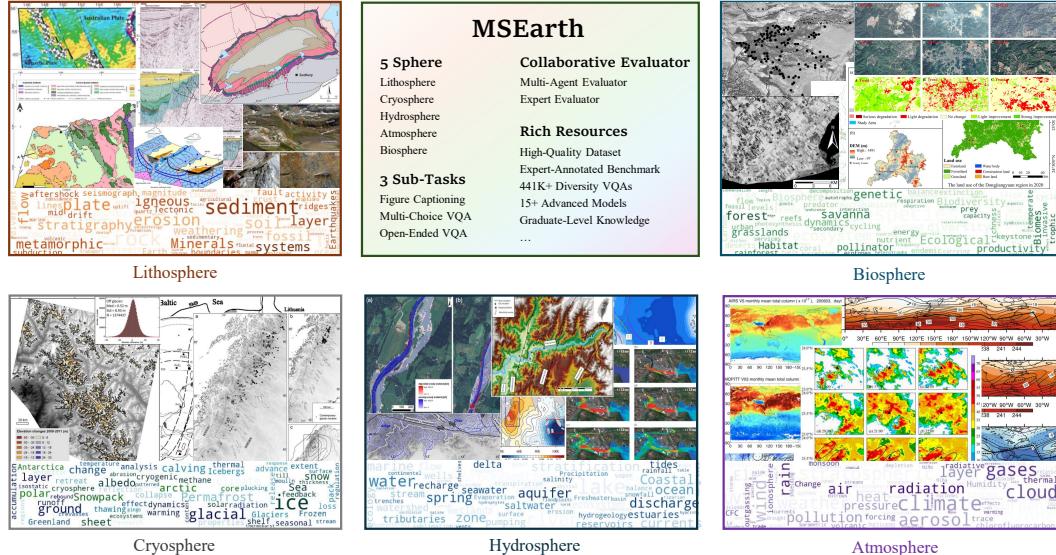


# 000 001 002 003 004 005 006 007 008 009 010 011 012 013 014 015 016 017 018 019 020 021 022 023 024 025 026 027 028 029 030 031 032 033 034 035 036 037 038 039 040 041 042 043 044 045 046 047 048 049 050 051 052 053 MSEARTH: A MULTIMODAL SCIENTIFIC DATASET AND BENCHMARK FOR PHENOMENA UNCOVERING IN EARTH SCIENCE

006  
007  
008  
009  
010  
011  
012  
013  
014  
015  
016  
017  
018  
019  
020  
021  
022  
023  
024  
025  
026  
027  
028  
029  
030  
031  
032  
033  
034  
035  
036  
037  
038  
039  
040  
041  
042  
043  
044  
045  
046  
047  
048  
049  
050  
051  
052  
053  
**Anonymous authors**  
Paper under double-blind review



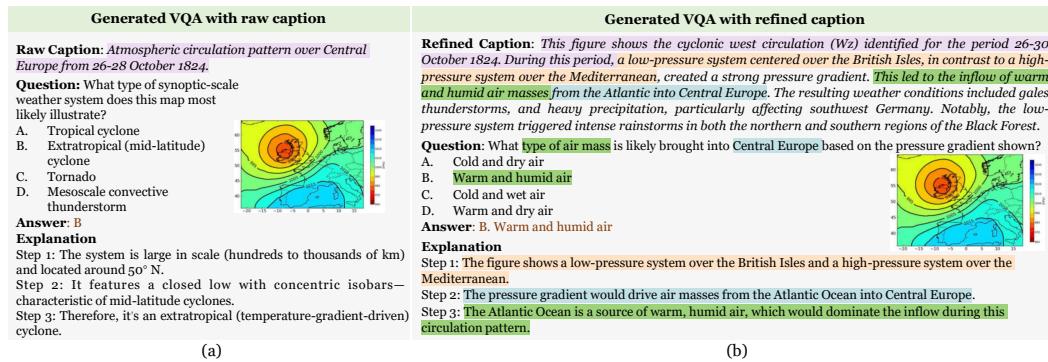
029  
030  
031  
032  
033  
034  
035  
036  
037  
038  
039  
040  
041  
042  
043  
044  
045  
046  
047  
048  
049  
050  
051  
052  
053  
Figure 1: Illustrative examples of the diverse types of scientific figures in MSEarth, sourced from open-access articles available from website.

## ABSTRACT

The rapid advancement of multimodal large language models (MLLMs) has unlocked new opportunities to tackle complex scientific challenges. Despite this progress, their application in addressing earth science problems, especially at the graduate level, remains underexplored. A significant barrier is the absence of benchmarks that capture the depth and contextual complexity of geoscientific reasoning. Current datasets and benchmarks often rely on synthetic datasets or simplistic figure-caption pairs, which do not adequately reflect the intricate reasoning and domain-specific insights required for real-world scientific applications. To address these gaps, we introduce MSEarth, a multimodal scientific dataset and benchmark curated from high-quality, open-access scientific publications. MSEarth encompasses the five major spheres of Earth science: atmosphere, cryosphere, hydrosphere, lithosphere, and biosphere, featuring over 289K figures with refined captions. These captions are crafted from the original figure captions and enriched with discussions and reasoning from the papers, ensuring the benchmark captures the nuanced reasoning and knowledge-intensive content essential for advanced scientific tasks. MSEarth supports a variety of tasks, including scientific figure captioning, multiple choice questions, and open-ended reasoning challenges. By bridging the gap in graduate-level benchmarks, MSEarth provides a scalable and high-fidelity resource to enhance the development and evaluation of MLLMs in scientific reasoning. The benchmark is publicly available to foster further research and innovation in this field. Resources related to this benchmark can be found at this anonymous link: <https://anonymous.4open.science/r/MSEarth-2B3F>.

## 054 1 INTRODUCTION

056 The advent of multimodal large language models (MLLMs) (Liu et al., 2023a; Liang et al., 2024b) has  
 057 revolutionized artificial intelligence, driving groundbreaking advancements across diverse scientific  
 058 disciplines. Prominent examples include ChemVLM (Li et al., 2025) in chemistry, GeoChat (Kuck-  
 059 reja et al., 2024) in geography, and WeatherQA (Ma et al., 2024) in atmospheric science. These  
 060 models excel in domain-specific visual question answering by integrating specialized knowledge.  
 061 For example, ChemVLM facilitates the analysis of molecular structures, chemical reactions, and  
 062 chemistry-related examination questions, while WeatherQA enables reasoning about severe weather  
 063 events in real-world scenarios.



072 Figure 2: Illustration of VQA generation methodologies: (a) VQA relying exclusively on figure  
 073 captions, and (b) VQA utilizing refined captions that integrate figure captions with content from  
 074 academic papers. Highlighted areas denote questions and answers supported by evidence.

076 Developing multimodal large language models (MLLMs) that understand advanced geoscientific  
 077 knowledge necessitates rigorous datasets and benchmarks (D&Bs) to enhance their ability to solve  
 078 complex, discipline-specific problems. Existing datasets and benchmarks often rely on synthetic  
 079 datasets or materials from high school and undergraduate textbooks (Lu et al., 2022; Yue et al.,  
 080 2024), which lack the depth required for professional, graduate-level tasks. Recent efforts (Li  
 081 et al., 2024c; Roberts et al., 2024; Li et al., 2024b) have turned to academic papers for constructing  
 082 multimodal scientific benchmarks, leveraging the complexity of graduate-level content. Yet, these  
 083 approaches typically extract only figures and captions, neglecting the critical scientific reasoning  
 084 and insights in the paper context. Consequently, tasks based on these benchmarks tend to be  
 085 oversimplified, focusing on basic *figure-caption matching*, which offers limited insight into a model’s  
 086 reasoning capabilities. Another obstacle to advancing current datasets and benchmarks is the difficulty  
 087 of designing questions that rigorously evaluate MLLMs’ data-analysis abilities to uncover Earth-  
 088 science phenomena from observational imagery. In scientific papers, images depicting observed  
 089 phenomena are often accompanied by information that reveals scientific hypotheses, supporting  
 090 evidence, analyses, and conclusions—critical insights that are typically embedded within the main  
 091 text rather than the figure captions alone. As shown in Figure 2 (a), question generation based  
 092 solely on captions oversimplifies tasks, as the quality and difficulty of questions are constrained  
 093 by the limitations of the generation models, and they lack the contextual support from the papers,  
 094 making verification challenging. Existing benchmarks focus solely on figure-caption matching  
 095 while neglecting the high knowledge density in scientific reasoning processes. This raises a new  
 096 challenge: *How to effectively align high-value Earth science images with long-context information  
 097 for phenomena uncovering?*

098 To address this challenge, we propose a novel approach to D&B construction that overcomes the  
 099 limitations of existing methods through two key innovations. First, we introduce the concept of the  
 100 *refined caption*. In scientific papers, observational images typically serve as visual representations of  
 101 phenomena, while the deeper scientific insights—such as hypotheses, supporting evidence, analytical  
 102 reasoning, and conclusions—are often embedded within the main body of the text. Raw captions  
 103 provide only brief descriptions and lack the necessary context for complex reasoning. The refined  
 104 caption bridges this gap by combining the raw caption with relevant, domain-specific information  
 105 extracted from the paper, resulting in a more comprehensive and scientifically meaningful description

108 of the image. As illustrated in Figure 2 (b), questions generated using refined captions not only  
 109 exhibit higher quality but are also supported and validated by the content of the academic papers,  
 110 ensuring the professionalism and accuracy of the generated questions. Second, we implement a  
 111 rigorous quality control process to validate the generated benchmark. This process combines a  
 112 multi-agent automated evaluation and expert evaluation, ensuring that the generated questions are  
 113 relevant, coherent, and aligned with the complexity of professional-level geoscientific reasoning. By  
 114 addressing the reliability and accuracy issues faced by current LLMs in scientific question generation,  
 115 our approach ensures that the benchmark is both robust and valuable for evaluating MLLMs in  
 116 advanced scientific domains.

117 Building on our adaptive annotation methodology, we present MSEarth, a comprehensive multimodal  
 118 D&B for graduate-level Earth science. This D&B is derived from 64,560 open-access publications  
 119 spanning five spheres, eight subjects, and 66 sub-subjects, from which we extract 289,891 figures.  
 120 We enhance these figures with “refined captions” that average 136.29 tokens in length, significantly  
 121 expanding from the original average of 37.56 tokens. The test set unifies scientific figure captioning  
 122 with multiple-choice and open-ended reasoning tasks, providing a holistic, interdisciplinary evaluation  
 123 framework. This rigorously assesses MLLMs within professional-level geoscientific contexts, filling  
 124 a critical gap in graduate-level multimodal benchmarks. Furthermore, MSEarth introduces a scalable,  
 125 high-fidelity pipeline for constructing domain-specific scientific D&B, establishing the test set as  
 126 a robust tool for advancing MLLMs in real-world Earth science applications. Leveraging this  
 127 framework, we produce a resource-rich training dataset covering captioning, open-ended QA, and  
 128 multiple-choice tasks for post-training (e.g., instruction tuning and GRPO-based reinforcement  
 129 learning). This training data delivers substantial improvements across a wide range of tasks for  
 130 open-source models. In summary, our contributions are as follows:

131 **Development of a Scalable Adaptive Framework:** We introduced an innovative semi-automated  
 132 tool, enabling the automatic generation and machine-assisted filtering of VQA tasks. This provides  
 133 a robust, scalable solution for creating high-fidelity, domain-specific D&B that can be extended to  
 134 other scientific fields.

135 **High-Quality D&B Resources for Earth Science MLLMs:** We present an expert-annotated bench-  
 136 mark for graduate-level Earth science, along with a diverse training corpus—including captioning,  
 137 open-ended QA, and multiple-choice tasks—to support advanced post-training of MLLMs.

138 **Comprehensive Evaluation and Validation of State-of-the-Art MLLMs:** Through extensive  
 139 evaluations of MLLMs on MSEarth, we not only provide crucial insights into current limitations and  
 140 future research directions but also demonstrate the effectiveness of our training data by developing a  
 141 state-of-the-art baseline model.

## 143 2 RELATED WORK

144 **Multimodal Scientific Datasets and Benchmarks.** Numerous multimodal D&Bs have been devel-  
 145 oped to evaluate scientific understanding across various domains. These benchmarks often integrate  
 146 text, images, and other modalities to assess models’ reasoning and cross-modal capabilities. However,  
 147 their creation typically requires significant manual effort in data collection and validation. Sci-  
 148 enceQA (Lu et al., 2022) is an early multimodal benchmark that features multiple-choice questions  
 149 (MCQs) collected from online resources and manually filtered for quality. It covers general science  
 150 topics such as physics, chemistry, and biology, with a focus on elementary and high school-level  
 151 reasoning. SceMQA (Liang et al., 2024a) and Mmmu (Liang et al., 2024a) extended this by incorpo-  
 152 rating both MCQs and open-ended questions (OE) from textbooks and online resources, targeting  
 153 pre-college and college-level difficulty. OlympiadBench (He et al., 2024a) introduced competition-  
 154 level problems in mathematics and physics, offering open-ended tasks sourced from Olympiad exams.  
 155 These problems are highly challenging but limited to specific domains. More recently, EMMA (Hao  
 156 et al., 2025) combined manually designed questions with existing benchmarks, covering a broader  
 157 range of topics with mixed difficulty levels. In contrast, our objective is to enhance models’ ability to  
 158 comprehend multimodal, complex scientific problems—drawn from high-quality research papers in  
 159 earth science—that demand graduate-level, domain-specific expertise.

160 **Paper-Based Multimodal Scientific Datasets and Benchmarks.** D&Bs based on academic papers  
 161 aim to leverage the rich, domain-specific content found in scientific literature. FigureSeer (Siegel

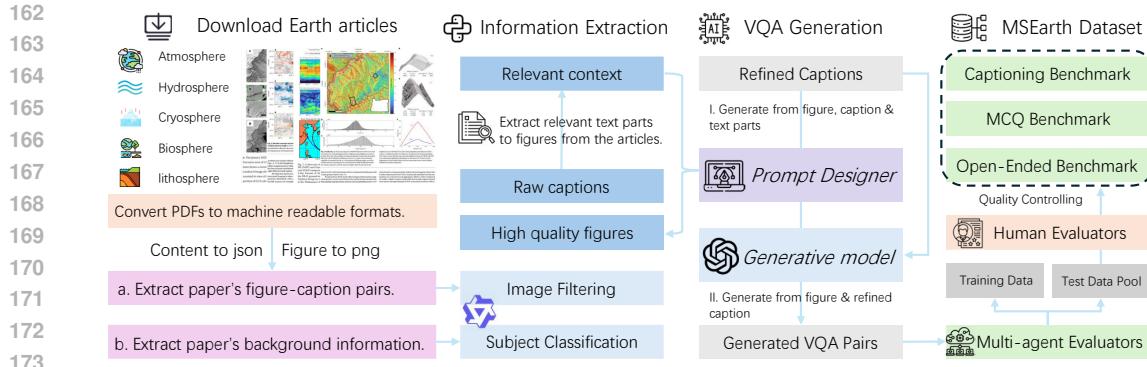


Figure 3: Data curation process for MSEarth. The two parts on the left represent data preprocessing, while the two parts on the right encompass the automated generation of VQA and expert-AI collaborative filtering.

et al., 2016) first extracts figures from academic papers, focusing on chart figures to evaluate the understanding of chart figures. SciFiBench (Roberts et al., 2024) extended this by introducing figure-to-caption and caption-to-figure matching tasks, while MMSci (Li et al., 2024c) further advanced this approach using figures from Nature papers. However, these benchmarks lack original questions, limiting their ability to assess advanced reasoning and contextual understanding. ArxivQA/Cap (Li et al., 2024b) expanded the scope by generating new questions for figures from 32 subjects on arXiv. However, these questions were generated solely using the inherent capabilities of GPT-4V (Achiam et al., 2023) and did not have contextual support from the relevant text in the papers, raising concerns about their scientific validity. In contrast, our proposed benchmark, MSEarth, addresses these limitations by introducing original, evidence-supported questions grounded in refined captions. This approach enables a rigorous evaluation of MLLMs in professional-level geoscientific applications.

### 3 MSEARTH - A MULTIMODAL SCIENTIFIC D&B FOR EARTH SCIENCE

This section provides a detailed overview of the construction process for MSEarth. As illustrated in Figure 3, we outline the framework used to develop MSEarth from open-access scientific publications. The section is organized into three main parts: first, we detail the data collection and preprocessing steps. Next, we elaborate on the construction procedures for the two D&B within MSEarth, namely MSEarthCap and MSEarthQA. Finally, we describe the process of ensuring the reliability of the test data, which involves expert annotation and manual screening of the sampled test data.

#### 3.1 DATA PREPARATION

The first part of the D&B construction focuses on data collection and preprocessing. The data collection begins with more than 400K Earth science papers obtained in PDF format. These are uniformly converted into structured JSON text using the MinerU (Wang et al., 2024) parser. To classify the papers, semantic similarity is calculated between the abstracts and keywords from the five Earth spheres: hydrosphere, biosphere, lithosphere, atmosphere, and cryosphere. Based on this, the papers are assigned to respective disciplinary categories. Details are provided in Appendix B.3. We then selected papers based on the criterion of containing high-quality, Earth science-related images, resulting in a subset of around 83k papers. Specifically, Qwen-2.5-VL-72B (Bai et al., 2025b) is utilized to filter and select images, with the filtering prompts detailed in the Appendix B.4.

