

OrderChain: Towards General Instruct-Tuning for Stimulating the Ordinal Understanding Ability of MLLM

Jinhong Wang^{1,2,3*} Shuo Tong^{2*} Jian Liu³ Dongqi Tang³ Weiqiang Wang³ Wentong Li¹
 Hongxia Xu² Danny Z. Chen⁴ Jintai Chen^{5†} Jian Wu^{2†}

¹College of Computer Science & Technology, Zhejiang University

²Transvascular Implantation Devices Research Institute and Liangzhu Laboratory

³Ant Group

⁴University of Notre Dame

⁵HKUST (Guangzhou)

*Contribute equally †Corresponding Authors

Abstract

Despite the remarkable progress of multimodal large language models (MLLMs), they continue to face challenges in achieving competitive performance on ordinal regression (OR; a.k.a. ordinal classification). To address this issue, this paper presents OrderChain, a novel and general prompting paradigm that improves the ordinal understanding ability of MLLMs by specificity and commonality modeling. Specifically, our OrderChain consists of a set of task-aware prompts to facilitate the specificity modeling of diverse OR tasks and a new range optimization Chain-of-Thought (RO-CoT), which learns a commonality way of thinking about OR tasks by uniformly decomposing them into multiple small-range optimization subtasks. Further, we propose a category recursive division (CRD) method to generate instruction candidate category prompts to support RO-CoT automatic optimization. Comprehensive experiments show that LLaVA model with our OrderChain improves baseline LLaVA significantly on diverse OR datasets, e.g., from 47.5% to 93.2% accuracy on the Adience dataset for age estimation, and from 30.0% to 85.7% accuracy on the Diabetic Retinopathy dataset. Notably, LLaVA with our OrderChain also remarkably outperforms state-of-the-art methods by 27% on accuracy and 0.24 on MAE on the Adience dataset. To our best knowledge, our OrderChain is the first work that augments MLLMs for OR tasks, and the effectiveness is witnessed across a spectrum of OR datasets. Project Page: <https://order-chain.github.io/>.

1. Introduction

With the success of Large language models (LLMs), e.g., GPT3 [6], LLaMA [49], Gemini [48], and Qwen [3], re-




Figure 1. Comparing traditional ordinal regression (OR) methods and modern MLLM methods. Traditional OR methods perform well but need to train separate models for different tasks. MLLMs can be unified for diverse tasks but struggle in performance.

searchers have been enhancing these models by incorporating visual understanding capabilities and led to emergence of multimodal large language models (MLLMs), such as LLaVA [34] and GPT-4V [40] and Qwen-VL [3, 4], exhibiting capabilities comparable to human intelligence in visual understanding tasks [53, 55]. Despite these advances, the potential of MLLMs on order understanding for ordinal regression tasks is not yet well explored.

Ordinal regression (OR) refers to classifying object instances into ordinal categories, and is crucial for applications in various areas like facial age estimation [26, 38, 47], image aesthetics assessment [23, 24, 42], medical disease grading [17, 36] and so on. The category labels of these tasks all follow a natural order. Unlike general classification tasks, ordinal understanding is a crucial issue for the representation learning of OR tasks. Mainstream methods in the past, including order distribution learning [29, 39], instance comparison [47], and CLIP-based [13, 30, 51], all revolve around this point. Although these methods are effective, it is still challenging to train a unified model for all OR tasks due to different specificities of diverse tasks (e.g., in terms of the number and range of categories). Therefore, separate models are still needed for different tasks (see Fig. 1).

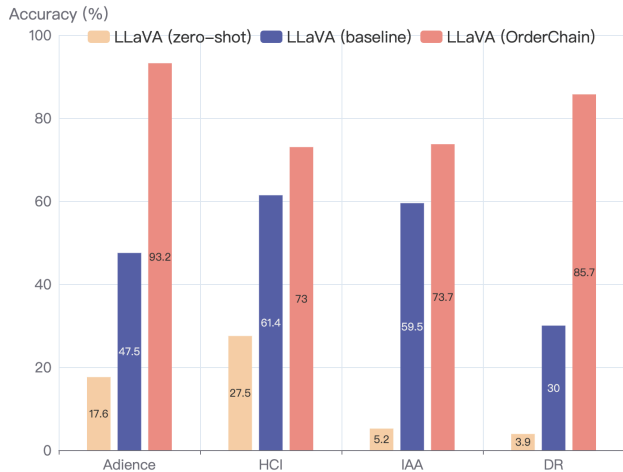


Figure 2. The performance of different versions of LLaVA on diverse datasets. The accuracy of Vanilla LLaVA is less than 60% on most datasets. Adience: Age estimation dataset. HCI: Historical color image dating dataset. IAA: Image aesthetics assessment dataset. DR: Diabetic retinopathy grading dataset.

MLLMs appear promising to address this challenge with their extensible language system. However, through investigation, we find that no matter whether MLLMs are training-free (zero-shot) for direct inference or are fine-tuned by image-label pairs (baseline), their performance is not satisfactory, as shown in Fig. 2. The main challenges that hinder the model performance can be attributed to: (1) Lack of specificity modeling. Especially, when using MLLM for zero-shot inference, the conceptual information of domain knowledge and prior knowledge of category boundaries and range are overlooked, which contain task-related hints for task-aware ordinal understanding. (2) Lack of commonality modeling. For OR tasks, their most striking similarity is that the categories have an order. Thus, it is crucial to learn a commonality way of thinking about OR tasks which can strengthen MLLMs’ understanding of category order.

