *complex reasoning* answer generation.

Conversation
GPT-4 System Message You are an AI visual assistant, and you are seeing a single image. What you see are provided with five sentences, describing the same image you are looking at. Answer all questions as you are seeing the image. The answers should be in a tone that a visual AI assistant is seeing the image and answering the question. Provide detailed answers when answering complex questions. For example, give detailed examples or reasoning steps to make the content more convincing and well-organized. You can include multiple paragraphs if necessary. <captions><bouding boxes> GPT-4V/CogVLM System Message You are an AI visual assistant, and you are seeing a single image. Answer all questions according to the image. High-quality answers should demonstrate a deep understanding of the visual information, drawing connections between different elements in the image and providing thoughtful insights. Your responses should be clear, concise, and logically structured. For complex questions, give detailed examples or reasoning steps to make the content more convincing and well-organized. You can include multiple paragraphs if necessary. Do not mention anything from the prompt in your response. <image>
Turn 1 <system message> Question: <question 1> === Answer:
Turn 2 <system message> Question: <question 1> === Answer: <generated answer 1> === Question: <question 2> === Answer:
...

Table 11: Prompts for *conversation* answer generation. Answers of different turns are generated autoregressively.

Detail Description
<p>GPT-4</p> <p>You are an AI visual assistant that can analyze a single image. You receive five sentences, each describing the same image you are observing. In addition, specific object locations within the image are given, along with detailed coordinates. These coordinates are in the form of bounding boxes, represented as (x1, y1, x2, y2) with floating numbers ranging from 0 to 1. These values correspond to the top left x, top left y, bottom right x, and bottom right y.</p> <p>Using the provided caption and bounding box information, describe the scene in a detailed manner. Instead of directly mentioning the bounding box coordinates, utilize this data to explain the scene using natural language. Include details like object counts, position of the objects, relative position between the objects.</p> <p>When using the information from the caption and coordinates, directly explain the scene, and do not mention that the information source is the caption or the bounding box. Always answer as if you are directly looking at the image.</p> <p><captions><bounding boxes></p> <p>GPT-4V/CogVLM</p> <p>You are an AI visual assistant that can analyze a single image.</p> <p>Given the image, your task is to describe the scene in a detailed manner.</p> <p>Include details like object counts, position of the objects, relative position between the objects.</p> <p>Do not mention anything from the prompt in your response.</p> <p><image></p>

Table 12: Prompts for *detail description* answer generation.

LLM Alignment
<p>Rewrite</p> <p>Given the following Question and Answer, you are required to revise the Question and Answer in your writing style without changing the semantic meaning. If you think the original question and original answer are clear and consistent with your writing style, just leave it unchanged. The response should contain just the revised question and revised answer and the explanation of revision, formatted as: 'Revised Question:', 'Revised Answer:', and 'Explanation:'.</p> <p>Question: <original question></p> <p>Answer: <original answer></p>
<p>Review</p> <p>Given the following Original Question, Revised Question, Original Answer, and the Revised Answer, if the Revised Question or Revised Answer is inconsistent with your writing style, or deviates from the semantics of Revised Question or Original Answer, or adds or lacks any information, output "There is something wrong with the Revised Question or Revised Answer." Otherwise, output "The Revised Question and Revised Answer are fine." After making your decision, please provide a detailed explanation of your reasoning.</p> <p>Original Question: <original question></p> <p>Original Answer: <original answer></p> <p>Revised Question: <revised question></p> <p>Revised Answer: <revised answer></p>

Table 13: Prompts for LLM alignment.



Human Alignment (Question)

Question 1 (Accepted): What possible purpose could the handheld object in the image serve for the person?

Question 2: How does the inclusion of traffic light in public exhibits reflect societal values and priorities concerning urban planning and safety?

Human Alignment (Answer)

Question: What possible purpose could the handheld object in the image serve for the person?

Answer 1 (Accepted): The handheld object is not distinctly clear in the image, but given the context of the person walking past a traffic light in what appears to be an indoor setting, it is possible that the object could serve as a guide or an educational tool. If we imagine a scenario, it might be an interactive exhibit or part of a demonstration to teach about road safety or traffic signals. The individual might be using it to engage with the exhibit, perhaps by pressing buttons to change the lights or to simulate crossing the street while following traffic signals. Such educational setups are often found in museums, especially in those tailored towards children, to provide a hands-on learning experience about everyday situations and safety.

Answer 2: The handheld object the person is carrying appears to be a handbag. Handbags serve multiple purposes, such as carrying personal belongings, essential items, and accessories. For example, the person might have their wallet, keys, mobile phone, and other items in the handbag to keep them organized and easily accessible while walking past the traffic light, be it in a busy urban setting like Walnut Depot in the Bronx, or elsewhere. Additionally, handbags can also serve as a fashion accessory or a style statement, complementing the person's outfit and overall appearance.

Answer 3: In the image, the person walking past the traffic light is carrying a handbag. This handheld object serves multiple purposes for the individual. It is likely being used to carry personal belongings, such as wallet, keys, mobile phone, or other essentials while they are out walking. It can also serve as a fashion accessory, complimenting the individual's overall outfit. Considering the location, where they are walking past a traffic light in a city setting, and the presence of a bus in the background, it's plausible that this person is commuting and the handbag might carry items required for their journey or destination, such as books, electronics, or work-related items.

LLM Rewriting

Revised Question: What could be the potential use of the handheld object the person in the image is interacting with?