#### 3.2 MSEARTHCap

**Figure-Caption Extraction:** Figures and their corresponding captions are extracted from the JSON files processed by MinerU. As shown in Appendix B.2, MinerU has already extracted the figures along with their original captions, which can be directly utilized for subsequent processing. To ensure accurate alignment between figures and their references within the text, we employ a regex-based method to identify the labels of each figure. This approach enables precise matching between the figures and the relevant sections of the articles.

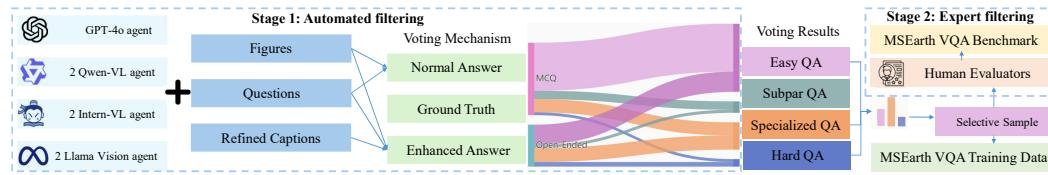


Figure 4: Overall approach of our multi-agent, voting-based approach to automate the validation of generated questions.

**Relevant Context Extraction:** To further enrich the captions with contextual information, we use the figure labels obtained in the previous step to perform approximate matching against the main body of the paper. Since MinerU processes papers with segmented paragraphs, we apply regular expression matching to each paragraph to extract contextual text that references the target figure. This ensures the inclusion of descriptions and reasoning associated with each figure within the paper. To guarantee that the extracted context provides sufficient detail about the target figure, only paragraphs exceeding two sentences were included in the final dataset. From this filtered subset, we selected around 64K papers that met the criteria for subsequent processing. For more details, refer to Appendix B.5.

**Refined Caption Generation:** To create professional-level figure descriptions, we employ GPT-4o for refined caption generation. The model takes as input the extracted figure, its original caption, and the contextual text from the relevant sections of the paper. Figure refinement is performed only for data that includes valid relevant contexts. The specific prompts used for this process are detailed in the Appendix B.6. After statistical analysis, we observe that the average word length of the raw captions is 37.56, while the average length of the refined captions increases to 136.29, reflecting the incorporation of richer, domain-specific content.

### 3.3 MSEARTHQA

To generate high-quality multiple-choice questions (MCQs) and open-ended questions, we use a question generation pipeline that takes the figure, its original caption, and the refined caption as input. The generation prompts, detailed in the Appendix B.6, are crafted to encourage the model to highlight differences between the original and refined captions, ensuring that the generated questions are grounded in evidence from the paper. The questions are constructed using GPT-4o to maintain high detail and relevance. However, due to inherent challenges such as self-inconsistency and uncertainty in the generation process, the generated questions undergo an automated and expert validation process to ensure quality.

#### 3.3.1 AUTOMATED VALIDATION

Inspired by the LLM Voting (Yang et al., 2024b; Lee et al., 2025; Kaesberg et al., 2025) method, we developed a multi-agent, voting-based approach to automate the validation of generated questions. Specifically, we employ a *Majority Voting* strategy, where multiple agents independently generate responses, and the final decision is based on the majority consensus of these agents. In our setup, we utilize the following MLLMs for decision-making: Qwen2.5-VL-72B, Qwen2.5-VL-7B, InternVL2.5-7B, InternVL2.5-78B, and GPT-4o. A key aspect of our evaluation process is the use of the *refined caption*, which incorporates scientists' reasoning and insights about the figure extracted from the paper. This refined caption provides additional context and domain-specific information that goes beyond the original caption. By comparing model performance with and without the refined caption, we can assess the quality of the questions and determine whether they test the model's ability to grasp scientific reasoning and insights. The detailed decision-making process is outlined below:

**Phase A:** The question and original caption are provided to a suite of models  $\{M_1, M_2, \dots, M_n\}$ . The correctness of the model responses is used to evaluate the types and quality of the questions. A threshold of 60% is defined for supermajority voting. Specifically, if more than 40% of the models produce incorrect responses, the question is flagged for further analysis. Additionally, we discard questions that all models answer correctly, as these questions do not contribute to the effective testing or training of the models. Questions identified through this process are categorized as either potentially difficult or of poor quality, with the distinction made in subsequent phases.

**Phase B:** In this phase, the question and refined caption are provided to the same suite of models. If more than 60% of the models answer the question correctly with the refined caption, it indicates that the question requires relatively specialized scientific knowledge to answer. Such questions are categorized as *specialized QA*, as their answers rely on the model's ability to understand and apply specific domain knowledge rather than simply perceiving the image or relying on commonsense reasoning. Questions that fail this phase proceed to the next stage for further evaluation.

**Phase C:** In this phase, only models with 70B+ parameters are used for voting, including GPT-4o (the same model used for question generation), InternVL2.5-78B, and Qwen2.5-VL-72B. The question and its refined caption are provided to these large-scale models. If more than 60% of the large models answer the question correctly, it suggests that the difficulty of the question likely lies in the model’s ability to perceive and interpret the image content. Such questions are categorized as *hard QA*. Subsequent human validation will involve sampling and additional annotation across QA filtered in all phases to ensure the overall quality and accuracy of the benchmark.

This pipeline identifies high-quality questions by filtering out overly simplistic or poorly constructed ones. In **Phase A**, approximately 70% of the questions were categorized as *easy*, as most models could answer them correctly without refined captions. After **Phase B**, around 20% were classified as *specialized QA*, where refined captions enabled correct answers, indicating the need for domain-specific knowledge. In **Phase C**, 5% were labeled as *hard QA*, requiring high-performing models to interpret image content accurately, suggesting that these questions test the model’s ability to perceive and interpret image content. The remaining 5% were deemed flawed and discarded. Detailed processes and examples are provided in Appendix C. For the training dataset, we sampled more than 150K VQA pairs, including both multiple-choice and open-ended questions. Of these, 20% were drawn from Phase A and 80% from Phase B, ensuring a balanced distribution of question difficulty.

Table 1: Main statistics in MSEarth-Bench.

Statistic	Number
Total subjects	66
Total articles	64,560
Total figures	289,891
Total questions	448,980
Training Set	441,785
Captioning as QA	289,891
MCQ	102,753
Open-Ended QA	49,141
Test Set	7,195
Captioning as QA	3,000
MCQ	2,784
* Reasoning Question	2117 (76.0%)
* Perception Question	667 (24.0%)
Open-Ended QA	1,411
Average caption length	37.56
Average refined caption length	136.29

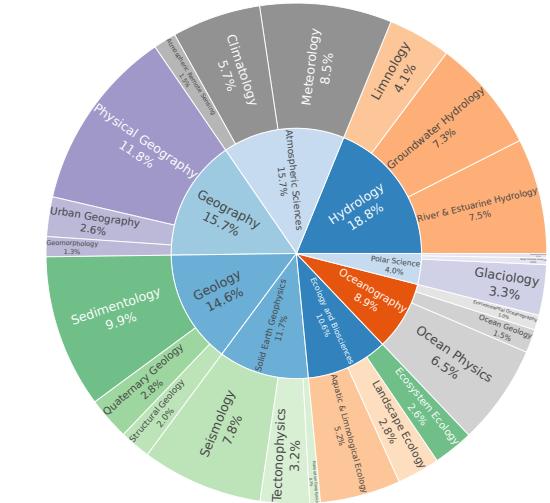


Figure 5: Subjects distribution in MSEarth.

### 3.3.2 EXPERT VALIDATION

Ensuring that synthetic data closely mirrors real-world distributions is critical for evaluation tasks. To achieve this, domain experts are engaged to review and annotate the curated QA pairs for accuracy and relevance. During this process, low-quality or invalid questions are identified and filtered out to ensure the overall quality of the dataset. The annotation process is conducted from two perspectives: image types and question types. For image types, we categorize the data into three groups: single-image question answering, single-image-focused question answering within multi-image figures, and multi-image relational question answering. For question types, we define two categories: scientific reasoning and perception questions. Scientific reasoning questions are constructed based on inferences or scientific discoveries presented in research papers, making them more specialized and challenging. In contrast, image perception questions focus on interpreting images and require less background knowledge of scientific concepts. This expert-AI collaborative process, combined with rigorous quality control, results in a high-quality dataset comprising 1,500 open-ended questions and 3,000

324 MCQs, forming the MSEarthQA benchmark. The annotated results are summarized in Table 1, with  
 325 further details and analysis provided in Appendix D and Table 10.  
 326

## 327 4 EXPERIMENTS

### 329 4.1 EVALUATED MODELS

331 We evaluate different families of MLLMs on our benchmark. We evaluate the following closed-  
 332 source models: GPT-4 series (Hurst et al., 2024), Gemini-2.5 series (Team et al., 2023) and  
 333 Claude-3 series (Anthropic, 2024). We also evaluate the following open-source models: LLaVA-  
 334 OneVision (Li et al., 2024a), Qwen-2.5-VL (Yang et al., 2024a), InternVL2.5/3/S1 (Chen et al., 2024;  
 335 Zhu et al., 2025; Bai et al., 2025a) and Llama-3.2-Vision-Instruct (Grattafiori et al., 2024). We use  
 336 chat/instruction-tuned variants of each model and compare the performance of multiple model sizes  
 337 where available. Details can be found in Appendix F. To validate the effectiveness of our training  
 338 data, we conducted post-training on the Intern-s1-mini and Qwen-2.5-VL-7b models. Specifically,  
 339 we employed instruct tuning method for fine-tuning on captioning and open-ended QA tasks. For the  
 340 MCQ task, we applied GRPO (Shao et al., 2024) reinforcement learning method.  
 341

### 342 4.2 EVALUATION METRICS

343 Both captioning and open-ended QA tasks require generating freeform textual outputs grounded  
 344 in complex scientific data and reasoning. To evaluate these tasks, we use lexical overlap-based  
 345 metrics such as BLEU (Papineni et al., 2002), ROUGE (Lin, 2004), and METEOR (Banerjee & Lavie,  
 346 2005) for surface-level similarity, while BERTScore (Zhang et al., 2019) assesses deeper semantic  
 347 alignment. Given the importance of factual correctness, we employ a factual entailment classifier  
 348 to measure the consistency of generated outputs with reference answers or captions. Additionally,  
 349 following G-Eval (Liu et al., 2023b), we utilize the Qwen2.5-VL-72B model with a specialized  
 350 prompt to compute a factual scientific score. For the captioning task, we define a Cap-Eval score  
 351 ranging from 1 to 5, where higher scores indicate better caption quality. For the open-ended QA task,  
 352 we introduce OE-Eval, which evaluates the reasonableness of generated answers using a binary 0/1  
 353 scoring system.  
 354

Evaluating MCQs is relatively straightforward, as these tasks require selecting the correct answer  
 from a predefined set of options. In our experiments, models were guided by carefully structured  
 prompts to ensure adherence to a specific output format. Regular expression rules were employed to  
 extract the selected choice, ensuring strict alignment with the predefined answer format. We observed  
 that some models occasionally failed to consistently follow the formatting instructions. To address  
 this issue and preserve the integrity of the evaluation process, we implemented a similarity-based  
 method to identify the closest matching option when the selected choice could not be extracted using  
 regular expressions. Detailed evaluation metrics are provided in Appendix G.  
 361

### 362 4.3 MAIN RESULTS

364 **Current models struggle with scientific question-answering tasks, particularly on questions**  
 365 **requiring specialized knowledge and reasoning across multiple images.** The performance of  
 366 MCQs is summarized in Table 2. The results reveal that most models do not perform exceptionally  
 367 well on scientific question-answering tasks, with proprietary models generally achieving better  
 368 results. Further analysis of the models' failure rates on reasoning and perception-based questions  
 369 is provided in the Appendix (Figure 8). This analysis shows that models are more prone to errors  
 370 on questions requiring specialized knowledge, underscoring significant room for improvement in  
 371 scientific reasoning question-answering. In contrast, for relatively simpler perception-based questions,  
 372 which require less domain-specific knowledge, the models tend to perform better. Similarly, when  
 373 analyzing performance across different image types, we observe that most models achieve their best  
 374 results on tasks involving single-image inputs. However, for tasks requiring multi-image inputs,  
 375 particularly those that demand reasoning across multiple images to derive an answer, the models  
 376 perform the worst. Additional experimental results can be found in Appendix H.  
 377

**Proprietary models consistently outperform open-source models in both Scientific Figure  
 Captioning and Open-Ended VQA tasks, with LLM-based metrics providing a more reasonable**

378  
379  
380  
381 Table 2: Accuracies (%) of different models on multiple-choice questions. The best results are  
382 highlighted in bold, with the second-best underlined.  
383

Model	Image-Type			Task Type		Overall ACC
	SINGLE	MULTI	CROSS	REASONING	PERCEPT	
<b><i>Open-source Models</i></b>						
LLaVA-onvision-72B	53.55	49.48	47.95	46.58	65.52	51.11
Qwen2.5-VL-7B	47.65	44.07	37.53	40.53	58.47	44.83
Qwen2.5-VL-32B	52.59	46.99	43.84	42.47	70.16	49.10
Qwen2.5-VL-72B	52.11	50.43	46.30	44.40	70.46	50.65
InternVL2-8B	44.86	43.99	38.36	38.97	58.47	43.64
InternVL2.5-78B	53.23	49.74	44.38	43.17	74.21	50.61
InternVL3-78B	57.53	51.37	45.48	47.00	73.61	53.38
Llama3.2-90B-Vision	45.98	40.46	38.90	38.26	56.97	42.74
DeepSeek-VL2	52.43	49.23	44.66	46.06	62.82	50.07
Intern-S1-mini	60.00	57.22	49.04	51.68	75.56	58.12
Intern-S1	67.01	65.62	64.11	61.22	79.61	65.63
MiMo-VL-7B-RL	51.31	47.51	32.33	42.70	61.62	47.23
GLM-4.1V-Thinking	54.18	49.83	38.36	46.29	62.97	50.29
Qwen-7B-MSEarth	57.61	52.75	45.20	50.68	64.32	53.95
Intern-S1-mini-MSEarth	<b>65.49</b>	<b>62.97</b>	<b>58.63</b>	<b>58.81</b>	<u>78.56</u>	<b>63.54</b>
<b><i>Proprietary Models</i></b>						
Gemini-2.5-Flash	58.33	54.55	53.42	49.98	75.56	56.11
Gemini-2.5-Flash-Thinking	60.64	54.64	53.70	51.35	75.86	57.22
Gemini-2.5-Pro-Thinking	64.78	59.36	55.34	<u>56.31</u>	77.06	61.28
Claude-3.5-Haiku	49.48	47.16	42.47	42.18	64.77	47.59
Claude-3.7-Sonnet	59.52	56.53	<u>57.53</u>	51.68	78.11	58.01
GPT-4o-mini	52.51	48.63	43.01	43.65	68.67	49.64
GPT-4o	63.03	55.76	47.67	50.45	<b>81.86</b>	57.97

405  
406 Table 3: Performance on scientific figure captioning. The best results are highlighted in bold, with  
407 the second-best underlined.  
408

Model	ROUGE1	ROUGE2	Overlap		BLEU	Similarity BERTSCORE	MLLM CAP-EVAL
			ROUGEL	METEOR			
<b><i>Open-source Models</i></b>							
LLaVA-onevision-72B	30.81	5.99	<b>17.82</b>	18.35	2.15	83.47	2.07
Qwen2.5-VL-7B	27.01	5.21	15.72	16.57	1.55	83.82	2.22
Qwen2.5-VL-32B	29.90	5.88	15.53	<u>25.81</u>	2.14	82.86	2.66
Qwen2.5-VL-72B	30.68	5.85	16.53	21.35	2.36	83.87	2.56
InternVL2.5-8B	29.07	5.16	16.71	19.28	1.58	83.25	1.91
InternVL2.5-78B	30.66	5.95	17.24	20.66	2.27	83.56	2.30
InternVL3-78B	30.73	5.87	16.95	20.95	2.32	83.72	2.43
Intern-S1-mini	31.05	7.13	17.42	24.18	2.26	83.65	2.65
Llama3.2-90B-Vision	19.93	4.32	12.98	21.21	1.49	78.89	1.82
DeepSeek-VL2	29.06	5.42	16.71	18.22	1.69	83.64	2.22
Intern-S1	33.46	7.37	17.93	27.97	3.29	83.95	3.41
MiMo-VL-7B-RL	29.84	5.72	15.24	<u>25.18</u>	2.11	83.15	2.61
Qwen-7B-MSEarth	32.36	<u>7.33</u>	17.74	26.59	2.86	<u>83.71</u>	2.73
Intern-S1-mini-MSEarth	<u>33.41</u>	<u>7.64</u>	<b>17.94</b>	<b>27.61</b>	<u>3.25</u>	83.86	3.02
<b><i>Proprietary Models</i></b>							
Gemini2.5-Flash	32.89	7.30	<u>17.79</u>	23.01	<u>3.10</u>	<b>83.96</b>	2.98
Gemini2.5-Flash-Thinking	32.67	7.00	17.42	23.47	2.97	83.85	<u>3.04</u>
Gemini2.5-Pro-Thinking	<b>33.47</b>	<b>7.69</b>	17.15	<u>27.45</u>	<b>3.33</b>	83.62	<b>3.35</b>
Claude-3.5-haiku	26.63	4.47	14.97	17.58	1.38	83.46	2.36
Claude-3.7-Sonnet	29.89	5.52	15.94	21.15	2.15	83.62	2.71
GPT-4o-mini	28.39	5.00	15.96	18.45	1.59	83.90	2.40
GPT-4o	29.89	5.57	16.33	20.19	2.15	83.93	2.72

429  
430 **evaluation.** The captioning results are presented in Table 3, where overlap-based, similarity-based,  
431 and LLM-based metrics exhibit similar trends, with no significant differences observed among the  
432 overlap-based and similarity-based metrics. The Gemini-2.5-Pro model achieves the best performance  
433 across most metrics. The LLM-based metric, designed to evaluate the professionalism and accuracy

of generated captions, demonstrates greater variance compared to similarity-based metrics, making it more suitable for assessing the Scientific Figure Captioning task. Open-source models still show a noticeable gap compared to proprietary models, consistent with the findings from the MCQ results, suggesting a close interconnection between a model’s understanding and reasoning capabilities. Similarly, the results for open-ended question answering, presented in Table 4, show that overlap-based and similarity-based metrics tend to yield higher scores due to the shorter nature of both ground truth answers and model-generated responses. However, for open-ended questions, the focus should be on the rationality and correctness of the answers, making the LLM-based metric a more reasonable evaluation method. This metric also reveals trends consistent with the previous tasks, further highlighting the performance gap between open-source and proprietary models.