To tackle these two challenges, this work investigates how the ordinal understanding capability of MLLMs can be realized by developing a novel Chain-of-Thought method, called OrderChain. OrderChain is driven by three pivotal motivations. **First**, domain-related knowledge needs to be utilized to guide the MLLM to extract critical imaging features and make predictions more effectively. **Second**, for different OR tasks, the knowledge on the number and range of candidate category labels needs to be injected to mitigate the out-of-bound predictions made by the MLLM. Both these two types of knowledge can help the MLLM to model the specificity of different OR tasks. **Third**, inspired by the idea of decomposition, OR tasks can be decomposed into sub-interval classification tasks. Through continuous decomposition and classification, the candidate range can be gradually optimized to obtain the final predic-

tion. This insight is applicable to all tasks whose categories have an order, which can help the MLLM to model the commonality of OR tasks. Based on these considerations, our proposed OrderChain consists of the following key components: (1) **Customized task-aware prompts to guide MLLMs in specificity modeling.** These include domain knowledge prompts, which provide prior knowledge to help understand the task, and category feature prompts, which offer information on the number and range of candidate labels to enable more accurate predictions. (2) **A Range Optimization Chain-of-Thought (RO-CoT) for progressive ordinal regression.** As shown in Fig. 3, RO-CoT transforms inputs into instruction templates, generates task definitions as domain knowledge prompts, recursively divides labels or coarse predictions into smaller subsets as category feature prompts, and iteratively refines predictions until a final response is achieved. This approach decomposes visual rating tasks into subtasks for gradual refinement. (3) **A category recursive division method for coarse-to-fine candidate label generation.** This method automatically splits RO-CoT predictions into refined subsets, which are used as category feature prompts to support continuous optimization in the prediction process.

Compared to traditional deep learning (DL) methods, Our OrderChain allows MLLMs to train a unified model for all OR tasks. Extensive experiments on OR datasets of various domains show the effectiveness of our OrderChain and its components. Notably, OrderChain achieves state-of-the-art performance in facial age estimation tasks, remarkably improving accuracy to 93.2% (a 27% improvement). In other domains of OR tasks, LLaVA with our OrderChain yields improvement of $\sim 12\%$ to $\sim 56\%$ compared to baseline LLaVA, showing highly competitive performance.

Our main contributions are summarized as follows.

- (A) For the first time, we explore the potential of MLLMs for ordinal regression tasks.
- (B) We propose a new prompting paradigm, called OrderChain, to instruct the MLLM to manage OR tasks in a general and progressive manner with a range optimization Chain-of-Thought.
- (C) We design a set of task-aware prompts, including domain knowledge prompts and category feature prompts for enhancing MLLM task-specific knowledge modeling.
- (D) Extensive experiments show the effectiveness of our OrderChain on various OR datasets, which provides a promising way to establish a unified OR model.

2. Related Work

2.1. Multimodal Large Language Models (MLLMs)

Large language models (LLMs), such as GPT-3 [1], Qwen [2], and LLaMA [49], have attracted lots of atten-

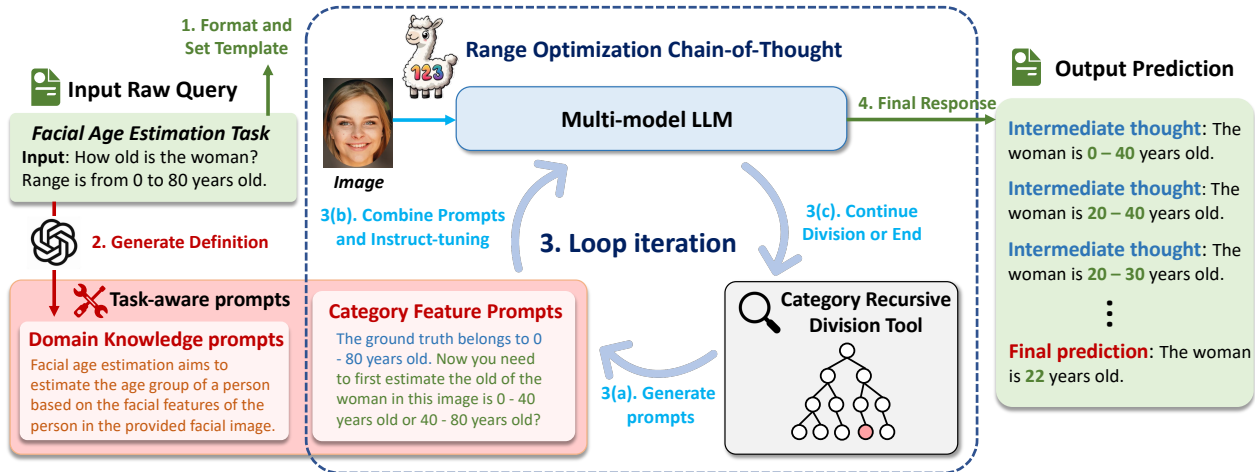


Figure 3. The overall framework of our proposed OrderChain for improving the ordinal understanding of Multi-modal LLMs (e.g., LLaVA).

tion for their remarkable capabilities across various linguistic tasks. This wave of interest has paved the way for the development of recent Multimodal Large Language Models (MLLMs), which integrate LLMs with visual encoders to enable an enriched comprehension and understanding of multimodal content. Prominent examples include the LLaVA series [32, 33, 35], GPT-4V [40], mPLUG-Owl [54], InstructBLIP [10], Qwen-VL [4], Google’s Gemini series [45, 48], etc. Based on LoRA [20], MLLMs can be fine-tuned for downstream generalization in various areas, such as detection [53] and segmentation [56]. These advances highlight the diverse and expanding landscape of MLLMs, which have remarkably impacted the landscape of Artificial General Intelligence (AGI). This work is the first to explore the potential of MLLMs for OR tasks, providing a promising way to build a unified OR model.

2.2. Chain-of-Thought (CoT)

CoT prompting is a specialized tool for inducing LLMs to produce intermediate reasoning steps that lead to a final answer and decision-making [52]. This technique elicits LLMs to generate a coherent series of intermediate reasoning steps that arrive at the final answer to a question. The traditional prompting method [6] performs poorly when it faces tasks that require reasoning abilities. Inspired by the concept of using intermediate steps to solve reasoning problems [9], the chain-of-thought method mimics a step-by-step thinking process and breaks a multi-step problem into intermediate steps, enabling the model to deduce more accurate results [52]. Moreover, CoT is also a very effective tool for applying LLMs to a variety of downstream scenarios. Diverse customized CoTs have emerged to be a powerful prompting paradigm in the vision domain, such as objective detection [53] and segmentation [27]. However, there is a lack of Chain-of-Thought methods tailored for OR tasks.

Since an OR task can be treated as a coarse-to-fine problem [50], our approach subtly designs a range optimization CoT, which is the first to propose using a CoT prompting paradigm to guide MLLMs for OR tasks.