Revised Answer: The handheld object appears to be an interactive exhibit or educational tool, possibly found in a museum or a child-oriented educational setting. The person might be using it to learn about road safety and traffic signals by pressing buttons to simulate crossing the street while following traffic signals or change the lights. This type of hands-on learning experience is designed to teach children and adults about everyday situations and safety in a fun and engaging way.

Table 14: An alignment example of *complex reasoning* type instruction.

Detail Description



Human Alignment

Answer 1: A red city bus is traveling down a busy city street, surrounded by parked cars on both sides. The bus is moving between the vehicles and takes up a significant portion of the available street space.

In addition to the parked cars, there are also a few people present in the scene. One person is walking near the middle of the scene, while two other individuals are located towards the right side of the image. The scene captures the typical hustle and bustle of a city street with various modes of transportation and pedestrians going about their daily activities.

Answer 2 (Accepted): This photo captures a red single-decker bus in service, designated as route 152, which is heading towards Liverpool Street. The bus has a digital display that shows the route number and destination, and it's driving down a city street lined with Georgian-style terraced houses. The scene is set on a day with clear skies.

The road is shared with other vehicles, including parked cars along the curb on the right side. The architectural uniformity of the terraced houses along the street provides a historical backdrop to the modern activity of urban transportation. The setting suggests a typical urban street scene, showcasing the integration of public transport within a residential area.

Answer 3: This photo depicts a red bus on a street lined with parked cars and buildings. The bus has a destination sign indicating "Liverpool Street" and a route number of "153". The location appears to be in an urban area, likely in the UK, as indicated by the style of the bus and the street architecture, which includes terraced houses with brick facades. The street is relatively narrow, and the bus is driving in what seems to be a typical residential or mixed-use area. The sky is partly cloudy, suggesting it might be either late afternoon or early morning.

LLM Rewriting

Revised Answer: This photo appears to have been captured in a residential area in a city, possibly in the UK. The red single-decker bus with the route number 152 and destination of Liverpool Street suggests that the bus is in service and heading towards the Liverpool Street station. The clear skies and the daytime suggest that the photo was taken during daylight hours. The Georgian-style terraced houses on either side of the road suggest that the area has a historical background. The presence of other vehicles, including parked cars along the curb on the right side, indicates that the road is shared with other transportation modes. Overall, the photo showcases the integration of public transport within a residential area.

Table 15: An alignment example of *detail description* type instruction.



Human Alignment (Question)

Questions 1:

What is the main dish on the plate?

Is the meal on a plate or a tray?

What is the state of the meal?

Can you give an example of how to create a delicious **homemade pizza**?**Questions 2 (Accepted):**

What is shown in the image?

Where is the dipping sauce located on the plate?

Are there any utensils visible in the image?

What dining setting does this image suggest?

Human Alignment (Answer)

Question: What is shown in the image?**Answer 1:** The image shows a plated dish of three crab cakes, accompanied by a small side salad with mixed greens and a mango salsa. There are also **two** dollops of a creamy dipping sauce on opposite sides of the plate.**Answer 2 (Accepted):** The image shows a plate of crab cakes, a popular seafood dish. The crab cakes are served with pineapple salsa, creamy sauce, and fresh greens.**Question:** Where is the dipping sauce located on the plate?**Answer 1:** The dipping sauce is located at **two spots** on the plate, one on the left and one on the right, symmetrically placed near the edges.**Answer 2 (Accepted):** The dipping sauce is located in **three small dollops** on the plate. They are positioned **around the crab cakes**.**Question:** Are there any utensils visible in the image?**Answer 1 (Accepted):** Yes, there is **a fork visible to the right** of the plate, along with **a knife to the left**. These utensils suggest that the dish is ready to be eaten.**Answer 2:** Yes, there are utensils visible in the image. You can see **a fork and a knife on the right side** of the plate.**Question:** What dining setting does this image suggest?**Answer 1:** The image suggests a casual yet refined dining setting. The plating is elegant but not overly formal, and the presence of a wine glass in the background indicates that this could be a meal served in a restaurant or at a home dinner party. The overall atmosphere seems relaxed, making it suitable for a pleasant, leisurely meal.**Answer 2 (Accepted):** The image suggests a dining setting that leans towards a casual to mid-range restaurant. The uncomplicated plate design, paired with the straightforward presentation of the crab cakes, pineapple salsa, and creamy sauce, hints at a relaxed atmosphere. The inclusion of a fork and knife, while standard, further supports the notion of a casual yet attentive dining experience. This type of setting often caters to a broad clientele seeking a satisfying meal without excessive formality.

LLM Rewriting

Revised Question: What is depicted in the image?**Revised Answer:** The image depicts a plate of crab cakes, a beloved seafood dish. The crab cakes are accompanied by pineapple salsa, creamy sauce, and a sprig of fresh greens.**Revised Question:** How is the dipping sauce presented on the plate?**Revised Answer:** The dipping sauce is presented in three small dollops on the plate, positioned around the crab cakes.**Revised Question:** What utensils can be seen in the image?**Revised Answer:** A fork is visible to the right of the plate, and a knife is visible to the left. These utensils suggest that the dish is ready to be eaten.**Revised Question:** In what type of dining setting does the image suggest the crab cakes and pineapple salsa are being served?**Revised Answer:** The image suggests that the crab cakes and pineapple salsa are being served in a casual to mid-range restaurant. The straightforward presentation of the dish and the inclusion of a fork and knife hint at a relaxed atmosphere. This type of setting caters to a broad clientele seeking a satisfying meal without excessive formality.Table 16: An alignment example of *conversation* type instruction.