Table 4: Performance on scientific open-ended question answering. The best results are highlighted in **bold**, with the second-best underlined.

Model	ROUGE1	ROUGE2	Overlap	METEOR	BLEU	Similarity	LLM
<i>Open-source Models</i>							
LLaVA-onevision-72B	38.07	20.62	37.91	27.88	1.94	89.72	41.56
Qwen2.5-VL-7B-Chat	35.52	20.71	<u>35.24</u>	29.00	2.40	88.62	40.68
Qwen2.5-VL-32B-Chat	36.82	21.33	36.51	29.33	1.83	89.20	41.96
Qwen2.5-VL-72B-Chat	38.57	23.49	38.34	30.65	2.20	89.22	44.82
InternVL2.5-8B	36.12	21.05	35.95	28.25	2.04	89.14	39.05
InternVL2.5-78B	40.59	24.14	40.34	31.15	2.23	<u>90.05</u>	45.64
InternVL3-78B	40.59	23.87	40.42	31.77	2.37	<u>89.98</u>	47.00
Llama3.2-90B-Vision	37.64	22.26	<u>37.53</u>	29.08	1.99	89.16	42.72
DeepSeek-VL2	36.49	20.44	36.37	27.24	1.83	89.38	40.68
Intern-S1-mini	36.43	22.04	37.46	29.72	2.34	89.25	43.69
Intern-S1	43.24	22.04	37.46	29.72	2.34	89.25	43.69
GLM-4.1V-Thinking	<u>3</u>	22.04	37.46	29.72	2.34	89.25	43.69
Qwen-7B-MSEarth	<u>39.55</u>	22.32	<u>38.47</u>	<u>30.10</u>	2.37	<u>89.82</u>	<u>48.19</u>
Intern-S1-mini-MSEarth	<b>41.82</b>	<b>25.11</b>	<b>41.36</b>	31.87	<u>2.44</u>	<b>90.42</b>	49.74
<i>Proprietary Models</i>							
Gemini-2.5-Flash	40.26	22.98	40.02	<u>32.34</u>	2.06	89.39	<b>52.00</b>
Gemini-2.5-Flash-Thinking	39.59	22.30	39.47	30.77	1.87	89.58	46.49
Gemini-2.5-Pro-Thinking	38.67	22.30	38.38	31.73	<b>2.50</b>	88.93	47.70
Claude-3.7-Sonnet	40.61	23.58	40.21	30.75	1.73	89.37	48.33
GPT-4o-mini	36.49	21.37	36.27	28.89	1.63	89.06	41.81
GPT-4o	41.30	<u>24.74</u>	41.03	<b>32.70</b>	2.04	89.78	48.55

### 4.3.1 ANALYSIS AND DISCUSSION

As shown in Table 2, most MLLMs exhibit a pronounced gap between perception-based and reasoning-based performance. On simple visual questions—where answers depend on direct feature extraction—these models routinely exceed 75% accuracy. However, on scientific reasoning tasks that demand domain-specific knowledge, their scores drop sharply (e.g. Gemini-2.5-Pro-Thinking: 77.06% perception vs. 56.31% reasoning). This divergence suggests that while robust perception is a necessary foundation, it is not sufficient for Earth-science inference: once a model surpasses the  $\approx 75\%$  perception threshold, further gains hinge on integrating specialized knowledge and enabling multi-step reasoning. By contrast, models pretrained or fine-tuned on scientific datasets, such as Intern-S1, demonstrate substantially higher accuracy in both perception and reasoning, thereby confirming that general MLLMs lack the requisite Earth-science expertise. Furthermore, we find that further training on the MSEarth training set boosts performance across the board, with the largest relative improvement appearing in reasoning tasks. Taken together, these results underscore the critical role of domain-focused data and architectures in closing the reasoning gap.

Key factors driving low reasoning-task performance include (1) insufficient coverage of Earth science content in general pretraining corpora, (2) the absence of iterative chain-of-thought reasoning modules in standard multimodal fusion backbones, and (3) the scarcity of annotated, multimodal datasets that provide step-by-step scientific rationales.

486 4.3.2 MLLM-BASED METRICS  
487

488 To further establish the correlation between LLMs and human judgment specifically in the domain  
489 of Earth Science VQA, we conducted a human evaluation with four Ph.D. candidates specializing  
490 in Earth sciences. They scored a random sample of 160 questions from our MSEarth Open  
491 Ended benchmark. The models evaluated included Gemini-2.5-Flash, GPT-4o, InternVL3-78B  
492 and QwenVL2.5-72B. Our inter-annotator agreement, measured by Krippendorff’s alpha, is 69.5.  
493 Following LAVE (Mañas et al., 2024), in order to assess the validity of OE-Eval, we calculated its  
494 correlation with human judgment using Spearman’s rank correlation coefficients. We derive a single  
495 “quality” score from the 4 binary ratings (correct/incorrect) per answer as follows: 1.0 if at least  
496 3 annotators rate the answer as correct, 0.5 if only 2 did so, and 0.0 otherwise. The results of this  
497 evaluation are presented in the following table:  
498

Metric	QwenVL2.5-72B	Gemini-2.5-Flash	GPT-4o	InternVL3-78B	Overall
BERTScore	61.16	60.34	63.06	59.79	61.09
ROUGE	67.04	66.90	70.01	63.29	66.81
METEOR	69.00	67.12	67.34	64.55	66.75
BLEU	59.32	58.96	60.57	57.14	59.00
OE-Eval	68.31	<b>67.13</b>	69.85	<b>63.99</b>	<b>67.32</b>

504  
505 Table 5: Spearman correlation across models.  
506

507 From the table’s results, OE-Eval demonstrates a higher consistency with human judgment compared  
508 to all the considered baselines. Traditional metrics are fundamentally ill-suited for evaluating the  
509 scientific nuance and factual correctness required by our benchmark. Our MLLM-based metrics were  
510 explicitly designed to address this limitation, and their superiority is empirically proven via alignment  
511 with human expert judgments.  
512

## 513 4.3.3 HUMAN PERFORMANCE BASELINE

514 To clarify benchmark difficulty and further justify its educational relevance, we also included human  
515 expert scores in MSEarth-Bench-mini. Specifically, we hired three human experts, all of whom are  
516 Ph.D. students with backgrounds in Earth sciences, to evaluate the tasks. We report their average  
517 scores to provide a clear baseline for human performance:  
518

Model	Atmospheric	Solid Earth Geophysics	Geography	Ecology	Geology	Hydrology	Oceanography	Polar	All
InternVL3-78B	50.70%	29.73%	28.57%	47.06%	25.00%	51.02%	25.00%	30.00%	47.33%
Gemini-2.5-Pro	46.48%	58.11%	35.00%	35.71%	47.06%	50.00%	59.18%	50.00%	51.33%
o4-mini	50.00%	45.00%	55.88%	71.43%	49.30%	48.98%	43.33%	62.50%	53.00%
Expert	86.49%	85.00%	85.29%	92.86%	87.32%	85.71%	86.67%	87.50%	87.00%

523 Table 6: Accuracy on MSEarth-Bench-mini across Earth science domains. Human expert scores are  
524 averages of three Ph.D.-level Earth science evaluators.  
525

526 The results clearly demonstrate that human experts consistently outperform the current MLLMs  
527 across all Earth science domains.  
528

529 5 CONCLUSION  
530

531 In this work, we introduce MSEarth, a graduate-level multimodal dataset designed for MLLMs  
532 in geoscientific applications. MSEarth not only serves as a robust test set but also includes rich  
533 training resources aimed at enhancing the geoscientific understanding and reasoning capabilities  
534 of existing MLLMs. Our evaluation reveals significant gaps in current MLLMs’ ability to handle  
535 complex, graduate-level geoscientific reasoning, highlighting opportunities for improvement. We  
536 believe MSEarth will serve as a valuable resource for advancing MLLMs in scientific reasoning and  
537 plan to expand its scope to other scientific domains in future work.  
538

540 ETHICS AND REPRODUCIBILITY STATEMENT  
541

542 All papers used were obtained from OpenDataLab (He et al., 2024b) under the CC BY 4.0 license,  
543 which permits adaptation and redistribution with attribution. We strictly adhered to all licensing  
544 terms and usage requirements specified by OpenDataLab. This work establishes a benchmark for  
545 evaluating the multimodal Earth scientific exploration capabilities of MLLMs in the field of Earth  
546 sciences. It has broader positive impacts, including promoting the responsible use of AI in scientific  
547 research and enhancing public understanding of Earth sciences. We believe MSEarth will serve as a  
548 valuable resource for advancing multimodal language models (MLLMs) in scientific reasoning, and  
549 we plan to expand its scope to other scientific domains in future work. The benchmark is publicly  
550 available to foster further research and innovation in this field. All data in MSEarth are released  
551 anonymously, including the complete dataset on HuggingFace (<https://huggingface.co/MSEarth-Data>) and all data-processing, training, and evaluation code on Anonymous Github  
552 (<https://anonymous.4open.science/r/MSEarth-2B3F>)

554 REFERENCES  
555

556 Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman,  
557 Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. Gpt-4 technical report.  
558 *arXiv preprint arXiv:2303.08774*, 2023.

559

560 AI Anthropic. The claudie 3 model family: Opus, sonnet, haiku. *Claude-3 Model Card*, 1:1, 2024.

561

562 Lei Bai, Zhongrui Cai, Maosong Cao, Weihan Cao, Chiyu Chen, Haojiong Chen, Kai Chen,  
563 Pengcheng Chen, Ying Chen, Yongkang Chen, et al. Intern-s1: A scientific multimodal foundation  
564 model. *arXiv preprint arXiv:2508.15763*, 2025a.

565 Shuai Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Sibo Song, Kai Dang, Peng Wang,  
566 Shijie Wang, Jun Tang, et al. Qwen2. 5-vl technical report. *arXiv preprint arXiv:2502.13923*,  
567 2025b.

568 Satanjeev Banerjee and Alon Lavie. Meteor: An automatic metric for mt evaluation with improved  
569 correlation with human judgments. In *Proceedings of the acl workshop on intrinsic and extrinsic*  
570 *evaluation measures for machine translation and/or summarization*, pp. 65–72, 2005.

571

572 Zhe Chen, Weiyun Wang, Yue Cao, Yangzhou Liu, Zhangwei Gao, Erfei Cui, Jinguo Zhu, Shenglong  
573 Ye, Hao Tian, Zhaoyang Liu, et al. Expanding performance boundaries of open-source multimodal  
574 models with model, data, and test-time scaling. *arXiv preprint arXiv:2412.05271*, 2024.

575

576 Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad  
577 Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, et al. The llama 3 herd of  
578 models. *arXiv preprint arXiv:2407.21783*, 2024.

579

580 Yunzhuo Hao, Jiawei Gu, Huichen Will Wang, Linjie Li, Zhengyuan Yang, Lijuan Wang, and  
581 Yu Cheng. Can mllms reason in multimodality? emma: An enhanced multimodal reasoning  
582 benchmark. *arXiv preprint arXiv:2501.05444*, 2025.

583

584 Chaoqun He, Renjie Luo, Yuzhuo Bai, Shengding Hu, Zhen Leng Thai, Junhao Shen, Jinyi Hu,  
585 Xu Han, Yujie Huang, Yuxiang Zhang, et al. Olympiadbench: A challenging benchmark for  
586 promoting agi with olympiad-level bilingual multimodal scientific problems. *arXiv preprint*  
587 *arXiv:2402.14008*, 2024a.

588

589 Conghui He, Wei Li, Zhenjiang Jin, Chao Xu, Bin Wang, and Dahu Lin. Opendatalab: Empowering  
590 general artificial intelligence with open datasets. *arXiv preprint arXiv:2407.13773*, 2024b.

591

592 Aaron Hurst, Adam Lerer, Adam P Goucher, Adam Perelman, Aditya Ramesh, Aidan Clark, AJ Os-  
593 trow, Akila Welihinda, Alan Hayes, Alec Radford, et al. Gpt-4o system card. *arXiv preprint*  
594 *arXiv:2410.21276*, 2024.

595

596 Lars Benedikt Kaesberg, Jonas Becker, Jan Philip Wahle, Terry Ruas, and Bela Gipp. Voting or  
597 consensus? decision-making in multi-agent debate. *arXiv preprint arXiv:2502.19130*, 2025.

594 Kartik Kuckreja, Muhammad Sohail Danish, Muzammal Naseer, Abhijit Das, Salman Khan, and  
 595 Fahad Shahbaz Khan. Geochat: Grounded large vision-language model for remote sensing.  
 596 In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp.  
 597 27831–27840, 2024.

598 Woosuk Kwon, Zhuohan Li, Siyuan Zhuang, Ying Sheng, Lianmin Zheng, Cody Hao Yu, Joseph  
 599 Gonzalez, Hao Zhang, and Ion Stoica. Efficient memory management for large language model  
 600 serving with pagedattention. In *Proceedings of the 29th Symposium on Operating Systems  
 601 Principles*, pp. 611–626, 2023.

602 Xian Yeow Lee, Shunichi Akatsuka, Lasitha Vidyaratne, Aman Kumar, Ahmed Farahat, and Chetan  
 603 Gupta. Reliable decision-making for multi-agent llm systems. 2025.

604 Bo Li, Yuanhan Zhang, Dong Guo, Renrui Zhang, Feng Li, Hao Zhang, Kaichen Zhang, Peiyuan  
 605 Zhang, Yanwei Li, Ziwei Liu, et al. Llava-onevision: Easy visual task transfer. *arXiv preprint  
 606 arXiv:2408.03326*, 2024a.

607 Junxian Li, Di Zhang, Xunzhi Wang, Zeying Hao, Jingdi Lei, Qian Tan, Cai Zhou, Wei Liu, Yaotian  
 608 Yang, Xinrui Xiong, et al. Chemvlm: Exploring the power of multimodal large language models in  
 609 chemistry area. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, pp.  
 610 415–423, 2025.

611 Lei Li, Yuqi Wang, Runxin Xu, Peiyi Wang, Xiachong Feng, Lingpeng Kong, and Qi Liu. Multimodal  
 612 arxiv: A dataset for improving scientific comprehension of large vision-language models. *arXiv  
 613 preprint arXiv:2403.00231*, 2024b.

614 Zekun Li, Xianjun Yang, Kyuri Choi, Wanrong Zhu, Ryan Hsieh, HyeonJung Kim, Jin Hyuk Lim,  
 615 Sungyoung Ji, Byungju Lee, Xifeng Yan, et al. Mmsci: A multimodal multi-discipline dataset for  
 616 phd-level scientific comprehension. In *AI for Accelerated Materials Design-Vienna 2024*, 2024c.

617 Zhenwen Liang, Kehan Guo, Gang Liu, Taicheng Guo, Yujun Zhou, Tianyu Yang, Jiajun Jiao, Renjie  
 618 Pi, Jipeng Zhang, and Xiangliang Zhang. Scemqa: A scientific college entrance level multimodal  
 619 question answering benchmark. *arXiv preprint arXiv:2402.05138*, 2024a.

620 Zijing Liang, Yanjie Xu, Yifan Hong, Penghui Shang, Qi Wang, Qiang Fu, and Ke Liu. A survey  
 621 of multimodel large language models. In *Proceedings of the 3rd International Conference on  
 622 Computer, Artificial Intelligence and Control Engineering*, pp. 405–409, 2024b.

623 Chin-Yew Lin. Rouge: A package for automatic evaluation of summaries. In *Text summarization  
 624 branches out*, pp. 74–81, 2004.

625 Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. Visual instruction tuning. *Advances in  
 626 neural information processing systems*, 36:34892–34916, 2023a.

627 Yang Liu, Dan Iter, Yichong Xu, Shuohang Wang, Ruochen Xu, and Chenguang Zhu. G-eval: Nlg  
 628 evaluation using gpt-4 with better human alignment. *arXiv preprint arXiv:2303.16634*, 2023b.