2.3. Ordinal Regression

Given an input image, ordinal regression in computer vision aims to map the input to a rank or a continuous value. Many popular methods [7, 14, 16, 28, 46] adopt a classification framework. Many recent studies [25, 29, 31, 39] proposed ordinal distribution constraints to exploit the ordinal nature of regression. Adding prior order knowledge to loss calculation, several methods [11, 15] created soft labels artificially by changing the distances between categories. A few advanced methods [29, 37, 38, 47] sorted tuples that are formed by two or three instances with ordinal categories to learn the rank information. Ord2Seq [50] proposed to transform OR tasks as sequence prediction and solve ordinal regression using autoregressive models. Recent works like OrdinalCLIP [30], L2RCLIP [51], and NumCLIP [13] used CLIP [43] for OR tasks, focusing on designing the text encoder to map numerical labels to a continuous space for improved image-text alignment. Although these DL methods are general and effective, they need to train separate models for different OR tasks. This work explores utilizing MLLMs with our RO-CoT prompting paradigm to construct a unified OR model.

3. Method

3.1. Overview

Our work is the first to explore the effectiveness of MLLMs for OR tasks. Based on MLLMs, vision-language-driven OR tasks can be formulated as: Given a multimodal input (including images I , text L , etc.), output a classifica-

tion result $r \in S$ within a candidate category set $S = \{C_1, C_2, \dots, C_n\}$, where the candidate category labels are ordered, that is, $C_1 < C_2 < \dots < C_n$. Based on our exploration that found the poor performance of MLLMs on OR tasks with training-free inference, we propose a novel OrderChain method to improve the ordinal understanding of MLLMs for effective ordinal regression. Fig. 3 shows an overview of OrderChain. OrderChain introduces a Range Optimization Chain-of-Thought (Sec. 3.2), a Category Recursive Division Method (Sec. 3.3), and Task-aware prompts (Sec. 3.4) to conduct and reason the coarse-to-fine ordinal understanding process.

3.2. Range Optimization Chain-of-Thought

Since OR tasks have continuous ordinal labels, a coarse-to-fine paradigm can be utilized to predict ordinal labels progressively. Inspired by this observation, we introduce a new prompting paradigm, called Range Optimization Chain-of-Thought (RO-CoT), to enable the MLLM to solve OR tasks in a progressive range optimization manner. Specifically, the procedure of our RO-CoT consists of four parts:

1. **Formatting and Setting Template:** Format individualized raw query inputs to a specific template of common OR tasks and set corresponding recursive CoT logic.
2. **Definition Generation as Domain Knowledge Prompts:** For different OR tasks and user requirements, definitions and ideas of the tasks are generated through preliminary identification by LLM, which act as domain knowledge prompts to guide ensuing MLLM prediction.
3. **Loop Iteration in Each RO-CoT Step:**
 - (a) **Generating category feature prompts by the category recursive division (CRD) method** (discuss in Sec. 3.3): Given the initial optional categories or last intermediate thought, CRD will continue to divide these categories into multiple consecutive subsets with a smaller range as new candidate categories for the next step. These candidate subsets will serve as category feature prompts for instruction.
 - (b) **Using combined task-aware prompts to instruct MLLM** (discuss in Sec. 3.4): The task-aware prompts, combined with default domain knowledge prompts and newly generated category feature prompts, are used to instruct the MLLM to propose a thought which a new candidate subset that the category of the image belongs to.
 - (c) **Continue or end:** If the final prediction is made, MLLM will return the outputs; otherwise, MLLM will repeat step (a) to continue making refined predictions with the new prompts generated by CRD.
4. **Returning the Final Response:** When the subset contains only one category, the procedure will end and return the final prediction.

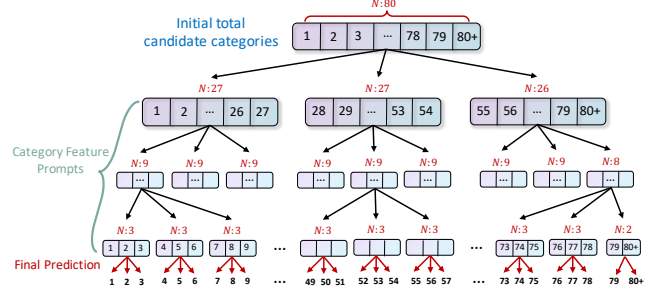


Figure 4. An example of the Category Recursive Division (CRD) process on the facial age estimation task. Assume that the total number of categories is 80 and CRD uses a trinomial balanced tree, that is, each node has three child nodes (subsets) and the difference in the number of categories per child node (sub candidate sets) is minimal. CRD first divides the entire range (80) of categories into three subsets with new category numbers 27, 27 and 26. It repeats this process until the number of categories for each subset is less than or equal to 3. Upon the last division, the final prediction is obtained. N at each node denotes the number of categories within the node (candidate set).

3.3. Category Recursive Division Method

To achieve range optimization to enhance the understanding of category order, we introduce a Category Recursive Division (CRD) Method to automatically plan the path of range optimization. CRD aims to divide the entire candidate categories into multiple more refined subsets. In this process, the range of candidate categories is optimized. Note that this strategy can be effective only when the categories are ordered, and is not suitable with non-order common classification tasks. Thus, this recursive division process can inject ordinal knowledge into the MLLM. Based on the automatic division, this method can limit as well as provide candidate options for each MLLM prediction. Specifically, assuming that the MLLM predicts a certain range, the CRD automatically queries the corresponding subsequent divisions, acting as a part of Category Feature Prompts to provide the MLLM with more refined candidate options and force the MLLM to focus on further refinement. To avoid the negative impact of class unbalance, we structure the division process into a balanced division tree. Based on the initial total number of categories, N_{init} , we use a k -tree for division. The total recursive steps, T , is calculated as:

$$T = \log_k(N_{init}). \quad (1)$$

For every step i , the maximum category number for each sub candidate set should be:

$$N_i = \frac{N_{init}}{k^i} + 1, \quad i = 1, 2, \dots, T. \quad (2)$$

Thus, the j -th candidate set $c_{i,j}$ of step i should be:

$$c_{i,j} = \{s_j, s_j + 1, \dots, \min(N_{init}, s_j + N_i - 1)\}, \quad (3)$$



Figure 5. An example of OrderChain for facial age estimation on the Adience dataset. **Brown:** Domain knowledge prompts. **Blue:** Description prompts. **Green:** Instruction prompts.

where the index of the starting category, s_j , for the candidate set $c_{i,j}$ is:

$$s_j = (j - 1) * N_i + 1, \quad j = 1, 2, \dots, k^i. \quad (4)$$

An example is shown in Fig. 4. In general, for datasets with fewer categories (e.g., image aesthetic assessment), we can use a binary tree. For datasets with more categories (e.g., face age estimation), we tend to use a trinomial tree to reduce the CoT length and number of thoughts. This method can be applied to all OR tasks, helping develop a general ordinal understanding approach for MLLMs to learn and model the commonality of ordinal regression, that is, the intrinsic ordinal relations among the categories.

