

# SeedBench: A Multi-task Benchmark for Evaluating Large Language Models in Seed Science

Jie Ying<sup>1\*</sup> Zihong Chen<sup>1\*</sup> Zhefan Wang<sup>1\*</sup> Wanli Jiang<sup>1</sup>

Chenyang Wang<sup>1</sup> Zhonghang Yuan<sup>1</sup> Haoyang Su<sup>1</sup>

Huanjun Kong<sup>1</sup> Fan Yang<sup>2</sup> Nanqing Dong<sup>1,3†</sup>

<sup>1</sup>Shanghai Artificial Intelligence Laboratory <sup>2</sup>Yazhouwan National Laboratory

<sup>3</sup>Shanghai Innovation Institute

## Abstract

Seed science is essential for modern agriculture, directly influencing crop yields and global food security. However, challenges such as interdisciplinary complexity and high costs with limited returns hinder progress, leading to a shortage of experts and insufficient technological support. While large language models (LLMs) have shown promise across various fields, their application in seed science remains limited due to the scarcity of digital resources, complex gene-trait relationships, and the lack of standardized benchmarks. To address this gap, we introduce SeedBench<sup>1</sup>—the first multi-task benchmark specifically designed for seed science. Developed in collaboration with domain experts, SeedBench focuses on seed breeding and simulates key aspects of modern breeding processes. We conduct a comprehensive evaluation of 26 leading LLMs, encompassing proprietary, open-source, and domain-specific fine-tuned models. Our findings not only highlight the substantial gaps between the power of LLMs and the real-world seed science problems, but also make a foundational step for research on LLMs for seed design.

## 1 Introduction

Food security is a fundamental global concern, with seeds serving as the foundation of agricultural production. However, the seed industry faces significant challenges, including its inherently interdisciplinary nature and low economic returns. These factors contribute to a persistent shortage of skilled breeding scientists, a trend expected to continue over the next decade (Egan et al., 2024). The critical shortage of agricultural scientists directly constrains productivity improvements and the sustainable growth of food production.

With the rise of artificial intelligence (AI), advanced AI techniques are transforming seed science. The integration of AI-driven solutions into seed breeding aligns with projections that the next agricultural revolution will be driven by smart, digital, and precision agricultural technologies (Iversen et al., 2021). Large language models (LLMs), in particular, offer the ability to process vast amounts of genetic, environmental, and agronomic data, optimizing crop development. However, effectively assessing and comparing LLM capabilities requires high-quality evaluation benchmarks.

Despite the availability of LLM benchmarks for general purpose, none have been specially developed for seed breeding, which is a field critical to agricultural production and food security. Progress in this domain has been slow due to a shortage of breeding experts and limited availability of online resources. While existing agricultural benchmarks, such as AgEval (plant stress phenotyping) (Arshad et al., 2024) and AgXQA (agricultural extension Q&A) (Kpodo et al., 2024), contribute to LLM evaluation in agriculture, they fail to address the complex decision-making and multi-step processes unique to seed breeding. The absence of a dedicated benchmark limits the ability to systematically measure LLM performance in this domain.

To bridge this gap, we introduce **SeedBench**, a multi-task benchmark designed to simulate expert decision-making across three essential seed breeding stages: (1) gene information retrieval; (2) gene function and regulation analysis; and (3) variety breeding with agronomic trait optimization. Each task category is carefully designed with information-rich, expert-validated question-answer pairs, ensuring the benchmark aligns with real-world seed breeding challenges. By providing a structured evaluation framework, SeedBench enables rigorous assessment of whether LLMs can assist human experts, accelerate breeding workflows, optimize outcomes, and advance towards

\*Equal contribution.

†Corresponding author.

<sup>1</sup><https://github.com/open-sciencelab/SeedBench>

autonomous intelligent seed breeding.

As the first and most comprehensive benchmark in this field, SeedBench systematically connects the capabilities of LLMs with real-world breeding challenges. Developed by experts with interdisciplinary backgrounds and Ph.D.-level expertise, SeedBench ensures both domain relevance and scientific rigor. Each question undergoes two-tier validation and refinement through machine-based assessment and human expert review, which guarantees accuracy and reliability. To accommodate diverse evaluation scenarios, SeedBench supports both one-shot and zero-shot formats, enabling a comprehensive assessment of LLMs in seed breeding.

We evaluate the performance of 26 leading LLMs on SeedBench, including 7 proprietary LLMs, 16 open-source LLMs, and 3 domain-specific models. Recognizing that LLM performance is highly sensitive to prompt structures, we further analyze multiple prompt templates to enhance robustness. Our evaluation aims to address the following research questions.