629 Pan Lu, Swaroop Mishra, Tanglin Xia, Liang Qiu, Kai-Wei Chang, Song-Chun Zhu, Oyvind Tafjord,  
 630 Peter Clark, and Ashwin Kalyan. Learn to explain: Multimodal reasoning via thought chains for  
 631 science question answering. *Advances in Neural Information Processing Systems*, 35:2507–2521,  
 632 2022.

633 Chengqian Ma, Zhanxiang Hua, Alexandra Anderson-Frey, Vikram Iyer, Xin Liu, and Lianhui  
 634 Qin. Weatherqa: Can multimodal language models reason about severe weather? *arXiv preprint  
 635 arXiv:2406.11217*, 2024.

636 Oscar Mañas, Benno Krojer, and Aishwarya Agrawal. Improving automatic vqa evaluation using large  
 637 language models. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pp.  
 638 4171–4179, 2024.

639 Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. Bleu: a method for automatic  
 640 evaluation of machine translation. In *Proceedings of the 40th annual meeting of the Association  
 641 for Computational Linguistics*, pp. 311–318, 2002.

648 Jonathan Roberts, Kai Han, Neil Houlsby, and Samuel Albanie. Scifibench: Benchmarking large  
 649 multimodal models for scientific figure interpretation. *arXiv preprint arXiv:2405.08807*, 2024.  
 650

651 Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Xiao Bi, Haowei Zhang,  
 652 Mingchuan Zhang, YK Li, Yang Wu, et al. Deepseekmath: Pushing the limits of mathematical  
 653 reasoning in open language models. *arXiv preprint arXiv:2402.03300*, 2024.

654 Noah Siegel, Zachary Horvitz, Roie Levin, Santosh Divvala, and Ali Farhadi. Figureseer: Parsing  
 655 result-figures in research papers. In *Computer Vision–ECCV 2016: 14th European Conference,  
 656 Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part VII 14*, pp. 664–680.  
 657 Springer, 2016.

658 Gemini Team, Rohan Anil, Sebastian Borgeaud, Jean-Baptiste Alayrac, Jiahui Yu, Radu Soricut,  
 659 Johan Schalkwyk, Andrew M Dai, Anja Hauth, Katie Millican, et al. Gemini: a family of highly  
 660 capable multimodal models. *arXiv preprint arXiv:2312.11805*, 2023.  
 661

662 Bin Wang, Chao Xu, Xiaomeng Zhao, Linke Ouyang, Fan Wu, Zhiyuan Zhao, Rui Xu, Kaiwen Liu,  
 663 Yuan Qu, Fukai Shang, et al. Mineru: An open-source solution for precise document content  
 664 extraction. *arXiv preprint arXiv:2409.18839*, 2024.

665 Wenhui Wang, Furu Wei, Li Dong, Hangbo Bao, Nan Yang, and Ming Zhou. Minilm: Deep self-  
 666 attention distillation for task-agnostic compression of pre-trained transformers. *Advances in neural  
 667 information processing systems*, 33:5776–5788, 2020.

668

669 An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan Li,  
 670 Dayiheng Liu, Fei Huang, Haoran Wei, et al. Qwen2. 5 technical report. *arXiv preprint  
 671 arXiv:2412.15115*, 2024a.

672 Joshua C Yang, Damian Dalisan, Marcin Korecki, Carina I Hausladen, and Dirk Helbing. Llm voting:  
 673 Human choices and ai collective decision-making. In *Proceedings of the AAAI/ACM Conference  
 674 on AI, Ethics, and Society*, volume 7, pp. 1696–1708, 2024b.

675

676 Xiang Yue, Yuansheng Ni, Kai Zhang, Tianyu Zheng, Ruoqi Liu, Ge Zhang, Samuel Stevens, Dongfu  
 677 Jiang, Weiming Ren, Yuxuan Sun, et al. Mmmu: A massive multi-discipline multimodal under-  
 678 standing and reasoning benchmark for expert agi. In *Proceedings of the IEEE/CVF Conference on  
 679 Computer Vision and Pattern Recognition*, pp. 9556–9567, 2024.

680 Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q Weinberger, and Yoav Artzi. Bertscore: Evaluating  
 681 text generation with bert. *arXiv preprint arXiv:1904.09675*, 2019.

682 Jinguo Zhu, Weiyun Wang, Zhe Chen, Zhaoyang Liu, Shenglong Ye, Lixin Gu, Yuchen Duan, Hao  
 683 Tian, Weijie Su, Jie Shao, et al. Internvl3: Exploring advanced training and test-time recipes for  
 684 open-source multimodal models. *arXiv preprint arXiv:2504.10479*, 2025.

685

686

687

688

689

690

691

692

693

694

695

696

697

698

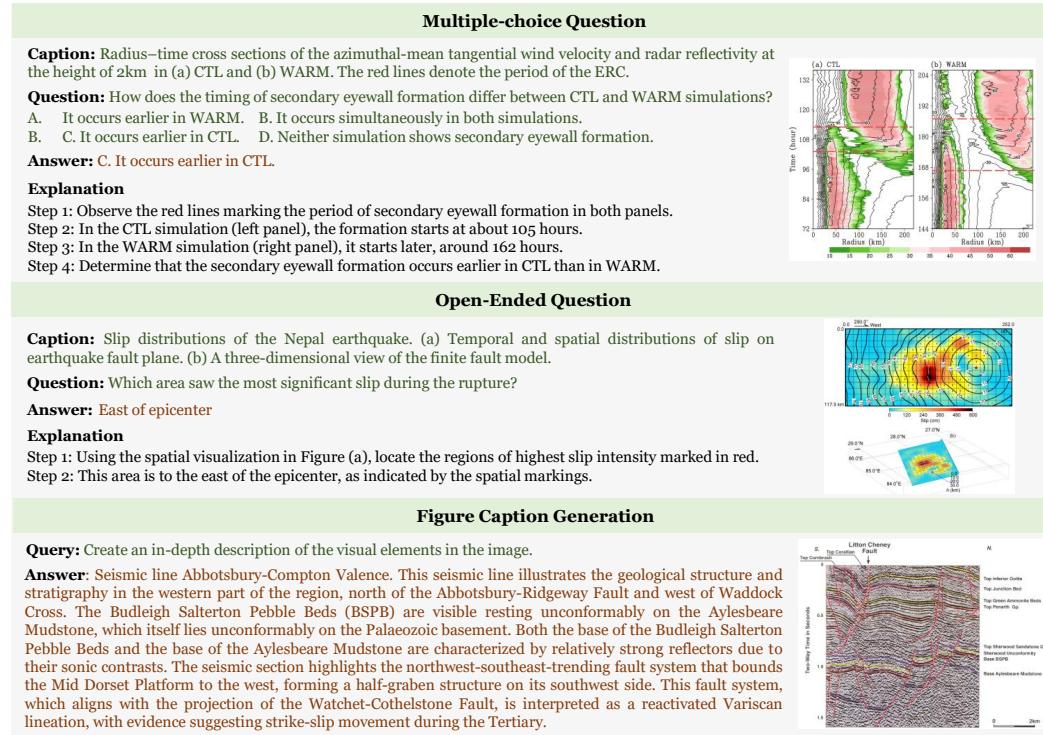
699

700

701

702 Table 7: Comparison with previous multimodal scientific benchmarks. Task types include OE (Open-  
 703 ended QA), MCQ (Multiple-choice QA), and CG (Caption Generation).

Benchmark Dataset	Science Topics	Tasks	Difficulty	Original Questions	Validated
<b>Multimodal Scientific Benchmarks</b>					
ScienceQA (Lu et al., 2022)	General Science	MCQ	Primary	✗	✓
SceMQA (Liang et al., 2024a)	General Science	MCQ,OE	Pre-College	✗	✓
Mmmu (Yue et al., 2024)	General Science	MCQ,OE	College-level	✗	✓
OlympiadBench (He et al., 2024a)	Math, Physics	OE	Competition	✗	✓
EMMA (Hao et al., 2025)	General Science	MCQ	Mixed	✓	✓
<b>Paper-Based Multimodal Scientific Benchmarks</b>					
SciFIBench (Roberts et al., 2024)	General Science	MCQ	Graduate-Level	✗	✓
ArxivCap/QA (Li et al., 2024b)	General Science	CG,MCQ	Graduate-Level	✓	✗
MMSci (Li et al., 2024c)	General Science	CG,MCQ	Graduate-Level	✗	✗
MSEarth(Ours)	Earth Science	CG,MCQ,OE	Graduate-Level	✓	✓



745 Figure 6: Examples of the three types of scientific question-answering tasks presented in our  
 746 benchmark.

## A USAGE OF LANGUAGE MODELS

747 We utilized a large language model (LLM) to aid in the preparation of this manuscript. Its use was  
 748 limited to editorial tasks, including proofreading for typographical errors, correcting grammar, and  
 749 improving the clarity and readability of the text.

## B BENCHMARK DETAILS

### B.1 FIELD EXPLANATION OF MSEARTH

750 In Table 8, we provide an explanation of each field for the three tasks in MSEarth.

756	Field Name	Input	Description	
757	<b>Multiple-choice Question</b>			
758	question_id	✗	The unique identifier for the question.	
759	query	✓	Contains the original caption, question, and options.	
760	response	✗	The correct answer to the question.	
761	images	✓	The file path to the associated image(s).	
762	refined_caption	✗	The enhanced image description based on the paper content.	
763	classification	✗	The classification of the question, including its domain and discipline.	
764	reasoning_chain	✗	The reasoning steps to arrive at the answer.	
765	<b>Open-Ended Question</b>			
766	question_id	✗	The unique identifier for the question.	
767	query	✓	Contains the original caption and the question.	
768	response	✗	The correct answer to the question.	
769	images	✓	The file path to the associated image(s).	
770	refined_caption	✗	The enhanced image description based on the paper content.	
771	classification	✗	The classification of the question, including its domain and discipline.	
772	<b>Caption Generation</b>			
773	question_id	✗	The unique identifier for the question.	
774	query	✓	The question.	
775	response	✗	The correct answer (refined caption).	
776	images	✓	The file path to the associated image(s).	
777	context	✗	The text from the paper that describes the image.	
778	original_caption	✗	The original caption of the image.	
779	classification	✗	The classification of the question, including its domain and discipline.	

779 **Table 8: Field Descriptions for Different Tasks.** The table provides details about each field, whether  
780 it is used as input, and its description. Fields are grouped by task type: MCQ, OE, and Caption  
781 Generation.

## 784 B.2 FORMAT CONVERSION

785 Specifically, our data source is a collection of papers gathered by OpenDataLab (He et al., 2024b)  
786 from online resources. These papers were processed using **MinerU**, which converted the textual  
787 content of the PDFs into JSON format and saved the images as PNG files. In Figure 7, we present  
788 a portion of the content list from a processed PDF paper, highlighting the original caption (raw  
789 caption) of the image and the corresponding discussion section. It is evident that the discussion of  
790 figures in the paper contains substantial scientific reasoning, which is crucial for a comprehensive  
791 understanding of scientific figures.

## 793 B.3 PAPER FILTERING

795 To classify scientific papers into relevant Earth system categories, we employ a similarity-based  
796 approach using pre-trained sentence embeddings and cosine similarity. The process begins by  
797 generating embeddings for both the paper’s title and predefined keywords using the pre-trained all-  
798 MiniLM-L6-v2 (Wang et al., 2020) model, which captures the semantic meaning of textual data. First,  
799 we calculate the similarity between the paper’s title and a set of general positive keywords, such as  
800 “Earth,” “Earth system,” “hydrosphere,” “biosphere,” “lithosphere,” “atmosphere,” and “cryosphere.”  
801 The cosine similarity is computed between the embedding of the paper’s title and the embedding of  
802 the general positive keywords. If the similarity score is below a predefined threshold (0.2), the paper  
803 is excluded from further analysis, as it is deemed irrelevant to the sciences of the Earth system.

804 To further filter out irrelevant papers, we calculate the similarity between the paper’s title and a  
805 set of general negative keywords, such as “cell biology,” “virus,” “pharmaceuticals,” “chemistry,”  
806 “physics,” and “astronomy.” If the similarity score exceeds a predefined threshold (0.1), the paper is  
807 excluded, as it is likely to belong to unrelated disciplines. For papers that pass the initial filtering,  
808 we calculate their similarity with predefined positive classification keywords for each Earth system  
809 category (e.g., hydrosphere, biosphere, lithosphere, atmosphere, cryosphere). Each category contains  
a list of domain-specific keywords. For example, the hydrosphere category includes keywords such

810  
 811  
 812  
 813  
 814  
 815  
 816  
 817  
 818  
 819  
 820  
 821  
 822  
 823  
 824  
 825  
 826  
 827  
 828  
 829  
 830  
 831  
 832  
 833  
 834  
 835  
 836  
 837  
 838  
 839  
 840  
 841  
 842  
 843  
 844  
 845  
 846  
 847  
 848  
 849  
 850  
 851  
 852  
 853  
 854  
 855  
 856  
 857  
 858  
 859  
 860  
 861  
 862  
 863

**Structure of Content List Field**

```

"content_list": [
  {
    "type": "text",
    "text": "LUMINESCENCE STUDIES ON NEOTECTONIC EVENTS IN SOUTH-CENTRAL KUMAUN HIMALAYA -- A FEASIBILITY STUDY",
    "text_level": 1
  },
  .....
  {
    "type": "text",
    "text": "The continued northward movement of the Indian plate has caused accumulation of stresses which get periodically released, resulting in earthquakes and neotectonic activity along faults. Geophysical and structural studies suggest that seismicity in the Himalaya is related to movements along three major longitudinal thrusts/faults (Fig. 1) viz. the Himalayan Frontal Fault (HFF), the Main Boundary Thrust (MBT) and the Main Central Thrust (MCT) (Valdiya, 1986, 1988; Nakata, 1989). These thrusts divide the Himalaya into three distinct lithotectonic zones. These zones are further dissected by numerous transverse faults (Valdiya, 1976; Khattri and Tyagi, 1983a). A concept of locked segments has been propounded to assess the seismogenic potential of various sectors and is related to accumulation of stresses and their eventual release along the faulted zones (Khattri and Tyagi 1983b). The repeat frequency of this process of locking of stresses and release of energy is not yet well established due to lack of dating methods. Needless to say, this is an aspect cardinal to the estimation of seismogenic hazards involved in planning large scale engineering and societal projects in the Himalaya."
    "text_level": 0
  },
  {
    "type": "image",
    "img_path": "s3://llm-pipeline-media/pdf-imgs/f5def6d72e7dbfe47ada87dobe9084997c67f989011db77a255d1817a33fe86b.png",
    "img_caption": "FIG. 1. Geological map indicating fault zones and locked segments in Himalaya"
  },
  .....
]

```

Figure 7: Examples of the content list field in a paper.

as “water cycle,” “ocean,” “rivers,” “lakes,” and “groundwater,” while the biosphere category includes “ecosystem,” “biodiversity,” “habitat,” and “species.” The cosine similarity is computed between the paper’s title and each keyword within a category, and the average similarity score for each category is calculated.

The category with the highest average similarity score is selected as the most relevant classification for the paper, provided the score exceeds a predefined threshold (0.15). To ensure robustness, we also calculate the similarity between the paper’s title and negative classification keywords for each category. For example, the hydrosphere negative keywords include “chemistry,” “universe,” “planets,” and “astronomy,” while the biosphere negative keywords include “cell biology,” “medicine,” and “pharmacology.” The final relevance score for each category is computed as the difference between the positive and negative similarity scores, ensuring that papers with high relevance to unrelated fields are excluded. The final classification of a paper is determined by passing the general positive and negative keyword thresholds, identifying the category with the highest positive relevance score adjusted by subtracting the negative relevance score, and ensuring the adjusted relevance score exceeds the classification threshold (0.15). This approach allows us to systematically classify papers into Earth system categories while filtering out irrelevant content, leveraging semantic embeddings and cosine similarity to ensure that the classification is both accurate and interpretable.

#### B.4 IMAGE FILTERING

Next, we further filtered the images in these 103,108 papers. Our goal was to retain Earth observation images, as our task focuses on evaluating the model’s ability to understand and reason about scientific phenomena in Earth sciences. These images include various types of visual data, such as those representing geophysical processes, atmospheric phenomena, geographic features, weather patterns, and cartographic representations.

To achieve this, we employed a systematic filtering pipeline based on the Qwen2-VL-7B-Instruct model. The filtering process was guided by a carefully designed prompt, which instructed the model to classify each image as either an Earth observation image or not. Specifically, the prompt defined Earth observation images as those depicting remote sensing imagery, atmospheric data visualizations, aerial views of geographical features (e.g., rivers, urban landscapes), weather-related images (e.g., precipitation maps, typhoon tracks), and cartographic representations. Conversely, the prompt explicitly excluded images containing biological entities (e.g., humans, plants, animals), artificial objects (e.g., vehicles, device structures), data visualizations (e.g., statistical charts, line graphs, scatter plots), text-based content, or blank images.