3.4. Task-aware Prompts

In the past, many general algorithms for ordinal regression have been proposed. Although effective in modeling common intrinsic order logic, they still required to train separate models for diverse OR tasks since the number and range of categories are objectively different. Subjectively, specific domain knowledge in different OR tasks is also difficult to model. Our observation is that the extensibility of MLLMs makes it possible to train a general OR model via prompt engineering. Based on this observation, we introduce task-aware prompts for OrderChain to enhance the modeling of the specificity of different OR tasks. Task-aware prompts

Method	Accuracy (%) \uparrow	MAE \downarrow
<i>Supervised SOTA</i>		
CNNPOR [38]	57.4	0.55
GP-DNNOR [39]	57.4	0.54
SORD [11]	59.6	0.49
POE [29]	60.5	0.47
OrdinalCLIP [30]	61.2	0.47
MWR [47]	62.6	0.45
Ord2Seq [50]	63.9	0.43
L2RCLIP [51]	66.2	0.36
<i>Zero-shot MLLM</i>		
LLaVA-1.5 [34]	17.6	1.48
<i>Lora Fine-tune MLLM [20]</i>		
LLaVA-1.5 (baseline) [34]	47.5	0.59
LLaVA-1.5 + OrderChain	93.2	0.12

Table 1. Accuracy and MAE comparison on the Adience dataset.

contain two types: Category Feature Prompts and Domain Knowledge Prompts, which are elaborated on below.

Category Feature Prompts. To allow MLLMs to know the prediction objective and operate following the range optimization Chain-of-Thought, we introduce category feature prompts with two parts of prompts, focusing on different aspects: **Description prompts** for describing the definition of the total categories and **Instruction prompts** for candidate categories instruction. The details are given below.

- Description prompts** include the range and number of categories for a specific task (see the **blue** part of Category Feature Prompts in Fig. 3), for overcoming the disadvantage that the dimension of fully connected layers of the traditional DL models cannot be changed. By setting different description prompts, it is promising to train a universal MLLM to handle all OR problems.
- Instruction prompts** include candidate categories need to be refined at this step (see the **green** part of Category Feature Prompts in Fig. 3). These prompts are obtained by the Category Recursive Division method based on previous predictions, and they act as the new query aiming to provide new limiting candidate categories for the MLLM to predict at this step.

Domain Knowledge Prompts. Though the internal ordinal logic of different OR tasks is the same, the data features, estimation criteria, and so on can be very different. To model the specificity of different OR tasks, we introduce domain knowledge prompts, which are obtained by the MLLM itself and provide prior domain knowledge for the following predictions. This is equivalent to a preliminary identification of the task, so that the MLLM can search for information related to the task as much as possible as a guide.

Method	Accuracy (%) \uparrow	MAE \downarrow
Palermo et al. [41]	44.9	0.93
CNNPOR [38]	50.1	0.82
GP-DNNOR [39]	46.6	0.76
SORD [11]	53.4	0.70
POE [29]	54.7	0.66
MWR [47]	57.8	0.58
Ord2Seq [50]	60.9	0.52
OrdinalCLIP [30]	56.4	0.67
L2RCLIP [51]	67.2	0.59
NumCLIP [13]	69.6	0.35
LLaVA-1.5 (zero-shot) [34]	27.5	1.20
LLaVA-1.5 (baseline) [34]	61.4	0.50
LLaVA-1.5 + OrderChain	73.0	0.32

Table 2. Accuracy and MAE comparison on the HCI dataset.

4. Experiments

4.1. Datasets and Setup

Datasets. To validate the effectiveness of our OrderChain, we conduct experiments on OR tasks in various domains, including Facial Age Estimation, Historical Image Dating, Image Aesthetics Assessment, and Diabetic Retinopathy Grading. The datasets for these tasks are as follows.

- **Facial Age Estimation.** We use the Adience dataset [26] for age group estimation that contains about 26,580 face images of 2,284 subjects from Flickr. Ages are annotated in 8 groups: 0-2, 4-6, 8-13, 15-20, 25-32, 38-43, 48-53, and over 60 years old.
- **Historical Image Dating.** The historical color image (HCI) dataset [41] is for estimating the decades of historical color photos. There are five decades from 1930s to 1970s, annotated as 1 to 5. Each decade has 265 images.
- **Image Aesthetics Assessment.** The Aesthetics dataset [12] contains 15,687 Flickr image URLs, 13,706 of which are available. The dataset is used to grade image aesthetics. There are four image classes: animal, urban, people, and nature. Each image was graded by at least 5 different graders in 5 ranking categories to evaluate the photographic aesthetic quality: unacceptable, flawed, ordinary, professional, and exceptional. The ground truth is defined as the median rank among all the gradings.
- **Diabetic Retinopathy Grading.** The Diabetic Retinopathy (DR) dataset [8] contains 35,126 high-resolution fundus images. In this dataset, images were annotated in five levels of diabetic retinopathy from 1 to 5, representing no DR, mild DR, moderate DR, severe DR, and proliferative DR, respectively.