Category	Type	Keywords
General	Positive	Earth, Earth system, hydrosphere, biosphere, lithosphere, atmosphere, cryosphere
General	Negative	cell biology, virus, pharmaceuticals, chemistry, physics, astronomy, food science, proteins, microbiology
Hydrosphere	Positive	water cycle, ocean, rivers, lakes, groundwater, ice caps, aquifers, precipitation, evaporation, humidity
Hydrosphere	Negative	chemistry, universe, planets, astronomy, astrophysics, space, stars, galaxy, cosmology
Biosphere	Positive	ecosystem, biodiversity, habitat, species, biomes, ecological balance, carbon cycle
Biosphere	Negative	cell biology, chemistry, medicine, pharmacology, microbiology, biochemistry, toxicology, pathology, clinical
Lithosphere	Positive	earthquake, tectonic plates, earth's crust, minerals, rocks, soil, sediments, mountains, volcanoes, landforms, geological processes
Lithosphere	Negative	ancient texts, archaeology, culture, history, artifacts, civilization, prehistoric, mythology, anthropology
Atmosphere	Positive	stratosphere, troposphere, weather, climate, greenhouse gases, ozone layer, air pressure, humidity, winds, carbon dioxide, temperature
Atmosphere	Negative	universe, galaxy, astronomy, astrophysics, space, stars, planets, cosmology, black holes, nebula, solar system
Cryosphere	Positive	glaciers, ice sheets, sea ice, permafrost, snowpack, icebergs, frozen ground, climate change, albedo effect, polar regions
Cryosphere	Negative	frozen food, ice cream, refrigeration, freezing, cold storage, ice cubes, food preservation, chilling, frost

Table 9: Keywords for positive and negative classifications across different Earth system categories. The table includes general keywords as well as specific keywords for hydrosphere, biosphere, lithosphere, atmosphere, and cryosphere.

The filtering process was implemented as follows: for each image, the model was provided with both the image and the prompt, and it generated a binary output (“1” for Earth observation images and “0” otherwise). To ensure robustness, the model’s output was validated through multiple sampling attempts with slight variations in generation parameters (e.g., temperature). If the model consistently classified an image as “1,” it was retained; otherwise, it was discarded. This iterative and robust classification approach allowed us to minimize false positives and negatives in the filtering process. After this step, we retained around 83K papers, which contained images classified as Earth observation images. These filtered images form the basis for subsequent analysis and evaluation of the model’s capabilities in understanding and reasoning about Earth science phenomena.

The following prompt was used to retain Earth observation images:

Analyze the provided image and classify it as an Earth observation image or not.

Earth observation images include, but are not limited to:

- Remote sensing imagery,
- Atmospheric data visualizations,
- Aerial views of geographical features (e.g., rivers, urban landscapes),
- Weather-related images (e.g., precipitation maps, typhoon tracks),
- Cartographic representations.

Exclude images depicting:

- Biological entities (humans, plants, animals),
- Artificial objects (vehicles, device structures),
- Data visualizations (statistical charts, line graphs, scatter plots),
- Text-based content or blank images.

Output format:

- Return “1” if the image is an Earth observation image.

918

919

920

921

922

923

## B.5 CONTENT FILTERING

924

925

926

927

928

929

To construct our benchmark, which requires generating VQA tasks supported by the content of the papers, we ensured that the selected figures not only had captions but were also discussed in detail within the text of the papers. Using regular expressions, we extracted the figure numbers and identified corresponding discussions in the main body of the papers. Figures with discussions exceeding two sentences were included in the final dataset. Finally, we selected 64,560 papers, resulting in a total of 289,891 figures for further processing.

930

931

## B.6 PROMPT DESIGNER FOR MSEARTH

932

The following prompt was used to generate a refined caption:

934

935

936

937

938

You are an expert assistant in scientific image analysis and caption generation. Your task is to rewrite or generate a new, detailed caption for the provided figure using the original caption and only the sentences or information from the Relevant Content that are directly associated with this figure.

**Please strictly follow these guidelines:**

939

940

941

942

943

944

- Assume the figure does not reference or depend on other figures in the document.
- Exclude any mention of other figures, their content, or references in the caption.
- If subfigures are present, provide specific descriptions for each subfigure accordingly. Otherwise, assume it represents a single figure.
- The new caption must be detailed, precise, and include only the relevant details from the provided content.

**Inputs for caption generation:**

945

946

947

- Original Caption: {caption}
- Relevant Content: {content}

Now write a detailed, high-quality caption for this figure below:

948

949

The following prompt was used to generate diverse VQAs:

950

951

952

953

954

You are an advanced AI model specialized in generating high-quality Visual Question Answering (VQA) tasks. Your role is to generate a diverse set of VQA questions, answers, and reasoning chains based on the provided visual input (a figure) and its captions.

**DEFINITIONS:**

955

956

957

958

959

960

961

962

963

964

965

966

967

968

969

970

971

1. **Figure:** A scientific or illustrative figure provided as the primary visual input. Test-takers will analyze this image to answer the questions.
2. **Caption:** A concise summary describing key aspects of the Figure.
3. **Supplementary:** In-depth information (e.g., summarized expert insight, detailed analysis, or background knowledge) that you can use to assist in designing advanced and meaningful questions. However, test-takers cannot access this information.

**INPUT INFORMATION PROVIDED:**

963

964

965

966

967

968

969

970

971

- **Caption:** {raw caption}
- **Supplementary:** {refined caption}

**TASK INSTRUCTIONS:**

Your task is to create a variety of advanced VQA tasks designed to test visual and contextual understanding based on the Figure and Caption. Below are key rules and guidelines you must follow:

972  
973  
974  
975  
976  
977  
978  
979

## 1. USE OF INPUT SOURCES:

- Ensure that no question can be answered entirely using Caption without observations from the Figure. Figure content should always serve as a primary source for reasoning.
- **Supplementary Usage:** The correct answers are encouraged to be derived from the Supplementary information. Focus on crafting questions where the Supplementary plays a crucial role in providing the answer.

## 2. QUESTION TYPES:

- **Multiple Choice Questions (MCQs):** At least 2 questions must be of this type, with 4 distinct options (A-D) and one correct answer.
  - Ensure that only one option can be logically correct based on the provided information (Figure, Caption, and/or Supplementary). Avoid creating options that could lead to ambiguous interpretations or alternate correct answers.
  - Incorrect options must be plausible and relevant to the context but should contain subtle logical flaws or lack supporting evidence when compared to the correct option.
- **Open-Ended Questions:** At least 2 questions must be open-ended, requiring concise and precise answers (no more than 4 words).

## 3. REASONING CHAINS:

- For every question, you must include a reasoning chain. The chain explains the logical process by which the correct answer can be determined.
- The reasoning chain must:
  - Be clear, step-by-step, and never explicitly mention Caption or Supplementary in reasoning\_chain (e.g., "According to the Supplementary" or "The Caption states").
  - Use different levels of reasoning complexity.

## 4. OUTPUT STRUCTURE:

The output must be written in **JSON format** using the structure below:

```
[  
  {  
    "question_type": MCQ or OE  
    "question": "Your question here",  
    "options": [A,B,C,D],  
    "answer": "Correct option or short answer",  
    "reasoning_chain": ["Step 1: ...", ...]  
  },  
  // Additional questions in the same format...  
]
```

## 5. TASK GUIDELINES:

1. Questions that are grounded in the Supplementary context are highly encouraged. These questions should require the test-taker to refer to in-depth knowledge and insights not immediately visible in the Figure or Caption.
2. Avoid referencing the Supplementary in any question and reasoning\_chain (e.g., "According to the Supplementary" or "The Supplementary states").

Provide your response below:

1019

1020

## C MULTI-AGENT VOTING

1022

## C.1 PROMPT

1024

1025 The following prompt was used to generate a normal answer for MCQ:

1026 You are tasked with answering a multiple-choice question about the given input image.  
 1027  
 1028  
 1029 **INSTRUCTIONS:**  
 1030   1. Carefully analyze the input image and the provided query.  
 1031   2. Based on the image, select the correct option (e.g., 'A', 'B', 'C') or directly state the  
 1032    correct option content.  
 1033   3. Provide reasoning explaining how to derive the correct answer.  
 1034  
 1035 **INPUT:**  
 1036   • **Query:** {query}  
 1037  
 1038 **OUTPUT FORMAT:**  
 1039  
 1040   The output must be written in **JSON format** using the structure below:  
 1041   {  
 1042       "answer": "Correct option or short answer",  
 1043       "Explanation": "Explaining how to derive correct answer."  
 1044   }

1045 The following prompt was used to generate a enhanced answer for MCQ:  
 1046

1047 You are tasked with answering a multiple-choice question about the given input image.  
 1048  
 1049

1050 **INPUT:**  
 1051   • **Question:** {question}  
 1052   • **Refined Caption:** {caption}  
 1053

1054 **INSTRUCTIONS:**

1055   1. Carefully analyze the input image and its caption.  
 1056   2. Based on the image and caption, select the correct option (e.g., 'A', 'B', 'C') or  
 1057    directly state the correct option content.  
 1058

1059 **OUTPUT FORMAT:**  
 1060

1061   The output must be written in **JSON format** using the structure below:  
 1062

1063   {  
 1064       "answer": "Correct option or short answer",  
 1065       "Explanation": "Explaining how to derive correct answer."  
 1066 }

## 1067 C.2 EXAMPLE OF DIFFERENT LEVELS OF QUESTIONS

1068  
 1069 Figure 9 illustrates an example of a simple problem in multi-agent voting, while Figure 10 presents  
 1070 an example of a domain-specific problem, and Figure 11 demonstrates an example of a challenging  
 1071 problem. The most notable distinction lies in the varying levels of perceptual ability required by  
 1072 the model: simple and challenging problems primarily differ in the model's ability to perceive and  
 1073 interpret images, whereas domain-specific problems emphasize the model's knowledge in specialized  
 1074 fields. Additionally, the answers to domain-specific questions are often supported by evidence found  
 1075 in the "refined caption" field provided in the paper.

1076 To construct the benchmark datasets, we sampled data from the multi-agent automated filtering  
 1077 process as follows: 900 questions from Phase A, 1800 questions from Phase B, and 300 questions  
 1078 from Phase C were selected to form the multiple-choice question (MCQ) set, while 500 questions  
 1079 from Phase A and 1000 questions from Phase B were selected to form the open-ended question set.  
 All sampled data were subsequently validated by experts to ensure accuracy and quality.

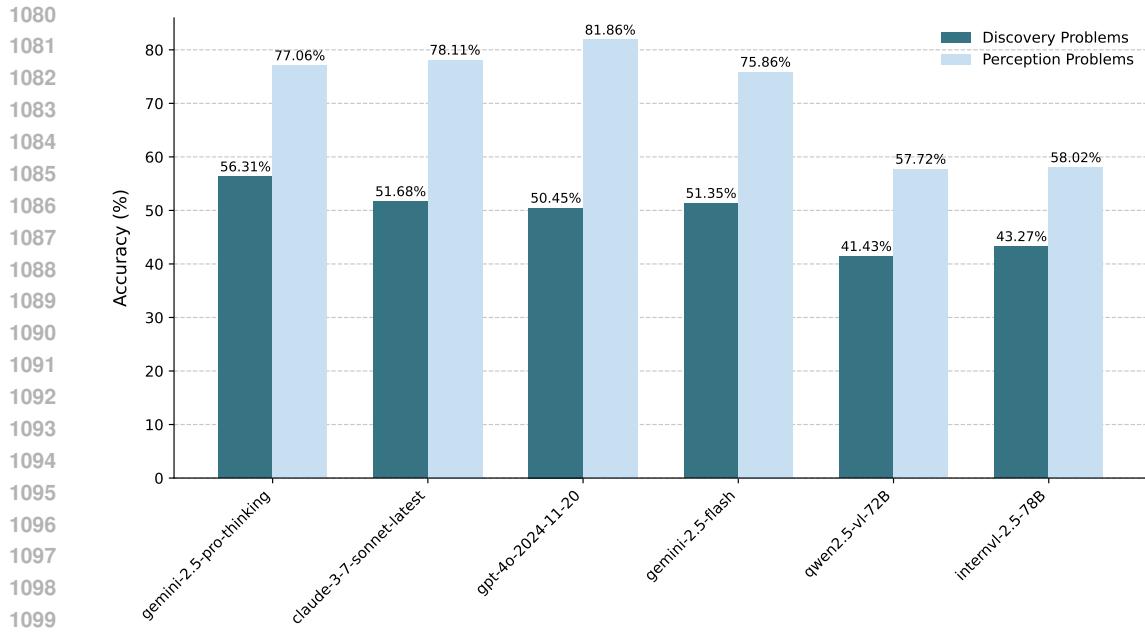


Figure 8: Models' accuracy on reasoning and perception problems.

## D EXPERT VALIDATION

### D.1 DETAILS

We recruited annotators with a background in Earth sciences and a master's degree through an annotation company to label the data. The annotated dataset consists of 3,000 MCQs and 1,500 open-ended QAs. We provided the annotators with figures, queries, reasoning chains, and our processed refined captions to assist them in evaluating whether the provided answers were reasonable. For questions where the answers could not be found in the refined captions, the annotators were required to use their own knowledge to determine the correctness of the answers. If they were unable to make a judgment, such questions were discarded to ensure that the filtered dataset contained only accurate and complete questions. The tasks assigned to the annotators are described below:

The evaluation framework categorizes questions based on several criteria. First, the **Image Type of Reasoning Required** distinguishes between questions involving a single image, where the input consists of just one image, and those with a **Single-image focus**, where multiple images are present but the question pertains to one specific image. Additionally, **Multi-image reasoning** questions require comparing or reasoning across multiple images.

Next, the **Type of Scientific Question** is considered. **Perception Questions** are those where answers can be derived through basic observation, such as identifying position or color. These questions do not have answers in the refined captions and require manual evaluation of their validity. In contrast, **Reasoning Questions** necessitate domain-specific knowledge for answering, and annotators must verify if the answer can be derived from the refined caption field.

The **Completeness of Questions** is another criterion, where questions are classified as **Complete** if all necessary information is provided in the question or image, and **Incomplete** if missing information makes it difficult or impossible to answer.

Finally, the **Correctness of Questions** assesses whether the provided answer is accurate, categorizing them as **Correct** or **Incorrect** based on the accuracy of the answer.

1134

1135

1136

1137

1138

1139

1140

1141

1142

1143

1144

1145

1146

1147

1148

1149

1150

1151

1152

1153

1154

1155

1156

1157

1158

1159

1160

1161

1162

1163

1164

1165

1166

1167

1168

1169

1170

1171

1172

1173

1174

1175

1176

1177

1178

1179

1180

1181

1182

1183

1184

1185

1186

1187

1188

**Easy Multiple-choice Question**

**Image Caption:**  
Coastline totally covered by seawall and concrete armor units on Shimizu coast (July 1995).

**Question:**  
How has the natural aesthetic of the coastal environment likely changed due to the structures visible in the figure?

**Options:**

- A. Enhanced beauty due to added greenery
- B. Preservation of the sandy beach's appearance
- C. Loss of natural shoreline aesthetics
- D. Creation of diverse wildlife habitats

**Response:**  
C. Loss of natural shoreline aesthetics

**Reasoning Chain:**  
Step 1: Use visual cues from the figure to observe the dominance of artificial structures, such as seawalls and concrete units.  
Step 2: Reflect on the visual impact of manmade elements completely covering the coastline.  
Step 3: Identify that no sandy beaches or natural aesthetics remain visible.  
Step 4: Determine that "Loss of natural shoreline aesthetics" best describes the impact based on both image and caption.

**Image Caption:**  
Hydrogeological cross section in the Grombalia basin

**Question:**  
Which layer lies directly below the shallow aquifer in most parts of the Grombalia basin?

**Options:**

- A. Topsoil
- B. Clayey sands aquiclude
- C. Semi-deep aquifer
- D. Sandstone

**Response:**  
B. Clayey sands aquiclude

**Reasoning Chain:**  
Step 1: The visual information from the figure shows a cross-sectional view of the stratigraphy.  
Step 2: Observing the labeling in the legend and the diagram, clayey sands aquiclude are consistently shown directly below the shallow aquifer.  
Step 3: Correlating these observations, the layer below the shallow aquifer is identified as the clayey sands aquiclude.

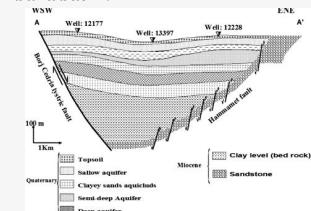


Figure 9: An example of easy multiple-choice VQA.

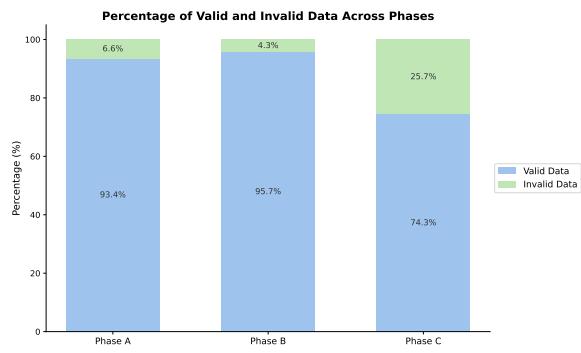


Figure 12: Proportion of valid and invalid data after manual screening across different phases. Phase B, or Specialized VQA, exhibits the highest quality.

1188

1189

1190 **Moderate Multiple-choice Question**

1191

1192 **Image Caption:**

1193 Change in sea level pressure between different pairs of sensitivity experiments with the sea surface temperature gradient altered in the Atlantic Ocean, Pacific Ocean, or in combinations in both oceans.

1194

1195 ....

1196 Bottom Panel: Difference between experiments AdPi and AiPd. Most changes outside of the range shown (\$-2.2\$ to \$2.2\$-\$\text{mb}\$) are significant at the 95% level.

1197

1198 **Question:**

1199 What mechanism best explains the Southern Hemisphere pressure changes observed in the bottom panel (AdPi-AiPd)?

1200

1201 **Options:**

1202 A. Pacific tropical warming amplifies wave refraction.

1203 B. High-latitude cooling destabilizes atmospheric conditions.

1204 C. Sea ice variations influence atmospheric circulation.

1205 D. Increased sensible heat transports planetary momentum.

1206

1207 **Response:**

1208 C. Sea ice variations influence atmospheric circulation.

1209

1210 **Reasoning Chain:**

1211 Step 1: Look at the bottom panel (AdPi-AiPd), which demonstrates significant pressure responses in the Southern Hemisphere.

1212 Step 2: Identify pressure anomalies during Southern Hemisphere summer that align with locations of seasonal sea ice.

1213 Step 3: Use context from the caption and figure dynamics to deduce influence from sea ice variations.

1214

1215 **Refined Caption:**

1216 Sea level pressure differences between sensitivity experiments with altered sea surface temperature (SST) gradients in the Atlantic and Pacific Oceans:

1217 ....

1218 •Bottom Panel: The difference between Experiment AdPi (Atlantic high-latitude warming and tropical cooling combined with Pacific tropical warming and high-latitude cooling) and Experiment AiPd (the reverse gradient alterations). This configuration produces a large positive NAO change, with warm tropical Pacific SSTs driving equatorward wave refraction and poleward angular momentum transport, while warm northern North Atlantic SSTs reduce low-altitude northward sensible heat transport and destabilize the local atmosphere. Notable responses are also observed in the Southern Hemisphere, even during summer, which are linked to sea ice variations. Most changes outside of the range of \$-2.2\$ to \$2.2\$-\$\text{mb}\$ are statistically significant at the 95% confidence level.

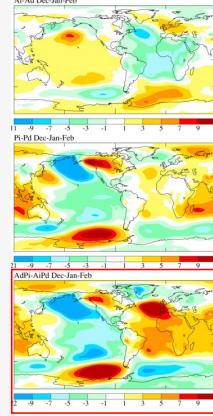


Figure 10: An example of specialized multiple-choice VQA.

1216

1217 **Hard Multiple-choice Question**

1218

1219 **Image Caption:**

1220 MERIS RGB image of scene over Mediterranean Sea from 2 May 2003 (left) and zoom into daily composite of TCWV (middle) and uncertainty of TCWV (right) from MERIS and SSM/I measurements.

1221

1222 **Question:**

1223 In which regions is TCWV density the highest based on the middle panel?

1224

1225 **Options:**

1226 A. Over the Mediterranean Sea

1227 B. Over land near coastal areas

1228 C. Over desert areas in the bottom left

1229 D. In the southernmost water regions

1230

1231 **Response:**

1232 D. In the southernmost water regions

1233

1234 **Reasoning Chain:**

1235 Step 1: Observe the middle panel, which shows TCWV values represented by a color gradient.

1236 Step 2: The highest values correspond to the red and yellow sections in the southernmost part of the image.

1237 Step 3: These sections lie over the near-equatorial regions of the southern water zones where higher water vapour levels are observed.

1238

1239 **Refined Caption:**

1240 MERIS RGB image of the Mediterranean Sea area from 2 May 2003 (left), zoomed-in view showing the daily composite of Total Column Water Vapour (TCWV) (middle), and the uncertainty of TCWV measurements (right). The middle panel illustrates the smooth transition of the water vapour field between land and ocean, with increased uncertainty in coastal areas. This uncertainty is linked to the use of MERIS data to fill gaps in the SSM/I measurements. The right panel highlights the regions of elevated uncertainty, particularly along the coast. Additionally, the figure emphasizes MERIS's high sensitivity to small-scale variations in the water vapour field, as seen over Western Turkey.

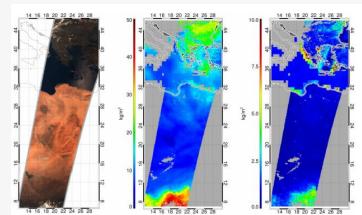


Figure 11: An example of hard multiple-choice VQA.

## E CLASSIFICATION OF RESEARCH PROBLEMS IN EARTH SCIENCES

Under the framework of the five major spheres, we further categorized the generated research problems into specific academic disciplines according to a standardized classification system. Within

Table 10: Main statistics in MSEarth-Bench.

Statistic	Number
Total questions	7,195
MCQ	2,784
Questions with single images	1,255 (45.1%)
Questions with multiple images	1,529 (54.9%)
* Single-image focus	≈1,164 (41.8%)
* Multi-image relational	≈365 (13.1%)
Reasoning Question	2,117 (76.0%)
Perception Question	667 (24.0%)
Open-Ended	1,411
Questions with single images	679 (45.1%)
Questions with multiple images	832 (54.9%)
* Single-image focus	≈619 (41.8%)
* Multi-image relational	≈113 (13.1%)
Captioning	3,000
Average caption length	37.71
Average refined caption length	137.47

the broad category of Earth Sciences, we refined the classification into detailed sub-disciplines or sub-fields. The classification process involves three main steps: first, identifying the primary sphere to which the research problem belongs, selecting from eight major disciplines (referred to as primary spheres), including Atmospheric Sciences, Ecology and Biosciences, Hydrology, Oceanography, Geology, Geography, Solid Earth Geophysics, and Polar Science. Second, the classification is further refined by selecting the most appropriate sub-discipline or sub-field from a detailed hierarchy. Third, for interdisciplinary problems, the primary classification is clearly stated, and any relevant secondary classifications are noted. This hierarchical approach ensures a systematic and precise categorization of research problems, enabling a deeper understanding of their academic and scientific context.

**Summary of Classification:** The classification system includes a total of 8 first-level disciplines and 66 second-level disciplines. Each research problem is assigned to one of the primary disciplines and further refined into a specific sub-discipline based on its characteristics and context.

You are tasked with classifying a research problem into one of the Earth’s spheres and refining it into a specific sub-discipline or sub-field.

#### INSTRUCTIONS:

1. Carefully analyze the input, which includes the research question, paper title, and any additional information derived from images (e.g., visual data descriptions or extracted features).
2. Identify the primary sphere (**Atmospheric Sciences, Ecology and Biosciences, Hydrology, Oceanography, Geology, Geography, Solid Earth Geophysics, or Polar Science**) that the problem belongs to.
3. Refine the classification by selecting the most appropriate sub-discipline or sub-field from the hierarchy.
4. If the problem spans multiple spheres or disciplines, clearly state the primary classification and mention any relevant secondary classifications.

#### CLASSIFICATION HIERARCHY:

**Atmospheric Sciences:** Atmospheric Chemistry, Meteorology, Climatology, Hydrometeorology, Paleoclimatology, Atmospheric Physics, Numerical Weather Prediction and Simulation, Atmospheric Remote Sensing.

**Ecology and Biosciences:** Regional Ecology, Population Ecology, Community Ecology, Ecosystem Ecology, Ecological Engineering, Restoration Ecology, Landscape Ecology, Aquatic Ecology and Limnological Ecology, Biogeochemistry, Biogeography.

1296  
 1297 **Hydrology:** Hydrology, Hydrogeology, Limnology, River Hydrology and Estuarine Hydrology, Groundwater Hydrology, Regional Hydrology, Ecohydrology, Hydrological Physics, Hydrological Geography, Hydrological Meteorology, Hydrological Measurement, Hydrological Cartography.  
 1298  
 1299  
 1300  
 1301 **Oceanography:** Ocean Chemistry, Ocean Physics, Ocean Biology, Ocean Geology, Remote Sensing Oceanography, Environmental Oceanography, Marine Resources Science.  
 1302  
 1303 **Geology:** Economic Geology, Engineering Geology, Environmental Geology, Quaternary Geology, Sedimentology, Stratigraphy, Paleogeography, Volcanology, Mineralogy and Petrology, Regional Geology, Remote Sensing Geology.  
 1304  
 1305 **Geography:** Physical Geography, Human Geography, Regional Geography, Urban Geography, Tourism Geography, World Geography, Historical Geography, Geomorphology, Biogeography, Chemical Geography, Other Disciplines in Geography.  
 1306  
 1307 **Solid Earth Geophysics:** Geodynamics, Seismology, Geomagnetism, Gravimetry, Geoelectricity, Geothermal Science, Tectonophysics, Exploration Geophysics, Computational Geophysics, Experimental Geophysics, Other Disciplines in Solid Earth Geophysics.  
 1308  
 1309 **Polar Science:** Polar Ecology, Polar Oceanography, Glaciology, Permafrost Science, Polar Climate Science.  
 1310  
 1311  
 1312

1313 **INPUT:**

1314  
 1315 • **Paper Title:** {paper\_title}  
 1316 • **Research Question:** {research\_question}  
 1317 • **Image Information:** {image\_caption}

1318 **OUTPUT FORMAT:**

1319 The output must be written in **JSON format** using the structure below:

1320 {  
 1321     "primary\_sphere": ,  
 1322     "primary\_sub\_discipline": ,  
 1323     "secondary\_sphere": "Ecology and Biosciences",  
 1324     "secondary\_sub\_discipline": "Aquatic Ecology"  
 1325 }  
 1326

1327  
 1328 **Table 11: Top Sub-disciplines in Various Scientific Subjects.** The table lists the top three sub-  
 1329 disciplines by count within each major scientific subject.  
 1330

1331

Subject	Top 1 Sub-subject	Top 2 Sub-subject	Top 3 Sub-subject
Hydrology	River Hydrology and Estuarine Hydrology	Groundwater Hydrology	Limnology
Ecology and Biosciences	Aquatic Ecology and Limnological Ecology	Landscape Ecology	Ecosystem Ecology
Geology	Sedimentology	Quaternary Geology	Structural Geology
Solid Earth Geophysics	Seismology	Tectonophysics	Exploration Geophysics
Geography	Physical Geography	Urban Geography	Geomorphology
Polar Science	Glaciology	Polar Climate Science	Permafrost Science
Atmospheric Sciences	Meteorology	Climatology	Atmospheric Remote Sensing
Oceanography	Ocean Physics	Ocean Geology	Environmental Oceanography

1337  
 1338 **F MLLMs' VERSIONS**

1339 For open-source models, we use vllm (Kwon et al., 2023) for local testing; for proprietary models,  
 1340 we conduct tests via API calls. The download paths for specific models and the versions of models  
 1341 accessed via API are provided in Figure 12.  
 1342

1343  
 1344 **G EVALUATION METRICS**

1345  
 1346 **G.1 MLLM-BASED METRICS**

1347 Following G-Eval (Liu et al., 2023b), we utilize MLLM (Qwen2.5-VL-72B) with a specialized  
 1348 prompt to compute a factual scientific score. For the captioning task, we define a Cap-Eval score  
 1349

1350 Table 12: Evaluated MLLMs in our experiments with their versions or Huggingface model paths.  
1351

<i>Open-source Models</i>	
<i>Model</i>	<i>Model path</i>
Qwen2.5-VL-7B	<a href="https://huggingface.co/Qwen/Qwen2.5-VL-7B-Instruct">https://huggingface.co/Qwen/Qwen2.5-VL-7B-Instruct</a>
Qwen2.5-VL-32B	<a href="https://huggingface.co/Qwen/Qwen2.5-VL-32B-Instruct">https://huggingface.co/Qwen/Qwen2.5-VL-32B-Instruct</a>
Qwen2.5-VL-72B	<a href="https://huggingface.co/Qwen/Qwen2.5-VL-72B-Instruct">https://huggingface.co/Qwen/Qwen2.5-VL-72B-Instruct</a>
InternVL2.5-8B	<a href="https://huggingface.co/OpenGVLab/InternVL2_5-8B">https://huggingface.co/OpenGVLab/InternVL2_5-8B</a>
InternVL2.5-78B	<a href="https://huggingface.co/OpenGVLab/InternVL2_5-78B">https://huggingface.co/OpenGVLab/InternVL2_5-78B</a>
InternVL3-78B	<a href="https://huggingface.co/OpenGVLab/InternVL3-78B">https://huggingface.co/OpenGVLab/InternVL3-78B</a>
LLaVA-onvision-72B	<a href="https://huggingface.co/llava-hf/llava-onevision-qwen2-72b-ov-hf">https://huggingface.co/llava-hf/llava-onevision-qwen2-72b-ov-hf</a>
Llama3.2-90B-Vision	<a href="https://huggingface.co/meta-llama/Llama-3.2-90B-Vision">https://huggingface.co/meta-llama/Llama-3.2-90B-Vision</a>
DeepSeek-VL2	<a href="https://huggingface.co/deepseek-ai/deepseek-vl2">https://huggingface.co/deepseek-ai/deepseek-vl2</a>
Intern-S1-mini	<a href="https://huggingface.co/internlm/Intern-S1-mini">https://huggingface.co/internlm/Intern-S1-mini</a>

<i>Proprietary Models</i>	
<i>Model</i>	<i>Model versioning</i>
GPT-4o	GPT-4o-20
Gemini-2.5-Pro-Thinking	gemini-2.5-pro-preview-05-06
Gemini-2.5-Flash	Gemini-2.5-Flash-preview-04-17
Gemini-2.5-Flash-Thinking	Gemini-2.5-Flash-preview-04-17
Claude-3.7-Sonnet	Claude-3.7-Sonnet-20250219
Claude-3.5-Haiku	claude-3-5-haiku-20241022
GPT-4o-mini	GPT-4o-mini-2024-07-18

1369  
1370  
1371 ranging from 1 to 5, where higher scores indicate better caption quality. For the open-ended QA task,  
1372 we introduce OE-Eval, which evaluates the reasonableness of generated answers using a binary 0/1  
1373 scoring system.

1374 The following prompt was used for Cap-Eval:  
1375

1376 Evaluate the quality of a generated caption for a geoscience research paper figure or image.  
1377

#### 1378 EVALUATION CRITERIA:

- 1380 **1. Scientific Accuracy:** Does the generated caption accurately describe the scientific  
1381 content of the figure or image?
- 1382 **2. Clarity and Coherence:** Is the caption well-structured, logically organized, and  
1383 easy to understand?
- 1384 **3. Relevance and Completeness:** Does the caption provide all necessary information  
1385 to understand the figure or image?

#### 1386 EVALUATION STEPS:

- 1388 1. Compare the **Generated Caption** to the **Standard Caption**. Assess whether the  
1389 generated caption aligns with the scientific content and intent of the standard caption.
- 1390 2. Assign a score for coherence on a scale of 1 to 5, where 1 is the lowest and 5 is the  
1391 highest, based on the Evaluation Criteria.

#### 1392 INPUT:

- 1394 • **Standard Caption:** {response}
- 1395 • **Generated Caption:** {generated\_caption}

#### 1397 IMPORTANT INSTRUCTIONS:

- 1398 • Only output the score in the specified JSON format.
- 1399 • Do not provide any explanations, comments, or additional text.

#### 1401 OUTPUT FORMAT:

1402 The output must be written in **JSON format** using the structure below:  
1403

```

1404
1405     {
1406         "score": 1-5
1407     }

```

1408  
1409 The following prompt was used for OE-Eval:

1410  
1411 You are tasked with evaluating the correctness of a generated answer to an open-ended  
1412 question about a given input image.

1413  
1414 INPUT:

1415 • **Question:** {query}  
1416 • **Refined Caption:** {refined caption}  
1417 • **Standard Answer:** {response}  
1418 • **Generated Answer:** {generated\_answer}

1419  
1420 INSTRUCTIONS:

1421 1. Based on the refined caption, question, and standard answer, determine if the generated  
1422 answer is correct.  
1423 2. Only output the determination in the specified JSON format.  
1424 3. Do not provide any explanations, comments, or additional text.

1425  
1426 OUTPUT FORMAT:

1427 The output must be written in **JSON format** using the structure below:

```

1428 {
1429     "is_correct": true or false
1430 }
1431

```

1432 To further establish the correlation between LLMs and human judgment specifically in the domain  
1433 of Earth Science VQA, we conducted a human evaluation with four Ph.D. candidates specializing  
1434 in Earth sciences. They scored a random sample of 160 questions from our MSEarth Open  
1435 Ended benchmark. The models evaluated included Gemini-2.5-Flash, GPT-4o, InternVL3-78B  
1436 and QwenVL2.5-72B. Our inter-annotator agreement, measured by Krippendorff's alpha, is 69.5.  
1437 Following LAVE (Mañas et al., 2024), in order to assess the validity of OE-Eval, we calculated its  
1438 correlation with human judgment using Spearman's rank correlation coefficients. We derive a single  
1439 "quality" score from the 4 binary ratings (correct/incorrect) per answer as follows: 1.0 if at least  
1440 3 annotators rate the answer as correct, 0.5 if only 2 did so, and 0.0 otherwise. The results of this  
1441 evaluation are presented in the following table:

Metric	QwenVL2.5-72B	Gemini-2.5-Flash	GPT-4o	InternVL3-78B	Overall
BERTScore	61.16	60.34	63.06	59.79	61.09
ROUGE	67.04	66.90	70.01	63.29	66.81
METEOR	69.00	67.12	67.34	64.55	66.75
BLEU	59.32	58.96	60.57	57.14	59.00
OE-Eval	68.31	<b>67.13</b>	69.85	<b>63.99</b>	<b>67.32</b>

1449  
1450 Table 13: Spearman correlation across models.