Experimental Setup. Our experiments are conducted on the PyTorch platform with an NVIDIA Tesla A100 GPU. We use LLaVA-1.5-7B as our MLLM backbone, in which the image encoder is ViT-L-16 pre-trained by CLIP. We

Method	Accuracy (%) \uparrow	MAE \downarrow
Poisson [5]	77.1	0.38
MT [44]	82.8	0.36
SORD [11]	78.2	0.73
POE [29]	80.5	0.30
CIG [8]	83.3	0.30
Ord2Seq [50]	84.2	0.25
LLaVA-1.5 (zero-shot) [34]	3.9	12.1
LLaVA-1.5 (baseline) [34]	30.0	0.99
LLaVA-1.5 + OrderChain	85.7	0.23

Table 3. Accuracy and MAE comparison on the DR dataset.

apply LoRA [20] to fine-tune MLLM. We employ the AdamW [22] optimizer with a learning rate of $2e-4$. We use a per-device batch size of 16. For fair comparison, all the known methods are implemented using their authors’ code or re-implemented based on the original papers. The division of all the datasets follows [30, 47, 50]. More details of the datasets, experimental settings, and codes are given in the Supplemental Document.

4.2. Experimental Results

Facial Age Estimation. Fig. 5 gives an example of OrderChain for facial age estimation on the Adience dataset, and Table 1 shows comparison results with state-of-the-art (SOTA) methods on the Adience dataset. One can see that the zero-shot version (training-free) and baseline version (LoRA fine-tuning) of LLaVA-1.5 perform worse than known supervised SOTA methods. Specifically, the zero-shot version of LLaVA-1.5 attains only 17.6% accuracy and baseline version of LLaVA -1.5 perform better but merely obtains a mediocre performance. In contrast, LLaVA with our proposed OrderChain achieves 93.2% accuracy and 0.12 MAE, outperforming SOTA methods L2RCLIP [51] by a remarkable margin of $\sim 27\%$ improvement in accuracy and 0.24 reduction in MAE. This indicates that when integrated with our OrderChain, the MLLM is instructed to think of the facial age estimation problem step by step with a few smaller refined subtasks, which allows the MLLM to learn the internal ordinal relationships between categories and effectively estimate the ages of faces with a wide range and a large number of categories. Our Ordchain achieves milestone performance on age estimation tasks, which further demonstrates the effectiveness of OrderChain in improving the ordinal understanding of MLLM.

Historical Image Dating. Table 2 compares the results on the HCI dataset. As can be seen, LLaVA with our OrderChain outperforms known methods and achieves state-of-the-art results, yielding improvements of 3.4% in Accuracy and 0.03 in MAE, which indicate the superiority of our new approach. Compared to the zero-shot and baseline versions

Method	Accuracy (%) \uparrow					MAE \downarrow				
	Nature	Animal	Urban	People	Overall	Nature	Animal	Urban	People	Overall
CNNPOR [38]	71.86	69.32	69.09	69.94	70.05	0.294	0.322	0.325	0.321	0.316
SORD [11]	73.59	70.29	73.25	70.59	72.03	0.271	0.308	0.276	0.309	0.290
POE [29]	73.62	71.14	72.78	72.22	72.44	0.273	0.299	0.281	0.293	0.287
Ord2Seq [50]	78.09	75.74	72.83	69.24	74.43	0.225	0.257	0.275	0.319	0.264
OrdinalCLIP [30]	73.65	72.85	73.20	72.50	73.05	0.273	0.279	0.277	0.291	0.280
L2RCLIP [51]	73.51	75.26	77.76	78.69	76.07	0.267	0.253	0.216	0.246	0.245
NumCLIP [13]	75.20	75.24	79.49	76.17	76.53	0.249	0.250	0.208	0.238	0.236
LLaVA-1.5 (zero-shot) [34]	3.68	8.01	10.8	1.47	5.21	1.422	1.109	1.439	0.901	1.275
LLaVA-1.5 (baseline) [34]	41.81	64.86	61.89	39.33	59.56	0.418	0.374	0.393	0.597	0.435
LLaVA-1.5 + OrderChain	73.91	75.61	77.87	66.26	73.83	0.229	0.260	0.252	0.297	0.260

Table 4. Results on the Image Aesthetics dataset. Accuracy and MAE are reported for each of the four image classes. The best three results are marked in **bold**, **orange**, and **blue**, respectively.

of LLaVA, LLaVA with our OrderChain achieves superior performance with remarkable improvement, validating the limitation of vanilla LLaVA and that the main improvement comes from our proposed OrderChain. In addition, we find that, like the facial age estimation task, the labels of the HCI dataset are also objective and true. For this kind of OR task, the MLLM has great potential and gains large improvement with our proposed OrderChain, indicating that OrderChain can learn a conceptual understanding of the essential order of objects in OR tasks.

Diabetic Retinopathy Grading. Table 3 shows the results on the DR dataset. Note that this dataset is unbalanced since the sample number decreases sharply as the severity DR level increases. Known methods yield poor performances due to the unbalanced data, like SORD [11], which is a modality-specific method utilizing modified soft labels, suffers serious errors in MAE. Worse, the zero-shot and baseline versions of LLaVA attain horrible performance. In comparison, LLaVA with our OrderChain still maintains competitive performances, achieving an Accuracy of 85.7% and an MAE of 0.23, which greatly outperforms the baselines and the other order learning methods, showing that our approach has better robustness on unbalanced data. It is mainly due to the category range division process of RO-CoT, which also enables better positive-negative distinction. Unlike one positive class against the other negative classes in previous work, it turns to classifying the first two categories against the last three categories in the first CoT step of OrderChain for the DR dataset (5 categories in total). In this way, the classification in a step is more category-balanced, which helps to better deal with unbalanced data. Moreover, our OrderChain provides a promising way to grade medical diseases of the OR type by a unified MLLM.

Image Aesthetics Assessment. Table 4 shows the results on the Image Aesthetics dataset. **We find that the MLLM with our OrderChain does not achieve optimal perfor-**



Figure 6. Quantitative visualization results of some examples of people aesthetics assessment.

mance, which we believe is due to the highly *subjective* nature of the labels made by human raters. The subjective differences between different people, and even between people and MLLMs for the definition of beauty, may be significant. Especially for the relatively low performance on the ‘People’ category, we suspect it is due to the pre-training imposed on the MLLM that tends to praise rather than demean people. To further illustrate this point, we provide quantitative visualization results as shown in the Fig. 6. Our model successfully predicted the aesthetic scores for the first and third images. However, for the second image, the model makes an incorrect prediction with a higher score, which also reflects the issue we described in the paper: **MLLMs have positive biases towards people.** In other relatively more objective categories, LLaVA with our OrderChain achieves higher performance, and thus the overall performance is mainly affected by the ‘People’ category. On the other hand, our proposed OrderChain remarkably improves vanilla LLaVA to a competitive level, demonstrating the effectiveness of our OrderChain.