1451 From the table's results, OE-Eval demonstrates a higher consistency with human judgment compared  
1452 to all the considered baselines.

1453  
1454 **G.2 SIMILARITY-BASED METRICS**

1455  
1456 In cases where some models fail to strictly follow instructions and only output the correct answer,  
1457 resulting in regular expression matching failures, we use the all-MiniLM-L6-v2 model (Wang et al.,

2020) to calculate the similarity between the model’s output and each option. The option with the highest similarity is then selected as the model’s answer.

## H DETAILED MSEARTH-MCQ RESULTS

The inputs to our model, MSEarth, consist of images, questions, and the original captions. The original captions provide contextual information about the images, such as the meanings of specific symbols. Therefore, we conducted tests on different models to evaluate their performance with and without the original captions.

Table 14: Accuracies (%) of different models on multiple-choice questions. The best results are highlighted in bold, with the second-best underlined. OC: original caption.

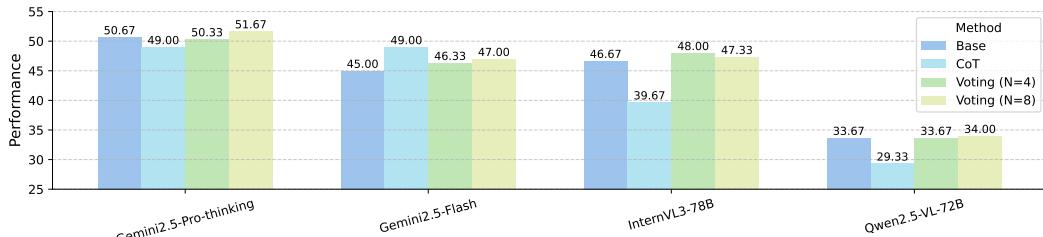
Model	Input OC	SINGLE	Image-Type MULTI	CROSS	Task Type REASONING	Task Type PERCEPT	Overall ACC
<b><i>Open-source Models</i></b>							
LLaVA-onvision-72B	✗	49.40	45.52	41.10	41.92	61.86	46.69
Qwen2.5-VL-7B	✗	39.12	35.65	39.18	37.27	38.98	37.68
Qwen2.5-VL-32B	✗	42.07	39.78	40.00	37.03	52.92	40.84
Qwen2.5-VL-72B	✗	47.65	43.30	43.84	41.43	57.72	45.33
InternVL2-8B	✗	35.94	34.11	32.33	34.25	36.13	34.70
InternVL2.5-78B	✗	48.13	45.88	45.21	43.27	58.02	46.80
InternVL3-78B	✗	51.95	44.85	45.75	44.54	59.67	48.17
Llama3.2-90B-Vision	✗	44.30	40.64	36.16	38.64	51.42	41.70
DeepSeek-VL2	✗	45.42	42.70	46.85	43.74	46.78	44.47
LLaVA-onvision-72B	✓	53.55	49.48	47.95	46.58	65.52	51.11
Qwen2.5-VL-7B	✓	47.65	44.07	37.53	40.53	58.47	44.83
Qwen2.5-VL-32B	✓	52.59	46.99	43.84	42.47	70.16	49.10
Qwen2.5-VL-72B	✓	52.11	50.43	46.30	44.40	70.46	50.65
InternVL2-8B	✓	44.86	43.99	38.36	38.97	58.47	43.64
InternVL2.5-78B	✓	53.23	49.74	44.38	43.17	74.21	50.61
InternVL3-78B	✓	57.53	51.37	45.48	47.00	73.61	53.38
Llama3.2-90B-Vision	✓	45.98	40.46	38.90	38.26	56.97	42.74
DeepSeek-VL2	✓	52.43	49.23	44.66	46.06	62.82	50.07
<b><i>Proprietary Models</i></b>							
Gemini-2.5-Flash	✓	58.33	54.55	<u>53.42</u>	49.98	75.56	56.11
Gemini-2.5-Flash-Thinking	✓	60.64	54.64	<u>53.70</u>	<u>51.35</u>	75.86	57.22
Gemini-2.5-Pro-Thinking	✓	<b>64.78</b>	<b>59.36</b>	55.34	<u>56.31</u>	77.06	<b>61.28</b>
Claude-3.5-Haiku	✓	49.48	47.16	42.47	42.18	64.77	47.59
Claude-3.7-Sonnet	✓	59.52	<u>56.53</u>	<b>57.53</b>	51.68	<u>78.11</u>	<u>58.01</u>
GPT-4o-mini	✓	52.51	48.63	43.01	43.65	68.67	49.64
GPT-4o	✓	<u>63.03</u>	55.76	47.67	50.45	<b>81.86</b>	57.97

For open source models, we performed experiments in settings: with and without the original caption. The results show that providing the original caption improves performance in all tasks. Notably, the improvement is more significant for perception tasks compared to reasoning tasks. This is likely because perception tasks rely more heavily on understanding the image content, and the original caption provides helpful contextual information for interpreting the image.

We have compiled several case studies to illustrate the necessity of the original caption when answering questions in certain situations. In example 14, if the original caption is not provided, InternVL3-78B will be unable to accurately determine that the geographical location is in Germany, resulting in an incorrect answer. In contrast, some proprietary models may possess stronger perceptual capabilities and can correctly identify the location as Germany even without the original caption. Similarly, in example 15, providing the original caption aids the model in understanding the image, thereby facilitating task completion. Both scenarios are prevalent in scientific question-answering contexts. To address this, we conducted separate experiments and explicitly integrated these settings into the design of the MSEarth-MCQ task.

## 1512 I RESULTS WITH COMPUTE SCALING

1514 From the main experiments, it is evident that the performance of various models declines significantly  
 1515 on questions requiring specialized knowledge. To explore whether existing methods can enhance  
 1516 model performance on such questions, we sampled 300 specialized questions from the MCQ dataset  
 1517 to create the MSEarth-mini set. We then evaluated the effectiveness of Chain-of-Thought (CoT)  
 1518 reasoning and majority voting mechanism, which selects the most frequent response among (N)  
 1519 candidate responses; in the case of a tie, one of the most frequent answers is randomly chosen. The  
 1520 results are presented in figure 13. Notably, for the Gemini-Pro-thinking model, which inherently  
 1521 incorporates a thinking mechanism, introducing CoT reasoning led to a decline in performance.  
 1522 Similarly, for some open-source models, such as Qwen and InternVL, the addition of CoT reasoning  
 1523 also resulted in performance degradation. However, the majority voting mechanism proved effective  
 1524 for most models.



1534 Figure 13: Performance of different strategies on MSEarth-mini.

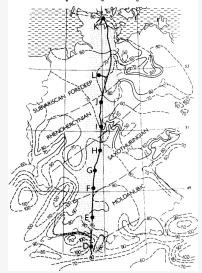
1536 **Raw Caption:** Heat-flow pattern in Germany and in adjacent areas (after Hurtig *et al.*, 1992) together with the position of the  
 1537 Central Segment of the European Geotraverse and shotpoints \$D-K\$

1538 **Question:** Which region along the EGT segment has the most sparse heat-flow data?

1539 A. North German Lowland  
 1540 B. Molasse Basin  
 1541 C. Saxothuringian  
 1542 D. Rhenohercynian

1543 **Answer:** A

1544 **Refined Caption:** Heat-flow pattern in Germany and adjacent areas, illustrating the position of the Central Segment of the  
 1545 European Geotraverse (EGT). The EGT extends from the southern rim of the Molasse Basin near Chur, northward to the Baltic  
 1546 Sea near Kiel. This section traverses significant geological formations, including the Alpine Molasse Basin, Hercynian (Variscan)  
 1547 units of the Swabian Jura, South German Basin, Mid-German Crystalline Rise, Hessian Depression, and Leine Graben,  
 1548 terminating in the North German Lowland. The heat-flow data, particularly sparse in the North German Basin, range from values  
 1549 north of Hamburg to east of Hannover. The pattern depicted is based on previously obtained temperature data. Shotpoints D-K  
 1550 are marked along the traverse.



1551 **InternVL3-78B w/o original caption:** A - The North German Lowland region shows fewer heat-flow contour lines compared to other regions, indicating  
 1552 sparser data. The contour lines are less dense, suggesting fewer measurements or less variation recorded, which implies sparse data collection in this area. ✓

1553 **InternVL3-78B w/o original caption:** B - The Molasse Basin region shows fewer heat-flow data points compared to other regions along the EGT segment.  
 1554 The contour lines are less dense, indicating sparser data collection in this area. ✗

1555 **gemini-2.5-flash w/o caption:** A - The image shows heat flow contours (dashed lines) and geological regions. The EGT segment is represented by the thick  
 1556 black line with points D through K. We need to identify the region along the EGT segment with the most sparse heat-flow data, which is indicated by the  
 1557 density of heat flow contours. Examining the map, the northernmost part of the EGT segment, around point K and extending further north, is likely  
 1558 within the North German Lowland. In this area, the heat flow contours are relatively sparse (e.g., 60, 70). In contrast, the Rhenohercynian region along  
 1559 the EGT (from G to J) shows a higher density of contours (60, 70, 80, 90). The Saxothuringian region along the EGT (around I and J) also has several  
 1560 contours (50, 60, 70, 80). The Moldanubic region (D to E) has a high density of contours (80, 90, 100, 110, 120). The Molasse Basin is not shown along  
 1561 the EGT segment. Comparing the density of contours along the EGT in the given options, the North German Lowland (northernmost part) exhibits the  
 1562 most sparse heat-flow data. ✓

1563 **gpt-40-2024-11-20 w/o caption:** A - The North German Lowland region along the EGT segment has the most sparse heat-flow data, as indicated by the  
 1564 lack of detailed contour lines and data points in this area compared to the other regions. The other regions, such as the Molasse Basin, Saxothuringian,  
 1565 and Rhenohercynian, show more detailed heat-flow contour lines and data points, indicating denser data coverage. ✓

1566 Figure 14: Performance comparison of different models under two settings: with and without the  
 1567 original caption. The results indicate that more powerful models exhibit less dependence on the  
 1568 original caption, highlighting their ability to interpret image content independently.

## 1569 J HUMAN PERFORMANCE BASELINE

1570 To clarify benchmark difficulty and further justify its educational relevance, we also included human  
 1571 expert scores in MSEarth-Bench-mini. Specifically, we hired three human experts, all of whom are  
 1572 Ph.D. students with backgrounds in Earth sciences, to evaluate the tasks. We report their average  
 1573 scores to provide a clear baseline for human performance:

1566

**Raw Caption:** Defining the mud volcano boundary, layers, and faults on the CDS stack migrated section. Black rectangle shows part of the section showing in subsection in Fig. 14 b.

1567

**Question:** What distinguishes the NW section of the Figure from the SE section?

1568

- A. The NW section shows more layered reflectors.
- B. The SE section contains more faults and disrupted layers
- C. The NW section has deeper visible structures
- D. The SE section is closer to the surface.

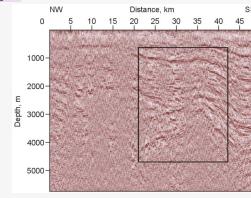
1571

**Answer:** B

1572

**Refined Caption:** Migrated section of the common diffraction surface (CDS) stacked data, illustrating the boundary, internal layers, and faults of the mud volcanoes. The left portion of the section, which is of particular interest, shows an improved imaging of the mud volcano boundary. Reflectors and other subsurface structures responsible for diffraction events in the stacked section are clearly visible, providing enhanced detail and clarity of the geological features.

1574



1575

**InternVL3-78B w original caption:** B - The SE section within the black rectangle shows more disrupted layers and faults compared to the NW section, which appears to have more continuous and less disturbed reflectors. This indicates that the SE section contains more structural disturbances.



1576

**InternVL3-78B w/o original caption:** A - The NW section of the figure shows more distinct and continuous layered reflectors compared to the SE section, which has more disrupted and less continuous layers. This indicates that the NW section has more pronounced stratification.



1577

**gemini-2.5-flash w/o caption:** A - The image displays a seismic section. The NW section, particularly in the upper part, shows more continuous and distinct layered reflectors compared to the SE section, which appears more complex with folding and potentially faulting within the boxed area. Option B is incorrect as the SE section, while complex, doesn't necessarily show more faults and disrupted layers than the NW section, and the NW section also exhibits some disruption. Option C is incorrect as both sections show structures down to similar depths. Option D is incorrect as the depth scale is the same for both sections, and the SE section is not closer to the surface than the NW section.



1578

**gemini-2.5-pro w/o caption:** B - Upon examining the seismic reflection profile, the SE section (roughly from 20 km to 45 km distance) displays prominent folding of the geological layers. Associated with this folding, there are also clear indications of faults, such as offsets and abrupt terminations of reflectors. This intense deformation results in disrupted layers. In contrast, the NW section (roughly from 0 km to 20 km distance) shows layers that are generally flatter or gently dipping, with less evidence of significant folding or faulting. Therefore, the SE section is distinguished by containing more faults and disrupted layers compared to the NW section.



1584

**gpt-4o-2024-11-20 w/o caption:** B - In the image, the NW section shows relatively continuous and layered reflectors, while the SE section displays more disrupted and faulted layers, as evidenced by the irregular and broken patterns in the SE region. This indicates the presence of faults and structural deformation in the SE section.



1586

Figure 15: Performance comparison of different models under two settings: with and without the original caption.

1590

Model	Atmospheric	Solid Earth Geophysics	Geography	Ecology	Geology	Hydrology	Oceanography	Polar	All
InternVL3-78B	50.70%	29.73%	28.57%	47.06%	25.00%	51.02%	25.00%	30.00%	47.33%
gemini-2.5-pro	46.48%	58.11%	35.00%	35.71%	47.06%	50.00%	59.18%	50.00%	51.33%
o4-mini	50.00%	45.00%	55.88%	71.43%	49.30%	48.98%	43.33%	62.50%	53.00%
Expert	86.49%	85.00%	85.29%	92.86%	87.32%	85.71%	86.67%	87.50%	87.00%

1594

Table 15: Accuracy on MSEarth-Bench-mini across Earth science domains. Human expert scores are averages of three Ph.D.-level Earth science evaluators.

1595

The results clearly demonstrate that human experts consistently outperform the current MLLMs across all Earth science domains.

1600

## K IMPACT OF EXPLICIT REASONING

1603

We assess the impact of explicit chain-of-thought (CoT) prompting on MSEarth-Bench-mini across both open-source and proprietary LVLMs. For open-source models, we compare InternVL3 and QwenVL2.5; for proprietary models, we examine the Gemini-2.5-Flash series. Within the proprietary family, variants with dedicated “thinking” capabilities (e.g., Gemini-2.5-Flash-Thinking) generally outperform counterparts without such capabilities (e.g., Gemini-2.5-Flash). In contrast, for open-source models, adding explicit CoT sometimes leads to performance declines, which we hypothesize stems from limited training for explicit reasoning behaviors (e.g., GRPO-style preference optimization).

1611

To further probe the role of explicit CoT, we include GPT-o4-mini, which exposes configuration options that control reasoning depth (low, medium, high), roughly corresponding to the length of the reasoning chain. Results are shown in Table 16.

1614

Overall, we observe the following:

1616

Models explicitly equipped and trained for “thinking” benefit from enabling CoT (e.g., Gemini-2.5-Flash-Thinking). When a model already exhibits strong inherent reasoning, additional explicit CoT can reduce performance, as seen in o4-mini at medium/high reasoning depths. Open-source models do not consistently benefit from CoT without targeted training for reasoning behaviors, suggesting a direction for future supervised and RL post-training.

Model	CoT (Accuracy %)	Non-CoT (Accuracy %)
Gemini-2.5-Pro	50.67%	52.33%
Gemini-2.5-Flash-no-think	42.00%	40.00%
Gemini-2.5-Flash-Thinking	52.00%	46.00%
o4-mini (low)	52.00%	51.00%
o4-mini (medium)	50.67%	53.00%
o4-mini (high)	50.33%	54.33%

Table 16: Effect of explicit chain-of-thought (CoT) prompting on MSEarth-Bench-mini. Higher is better; values are accuracy (%). For o4-mini, low/medium/high denote shorter-to-longer reasoning traces.

Model	Atmospheric Sciences			Ecology and Biosciences		
	Meteor.	Climat.	Atmos. RS	Ecosys. Ecol.	Landsc. Ecol.	Aquat. Ecol.
InternVL-8B	0.4039	0.3643	0.5238	0.3833	0.4792	0.3030
InternVL-78B	0.4624	0.3643	0.5476	0.6333	0.5208	0.6061
InternVL3-78B	0.4847	0.3857	0.4762	0.6333	0.5417	0.6364
Qwen2.5-VL-72B	0.4903	0.4071	0.4524	0.4500	0.4375	0.5455
Claude-3.7-Sonnet	0.5599	0.4714	0.6429	0.6500	0.6042	0.5758
Gemini-2.5-Pro-Thinking	0.5877	0.5071	0.5714	0.7000	0.5833	0.6667
GPT-4o	0.5097	0.4429	0.5714	0.5667	0.5208	0.6667
GPT-4o-mini	0.4457	0.3857	0.5000	0.6333	0.4375	0.5455
Gemini-2.5-Flash-Thinking	0.5265	0.3929	0.5476	0.7500	0.5833	0.6364
Gemini-2.5-Flash	0.5153	0.4000	0.5952	0.6167	0.5625	0.5455
Intern-S1	0.6320	0.5870	0.6750	0.7350	0.6580	0.7020
Intern-S1-mini-MSEarth	0.6080	0.5630	0.6420	0.7120	0.6350	0.6780
Intern-S1-mini	0.5850	0.5310	0.6180	0.6870	0.6090	0.6450

Table 17: Model Performance on Primary and Sub-Disciplines of Earth Science (Accuracy)

## L MORE RESULTS

For MCQ and OE questions, we used radar charts to illustrate the performance of various models across different disciplines. We also give some case studies in Figure 17 and Figure 18. We also present detailed performance breakdowns of all models across every sub-discipline of Earth science in Tables 20.

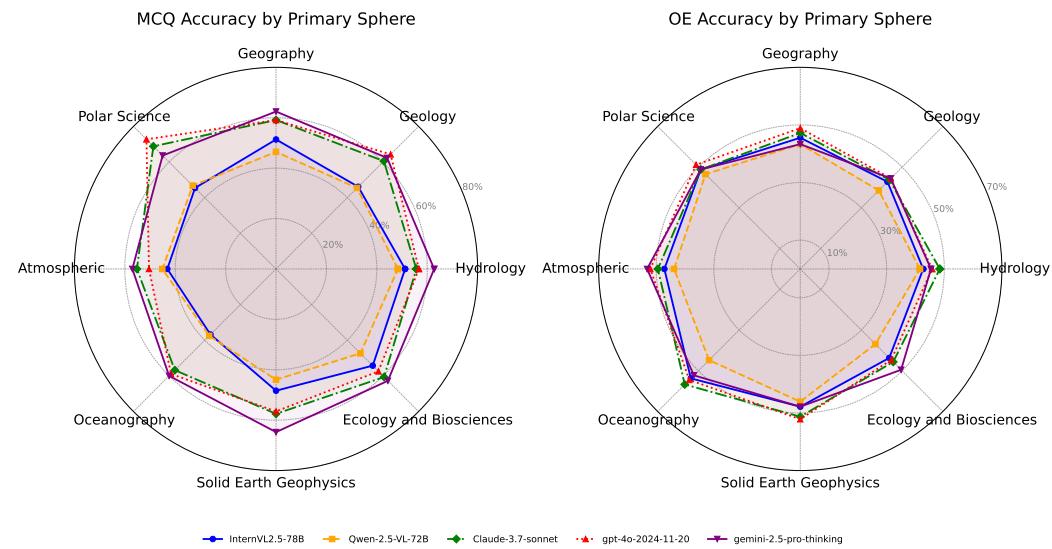


Figure 16: Performance comparison of different models across various subjects.

1674

1675

Model	Geography			Geology		
	Phys. Geog.	Urban Geog.	Reg. Geog.	Sediment.	Struct. Geol.	Quat. Geol.
InternVL-8B	0.4765	0.4444	0.4444	0.5090	0.4268	0.4878
InternVL-78B	0.5451	0.4815	0.5556	0.5663	0.4634	0.5610
InternVL3-78B	0.5884	0.3333	0.6667	0.5986	0.5366	0.6585
Qwen2.5-VL-72B	0.5307	0.3704	0.6667	0.5305	0.5000	0.6098
Claude-3.7-Sonnet	0.5993	0.4074	0.8889	0.6057	0.6341	0.4878
Gemini-2.5-Pro-Thinking	0.6354	0.5556	0.5556	0.6487	0.5854	0.6341
GPT-4o	0.5884	0.5556	0.7778	0.6703	0.6341	0.6829
GPT-4o-mini	0.5090	0.4815	0.6667	0.5269	0.5488	0.5122
Gemini-2.5-Flash-Thinking	0.6426	0.5556	0.7778	0.5986	0.5976	0.5122
Gemini-2.5-Flash	0.5921	0.5185	0.7778	0.5771	0.5366	0.5854
Intern-S1	0.6830	0.6250	0.9120	0.6970	0.6780	0.7250
Intern-S1-mini-MSEarth	0.6590	0.5980	0.8850	0.6730	0.6540	0.6980
Intern-S1-mini	0.6320	0.5640	0.8530	0.6450	0.6210	0.6670

Table 18: Model Performance on Primary and Sub-Disciplines of Earth Science (Accuracy, Continued)

1689

1690

1691

Model	Hydrology			Oceanography		
	River Hydrol.	Groundw. Hydrol.	Limnol.	Ocean Phys.	Ocean Geol.	Env. Oceanogr.
InternVL-8B	0.4550	0.4297	0.4348	0.3800	0.4762	0.2941
InternVL-78B	0.5500	0.4688	0.5652	0.4800	0.5714	0.3529
InternVL3-78B	0.5400	0.4766	0.5652	0.5250	0.6190	0.5882
Qwen2.5-VL-72B	0.5850	0.5078	0.5435	0.4900	0.6667	0.3529
Claude-3.7-Sonnet	0.6150	0.5391	0.5435	0.5650	0.7143	0.3529
Gemini-2.5-Pro-Thinking	0.6700	0.6172	0.6739	0.5750	0.6667	0.5882
GPT-4o	0.6200	0.5625	0.5217	0.5800	0.7619	0.4706
GPT-4o-mini	0.5900	0.4766	0.5435	0.4650	0.5238	0.4706
Gemini-2.5-Flash-Thinking	0.6250	0.5625	0.5435	0.5000	0.6667	0.4706
Gemini-2.5-Flash	0.5700	0.4922	0.6304	0.4800	0.5238	0.3529
Intern-S1	0.7050	0.6680	0.6970	0.6320	0.7950	0.6230
Intern-S1-mini-MSEarth	0.6820	0.6430	0.6720	0.6080	0.7680	0.5950
Intern-S1-mini	0.6570	0.6150	0.6450	0.5830	0.7320	0.5680

Table 19: Model Performance on Primary and Sub-Disciplines of Earth Science (Accuracy, Continued)

1705

1706

1707

Model	Polar Science			Solid Earth Geophysics		
	Glaciol.	Permafrost Sci.	Polar Ocean	Seismol.	Tectonophys.	Geomagn.
InternVL2.5-8B	0.4571	0.7500	0.0000	0.4248	0.5625	0.5455
InternVL2.5-78B	0.4571	0.5000	1.0000	0.4902	0.4375	0.5455
InternVL3-78B	0.5286	0.5000	1.0000	0.5033	0.5000	0.6364
Qwen2.5-VL-72B	0.6286	0.5000	1.0000	0.4706	0.3750	0.7273
Claude-3.7-Sonnet	0.7000	0.5000	1.0000	0.5752	0.5625	0.6364
Gemini-2.5-Pro-Thinking	0.6286	1.0000	1.0000	0.6863	0.5000	0.8182
GPT-4o	0.7286	0.7500	1.0000	0.5163	0.6250	0.9091
GPT-4o-mini	0.5571	0.7500	1.0000	0.4379	0.4375	0.6364
Gemini-2.5-Flash-Thinking	0.6571	0.7500	1.0000	0.5359	0.6250	0.8182
Gemini-2.5-Flash	0.6812	0.7500	1.0000	0.6013	0.5000	0.5455
Intern-S1	0.7650	1.0000	1.0000	0.7230	0.6850	0.9320
Intern-S1-mini-MSEarth	0.7380	0.9500	1.0000	0.6970	0.6580	0.9050
Intern-S1-mini	0.7120	0.9000	1.0000	0.6650	0.6230	0.8780

Note: 1. Sub-disciplines listed are the top 3 with the largest sample size in each primary discipline;

2. Abbreviations: Meteor.=Meteorology, Climat.=Climatology, Atmos. RS=Atmospheric Remote Sensing,

Ecosys. Ecol.=Ecosystem Ecology, Landsc. Ecol.=Landscape Ecology, Aquat. Ecol.=Aquatic &amp; Limnological Ecology,

Phys. Geog.=Physical Geography, Reg. Geog.=Regional Geography, Sedimentol.=Sedimentology,

Struct. Geol.=Structural Geology, Quat. Geol.=Quaternary Geology, River Hydrol.=River &amp; Estuarine Hydrology,

Groundw. Hydrol.=Groundwater Hydrology, Limnol.=Limnology, Env. Oceanogr.=Environmental Oceanography,

Glaciol.=Glaciology, Seismol.=Seismology, Tectonophys.=Tectonophysics, Geomagn.=Geomagnetism.

Table 20: Model Performance on Primary and Sub-Disciplines of Earth Science (Accuracy, Continued)

1724

1725

1726

1727

1728

1729

1730

1731

1732

1733

1734

1735

1736

1737

1738

1739

1740

1741

1742

1743

1744

1745

1746

1747

1748

1749

1750

1751

1752

1753

1754

1755

1756

1757

1758

1759

1760

1761

1762

1763

1764

1765

1766

1767

1768

1769

1770

1771

1772

1773

1774

1775

1776

1777

1778

1779

1780

1781

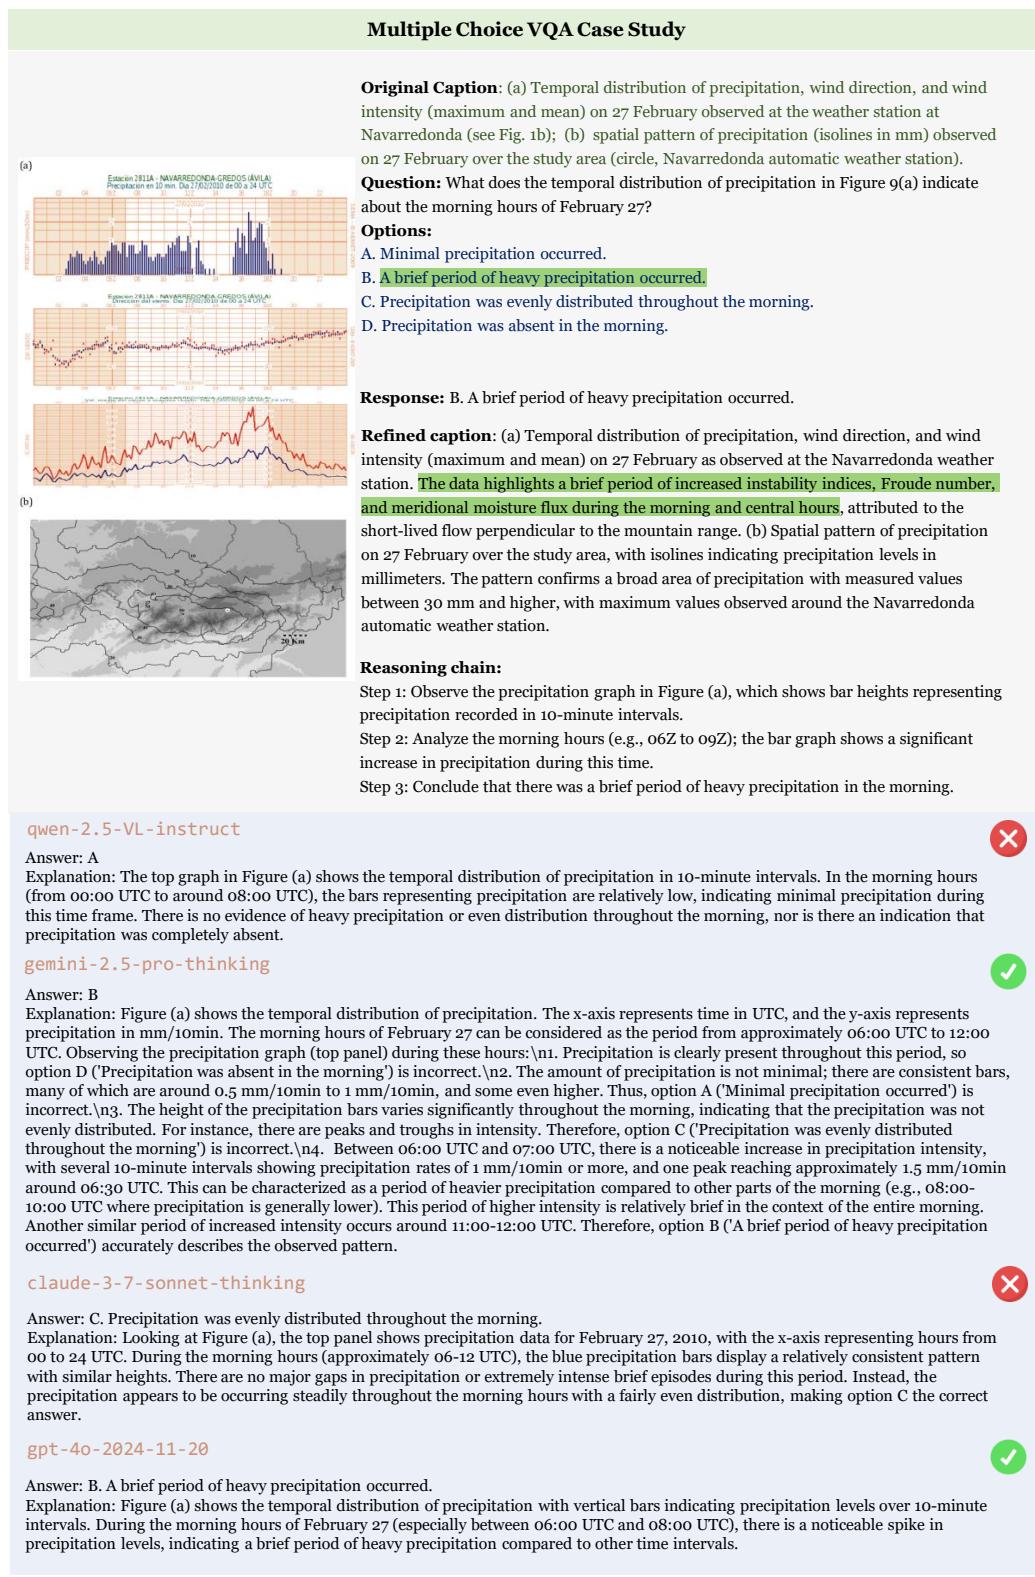


Figure 17: Case Study of Multiple Choice VQA.

1782  
 1783  
 1784  
 1785 **Open-Ended VQA Case Study 1**

1786 

1787 **Original Caption:** Schematic illustration of locations of the photometric stations.

1788 **Question:** Which station lies in a semi-arid steppe zone?

1789 **Response:** Baruun-Urt

1790 **Reasoning chain:**

1791 Step 1: From the Figure, Baruun-Urt is outside the Gobi Desert's shaded area, suggesting it is not in the arid region.

1792 Step 2: The semi-arid steppe zone is consistent with areas having denser vegetation, implied by Baruun-Urt's location in the Figure.

1793  
 1794  
 1795 **Refined caption:** Schematic illustration of the locations of the photometric stations used for atmospheric AOD measurements in the Eastern Gobi region of Mongolia. The stations are situated at approximately the same latitude and include Sainshand (located in an arid zone), Dalanzadgad (also in an arid zone), and Baruun-Urt (located in a semi-arid steppe zone with denser vegetation cover).

1796  
 1797  
 1798  
 1799 **qwen-2.5-VL-instruct** 

1800 Answer: Dalanzadgad

1801 Explanation: Located in semi-arid area 

1802 **gemini-2.5-pro-thinking** 

1803 Answer: Baruun-Urt

1804 Explanation: The Gobi Desert is an arid zone. Baruun-Urt is located north of the Gobi Desert, in an area likely to be a semi-arid steppe transition zone. 

1805 **claude-3-7-sonnet** 

1806 Answer: Baruun-Urt

1807 Explanation: Located at northern edge of Gobi Desert region where desert transitions to semi-arid steppe ecological zone. 

1808 **gpt-4o-2024-11-20** 

1809 Answer: Sainshand,

1810 Explanation: The map shows Sainshand located within the semi-arid steppe zone of the Gobi Desert. 

1811 **Open-Ended VQA Case Study 2**

1812 

1813 **Original Caption:** Map of Gösing with marked earth buildings from the cadastre.

1814 **Question:** Which natural feature primarily impacts settlement patterns visible in the map?

1815 **Response:** Terrain slope 

1816 **Reasoning chain:**

1817 Step 1: Observing the map shows that areas near slopes have higher earth building clusters.

1818 Step 2: Caption reveals that terrain influences settlement and construction choices.

1819 Step 3: Terrain slope emerges as the most significant natural factor influencing patterns.

1820  
 1821 **Refined caption:** Detailed map of Gösing illustrating the distribution of earth buildings as recorded in the earth building cadastre. The map highlights how the topography of the terrain influences the settlement patterns of these structures, providing insight into the construction techniques and methods that shape the appearance of entire streets within the village.

1822  
 1823 **gemini-2.5-pro-thinking** 

1824 Answer: Terrain or topography

1825 Explanation: The map shows buildings (settlements) concentrated along roads, and the presence of dashed lines, likely representing contour lines, suggests varying elevation. Settlements often follow flatter areas or valleys in hilly terrain for easier construction and access. Therefore, the terrain/topography is the primary natural feature influencing where settlements are located. 

1826  
 1827 **qwen-2.5-VL-instruct** 

1828 Answer: River proximity

1829 Explanation: Settlements cluster near river. 

1830 **gpt-4o-2024-11-20** 

1831 Answer: Terrain elevation differences

1832 Explanation: The map shows settlement patterns concentrated along roads and flatter areas, avoiding steeper or elevated terrain, which is typical in areas with significant elevation differences. 

1833  
 1834  
 1835 **Figure 18: Case Study of Open-Ended VQA.**