4.3. Ablation Study

We conduct a comprehensive ablation study to examine the effectiveness of each key component in our OrderChain, including Domain Knowledge Prompts, Category Feature Prompts, and Range Optimization Chain-of-Thought. Except for the zero-shot version, all the experiments presented in this subsection are conducted on the Adience dataset

Method	Accuracy (%) \uparrow	MAE \downarrow
(a) LLaVA (zero-shot)	17.6	1.48
(b) LLaVA (baseline)	47.5	0.59
(c) LLaVA + Domain Knowledge Prompts	58.0 (+10.5)	0.49 (-0.10)
(d) LLaVA + Category Feature Prompts	32.5 (-15.0)	1.42 (+0.83)
(e) LLaVA + Domain Knowledge Prompts + Category Feature Prompts	38.7 (-8.8)	1.35 (+0.76)
(f) LLaVA + Category Feature Prompts + RO-CoT	84.6 (+37.1)	0.18 (-0.41)
(g) LLaVA + OrderChain	93.2 (+45.7)	0.12 (-0.47)

Table 5. Ablation experiments on the Adience dataset. RO-CoT denotes range optimization Chain-of-Thought.

and use LoRA to fine-tune. Specifically, LLaVA (zero-shot) denotes directly using LLaVA for training-free inference. LLaVa (baseline) denotes using LLaVA for fine-tuning based on standard image-label pair samples. Table 5 shows the results, from which several observations can be drawn. **(1)** LLaVA (zero-shot) attains very limited performance, which indicates that the MLLM is difficult to exert the ability of ordinal understanding in the training-free situation. **(2)** LLaVA (baseline), which is fine-tuned by image-label pairs, achieves normal performance but still leaves potential for improvement to be desired. **(3)** Compared to the baseline LLaVA, the addition of Domain Knowledge Prompts provides considerable performance gains of nearly 10% in accuracy, demonstrating the importance of task-specific modeling. **(4)** By comparing (b) and (d) in Table 5, as well as (c) and (f), we find that merely adding Category Feature Prompts without RO-CoT could not help the MLLM to improve, which we hypothesize is due to the lack of relationships of multi-stage prompts that confuses the MLLM in the absence of CoT. **(5)** Comparing (f) and (d), it shows the remarkable improvement of the MLLM brought about by RO-CoT based on Category Feature Prompts, which demonstrates that our RO-CoT can fully utilize these prompts for ordinal understanding, since RO-CoT can rigorously connect multi-stage refined prompts to build a common thinking paradigm for OR tasks. The full version (g), LLaVA with our OrderChain, achieves the best performance, proving our OrderChain’s effectiveness in commonality and specificity modeling for endowing more powerful ordinal understanding to the MLLM.

4.4. Zero-shot Generalization

To further validate that our OrderChain learned the general way of thinking for ordinal understanding, we conduct inference-only experiments to show the zero-shot performance of baseline LLaVA and our OrderChain, as shown in Tab. 6. We choose three additional datasets including two in-domain datasets (APTOS [21] dataset for DR grading and AVA dataset for aesthetic assessment [18]) and one out-of-domain dataset (KonIQ [19] for image quality assessment). As we can see, our OrderChain can bring a re-

Method	Metric	APTOS [21]	AVA [18]	KonIQ [19]
Baseline	ACC \uparrow	23.2	49.6	36.3
	MAE \downarrow	1.02	0.56	0.81
OrderChain	ACC \uparrow	70.0	57.7	51.5
	MAE \downarrow	0.30	0.44	0.49

Table 6. Zero-shot Performance of baseline LLaVA and our Orderchain on in-domain and out-of-domain datasets.

markable zero-shot improvement to the baseline LLaVA on both in-domain and out-of-domain datasets, substantiating the generalization ability of OrderChain. Specifically, the significant improvement on the APTOS dataset can be attributed to the model’s prior training on similar tasks, enabling it to better adapt to the understanding of this specific task. However, even for unseen scenarios (*i.e.* image quality assessment), our OrderChain can still obtain a 15% improvement compared to baseline LLaVA. This demonstrates that our approach indeed stimulates the ordinal understanding ability of MLLMs, equipping them with a generalized way of thinking for ordinal regression tasks.

5. Conclusions

In this paper, we presented a novel and general prompting paradigm, OrderChain, to improve the ordinal understanding of MLLMs for ordinal regression. We first pointed out two major reasons for the limited performance of vanilla MLLMs, *i.e.*, lack of specificity modeling and commonality modeling. We adopted a range optimization Chain-of-Thought to learn a commonality way of thinking about ordinal regression tasks and task-aware prompts to inject task-specific information into MLLMs. We also introduce a category recursive division tool to generate refined candidate category subsets for supporting range optimization. Extensive experiments showed that our OrderChain significantly improves the performance of the MLLM and achieves optimal performance in most ordinal regression tasks and considerable zero-shot generalization performance, demonstrating that OrderChain can effectively improve the ordinal understanding of MLLMs and provides a promising paradigm to build a unified ordinal regression MLLM.

Acknowledgements This research was partially supported by National Natural Science Foundation of China under grants No. 82202984, No. 12326612, Zhejiang Key R&D Program of China under grant No. 2024SSYS0026, and the Transvascular Implantation Devices Research Institute (TIDRI).

References

- [1] Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altschmidt, Sam Altman, Shyamal Anadkat, et al. GPT-4 technical report. *arXiv preprint arXiv:2303.08774*, 2023. 2
- [2] Jinze Bai, Shuai Bai, Yunfei Chu, Zeyu Cui, Kai Dang, Xiaodong Deng, Yang Fan, Wenbin Ge, Yu Han, Fei Huang, Binyuan Hui, Luo Ji, Mei Li, Junyang Lin, Runji Lin, Dayiheng Liu, Gao Liu, Chengqiang Lu, Keming Lu, Jianxin Ma, Rui Men, Xingzhang Ren, Xuancheng Ren, Chuanqi Tan, Sinan Tan, Jianhong Tu, Peng Wang, Shijie Wang, Wei Wang, Shengguang Wu, Benfeng Xu, Jin Xu, An Yang, Hao Yang, Jian Yang, Shusheng Yang, Yang Yao, Bowen Yu, Hongyi Yuan, Zheng Yuan, Jianwei Zhang, Xingxuan Zhang, Yichang Zhang, Zhenru Zhang, Chang Zhou, Jingren Zhou, Xiaohuan Zhou, and Tianhang Zhu. Qwen technical report. *arXiv preprint arXiv:2309.16609*, 2023. 2
- [3] Jinze Bai, Shuai Bai, Yunfei Chu, Zeyu Cui, Kai Dang, Xiaodong Deng, Yang Fan, Wenbin Ge, Yu Han, Fei Huang, et al. Qwen technical report. *arXiv preprint arXiv:2309.16609*, 2023. 1
- [4] Jinze Bai, Shuai Bai, Shusheng Yang, Shijie Wang, Sinan Tan, Peng Wang, Junyang Lin, Chang Zhou, and Jingren Zhou. Qwen-vl: A frontier large vision-language model with versatile abilities. *arXiv:2308.12966*, 2023. 1, 3
- [5] Christopher Beckham and Christopher Pal. Unimodal probability distributions for deep ordinal classification. In *International Conference on Machine Learning*, pages 411–419. PMLR, 2017. 6
- [6] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *Advances in Neural Information Processing Systems*, 33:1877–1901, 2020. 1, 3
- [7] Shixing Chen, Caojin Zhang, Ming Dong, Jialiang Le, and Mike Rao. Using ranking-CNN for age estimation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5183–5192, 2017. 3
- [8] Yi Cheng, Haochao Ying, Renjun Hu, Jinhong Wang, Wenhao Zheng, Xiao Zhang, Danny Chen, and Jian Wu. Robust image ordinal regression with controllable image generation. *arXiv preprint arXiv:2305.04213*, 2023. 6
- [9] Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, et al. Training verifiers to solve math word problems. *arXiv preprint arXiv:2110.14168*, 2021. 3
- [10] Wenliang Dai, Junnan Li, Dongxu Li, Anthony Meng Huat Tiong, Junqi Zhao, Weisheng Wang, Boyang Li, Pascale Fung, and Steven Hoi. InstructBLIP: Towards general-purpose vision-language models with instruction tuning. In *NeurIPS*, 2023. 3
- [11] Raul Diaz and Amit Marathe. Soft labels for ordinal regression. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4738–4747, 2019. 3, 5, 6, 7
- [12] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020. 6
- [13] Yao Du, Qiang Zhai, Weihang Dai, and Xiaomeng Li. Teach CLIP to develop a number sense for ordinal regression. In *European Conference on Computer Vision*, pages 1–17. Springer, 2024. 1, 3, 6, 7
- [14] Eibe Frank and Mark Hall. A simple approach to ordinal classification. In *European Conference on Machine Learning*, pages 145–156. Springer, 2001. 3
- [15] Huan Fu, Mingming Gong, Chaohui Wang, Kayhan Batmanghelich, and Dacheng Tao. Deep ordinal regression network for monocular depth estimation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2002–2011, 2018. 3
- [16] Xin Geng, Chao Yin, and Zhi-Hua Zhou. Facial age estimation by learning from label distributions. *TPAMI*, 2013. 3
- [17] Amanda Elswick Gentry, Colleen K Jackson-Cook, Debra E Lyon, and Kellie J Archer. Penalized ordinal regression methods for predicting stage of cancer in high-dimensional covariate spaces. *Cancer Informatics*, 14:CIN-S17277, 2015. 1
- [18] Chunhui Gu, Chen Sun, David A Ross, Carl Vondrick, Caroline Pantofaru, Yeqing Li, Sudheendra Vijayanarasimhan, George Toderici, Susanna Ricco, Rahul Sukthankar, et al. AVA: A video dataset of spatio-temporally localized atomic visual actions. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6047–6056, 2018. 8
- [19] Vlad Hosu, Hanhe Lin, Tamas Sziranyi, and Dietmar Saupe. KonIQ-10k: An ecologically valid database for deep learning of blind image quality assessment. *IEEE Transactions on Image Processing*, 29:4041–4056, 2020. 8
- [20] Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, Weizhu Chen, et al. LoRA: Low-rank adaptation of large language models. *ICLR*, 1(2):3, 2022. 3, 5, 6
- [21] Maggie Karthik and Sohier Dane. APTOS 2019 blindness detection. *Kaggle <https://kaggle.com/competitions/aptos2019-blindness-detection> Go to reference in*, 5, 2019. 8
- [22] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. 6
- [23] Shu Kong, Xiaohui Shen, Zhe Lin, Radomir Mech, and Charles Fowlkes. Photo aesthetics ranking network with

- attributes and content adaptation. In *European Conference on Computer Vision*, pages 662–679. Springer, 2016. 1
- [24] Jun-Tae Lee and Chang-Su Kim. Image aesthetic assessment based on pairwise comparison a unified approach to score regression, binary classification, and personalization. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1191–1200, 2019. 1
- [25] Seon-Ho Lee and Chang-Su Kim. Deep repulsive clustering of ordered data based on order-identity decomposition. In *International Conference on Learning Representations*, 2020. 3
- [26] Gil Levi and Tal Hassner. Age and gender classification using convolutional neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pages 34–42, 2015. 1, 6
- [27] Lei Li. CPSeg: Finer-grained image semantic segmentation via chain-of-thought language prompting. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 513–522, 2024. 3
- [28] Ling Li and Hsuan-Tien Lin. Ordinal regression by extended binary classification. *Advances in Neural Information Processing Systems*, 19, 2006. 3
- [29] Wanhua Li, Xiaoke Huang, Jiwen Lu, Jianjiang Feng, and Jie Zhou. Learning probabilistic ordinal embeddings for uncertainty-aware regression. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13896–13905, 2021. 1, 3, 5, 6, 7
- [30] Wanhua Li, Xiaoke Huang, Zheng Zhu, Yansong Tang, Xiu Li, Jie Zhou, and Jiwen Lu. OrdinalCLIP: Learning rank prompts for language-guided ordinal regression. *Advances in Neural Information Processing Systems*, 35:35313–35325, 2022. 1, 3, 5, 6, 7
- [31] Kyungsun Lim, Nyeong-Ho Shin, Young-Yoon Lee, and Chang-Su Kim. Order learning and its application to age estimation. In *International Conference on Learning Representations*, 2019. 3
- [32] Haotian Liu, Chunyuan Li, Yuheng Li, and Yong Jae Lee. Improved baselines with visual instruction tuning. *arXiv:2310.03744*, 2023. 3
- [33] Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. Visual instruction tuning. In *NeurIPS*, 2023. 3
- [34] Haotian Liu, Chunyuan Li, Yuheng Li, and Yong Jae Lee. Improved baselines with visual instruction tuning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 26296–26306, 2024. 1, 5, 6, 7
- [35] Haotian Liu, Chunyuan Li, Yuheng Li, Bo Li, Yuanhan Zhang, Sheng Shen, and Yong Jae Lee. LLaVA-NeXT: Improved reasoning, OCR, and world knowledge, 2024. 3
- [36] Xiaofeng Liu, Yang Zou, Yuhang Song, Chao Yang, Jane You, and BV K Vijaya Kumar. Ordinal regression with neuron stick-breaking for medical diagnosis. In *Proceedings of the European Conference on Computer Vision (ECCV) Workshops*, pages 0–0, 2018. 1
- [37] Yanzhu Liu, Adams Wai-Kin Kong, and Chi Keong Goh. Deep ordinal regression based on data relationship for small datasets. In *IJCAI*, pages 2372–2378, 2017. 3
- [38] Yanzhu Liu, Adams Wai Kin Kong, and Chi Keong Goh. A constrained deep neural network for ordinal regression. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 831–839, 2018. 1, 3, 5, 6, 7
- [39] Yanzhu Liu, Fan Wang, and Adams Wai Kin Kong. Probabilistic deep ordinal regression based on Gaussian processes. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5301–5309, 2019. 1, 3, 5, 6
- [40] OpenAI. GPT-4V(ision) system card. https://cdn.openai.com/papers/GPTV_System_Card.pdf, 2023. 1, 3
- [41] Frank Palermo, James Hays, and Alexei A Efros. Dating historical color images. In *European Conference on Computer Vision*, pages 499–512. Springer, 2012. 6
- [42] Bowen Pan, Shangfei Wang, and Qisheng Jiang. Image aesthetic assessment assisted by attributes through adversarial learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 679–686, 2019. 1
- [43] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *ICML*, 2021. 3
- [44] Vadim Ratner, Yoel Shoshan, and Tal Kachman. Learning multiple non-mutually-exclusive tasks for improved classification of inherently ordered labels. *arXiv preprint arXiv:1805.11837*, 2018. 6
- [45] Machel Reid, Nikolay Savinov, Denis Teplyashin, Dmitry Lepikhin, Timothy Lillicrap, Jean-baptiste Alayrac, Radu Soricut, Angeliki Lazaridou, Orhan Firat, Julian Schrittwieser, et al. Gemini 1.5: Unlocking multimodal understanding across millions of tokens of context. *arXiv preprint arXiv:2403.05530*, 2024. 3
- [46] Rasmus Rothe, Radu Timofte, and Luc Van Gool. Deep expectation of real and apparent age from a single image without facial landmarks. *IJCV*, 2018. 3
- [47] Nyeong-Ho Shin, Seon-Ho Lee, and Chang-Su Kim. Moving window regression: A novel approach to ordinal regression. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18760–18769, 2022. 1, 3, 5, 6
- [48] Gemini Team, Rohan Anil, Sebastian Borgeaud, Jean-Baptiste Alayrac, Jiahui Yu, Radu Soricut, Johan Schalkwyk, Andrew M Dai, Anja Hauth, Katie Millican, et al. Gemini: A family of highly capable multimodal models. *arXiv preprint arXiv:2312.11805*, 2023. 1, 3
- [49] Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. LLaMA: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*, 2023. 1, 2
- [50] Jinhong Wang, Yi Cheng, Jintai Chen, TingTing Chen, Danny Chen, and Jian Wu. Ord2Seq: Regarding ordinal regression as label sequence prediction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5865–5875, 2023. 3, 5, 6, 7

- [51] Rui Wang, Peipei Li, Huaibo Huang, Chunshui Cao, Ran He, and Zhaofeng He. Learning-to-rank meets language: Boosting language-driven ordering alignment for ordinal classification. *Advances in Neural Information Processing Systems*, 36, 2023. [1](#), [3](#), [5](#), [6](#), [7](#)
- [52] Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. Chain-of-thought prompting elicits reasoning in large language models. *Advances in Neural Information Processing Systems*, 35:24824–24837, 2022. [3](#)
- [53] Yixuan Wu, Yizhou Wang, Shixiang Tang, Wenhao Wu, Tong He, Wanli Ouyang, Philip Torr, and Jian Wu. DetToolChain: A new prompting paradigm to unleash detection ability of MLLM. In *European Conference on Computer Vision*, pages 164–182. Springer, 2024. [1](#), [3](#)
- [54] Qinghao Ye, Haiyang Xu, Guohai Xu, Jiabo Ye, Ming Yan, Yiyang Zhou, Junyang Wang, Anwen Hu, Pengcheng Shi, Yaya Shi, et al. mPLUG-Owl: Modularization empowers large language models with multimodality. *arXiv preprint arXiv:2304.14178*, 2023. [3](#)
- [55] Yuqian Yuan, Wentong Li, Jian Liu, Dongqi Tang, Xinjie Luo, Chi Qin, Lei Zhang, and Jianke Zhu. Osprey: Pixel understanding with visual instruction tuning. In *CVPR*, 2024. [1](#)
- [56] Yuqian Yuan, Wentong Li, Jian Liu, Dongqi Tang, Xinjie Luo, Chi Qin, Lei Zhang, and Jianke Zhu. Osprey: Pixel understanding with visual instruction tuning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 28202–28211, 2024. [3](#)