- RQ1.** What is the relationship between the reasoning ability of LLMs and the performance on seed breeding tasks?
- RQ2.** Do domain-specific fine-tuned models outperform general models in seed breeding?
- RQ3.** What is the ideal model size for seed breeding tasks?

Answering these research questions can not only facilitate the research on LLMs on seed breeding, an emerging topic, but also the understanding of usage of LLMs in the field of AI for Science.

Our main contributions are as follows:

- SeedBench is the first benchmark designed to evaluate LLMs in seed science.
- SeedBench covers key seed breeding processes, ensuring reliability and accuracy through expert validation.
- Extensive evaluations of LLMs are conducted to identify their strengths and limitations, providing insights for future AI advances in breeding.

In the following sections, Section 2 reviews related work on LLMs and agricultural benchmarks. Section 3 describes the construction and methodology behind SeedBench, outlining its design, task types, and validation process. Section 4 presents the experimental setup, including evaluated models and performance comparisons. Finally, we discuss findings, limitations, and future research directions for LLMs in seed breeding.

## 2 Related Work

### 2.1 Domain-Specific LLMs

Based on the Transformer architecture (Vaswani et al., 2017), language models have rapidly advanced, achieving key milestones in their development. Starting with foundational models such as BERT (Devlin, 2018) and GPT (Radford, 2018), subsequent breakthroughs like GPT-4 (Achiam et al., 2023) and DeepSeek-R1 (Guo et al., 2025) have demonstrated exceptional text generation. By integrating specialized domain knowledge with continual pre-training and supervised fine-tuning, LLMs have shown potential in domains such as education (Gan et al., 2023), finance (Li et al., 2023b), and healthcare (Mumtaz et al., 2024). In the agricultural domain, LLMs are increasingly recognized for their potential to enhance food production and optimize agricultural management (De Clercq et al., 2024; Kuska et al., 2024). However, in contrast to the progress in these domains, leveraging LLMs for seed breeding remains an under-explored challenge in the field of AI for Science.

### 2.2 Domain-Specific Benchmarks

To effectively assess and compare LLM capabilities, high-quality evaluation benchmarks are essential. While general benchmarks now cover a wide range of areas, including causal inference (Wang, 2024), instruction following (Zhou et al., 2023), and safety (Lin et al., 2022), domain-specific benchmarks have also emerged. These specialized benchmarks are found in fields such as finance (Xie et al., 2024), geography (Li et al., 2023a), healthcare (Chen et al., 2024a), and law (Fei et al., 2024). These benchmarks help drive improvements by creating diverse datasets and tasks tailored to specific applications. However, many of these benchmarks rely heavily on website data and GPT-based re-annotation, which may limit the data diversity and depth of expertise, raising concerns about the reliability and accuracy of the results.

In addition to data limitation, another challenge is the difficulty of tasks, especially scientific tasks that require additional domain knowledge. Existing agricultural benchmarks does not cover the topic of seed breeding, a critical task in agriculture. For example, AgEval (Arshad et al., 2024) aims for plant stress phenotyping, AgXQA (Kpodo et al., 2024) focuses on agricultural extension, and CROP (Zhang et al., 2024) evaluates the crop knowledge. The disparity underscores the necessity

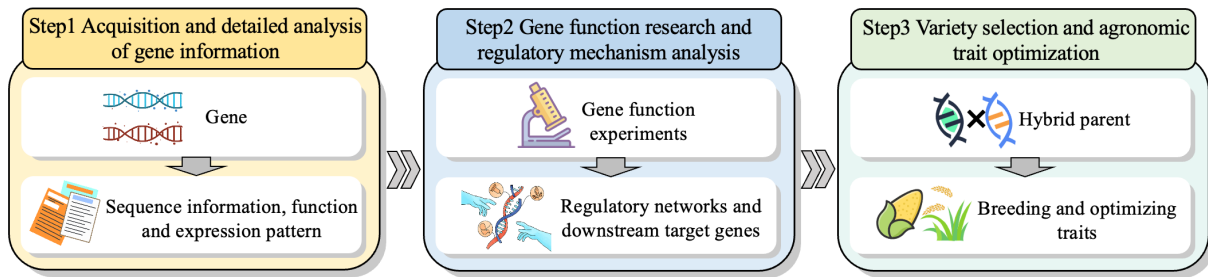


Figure 1: Breeding Expert Workflow Framework. We establish benchmark construction principles by consulting domain experts to replicate real-world seed breeding decision-making processes. (1) Gene Information Retrieval, utilizing established databases to obtain gene sequences and expression patterns; (2) Gene Function & Regulation, employing experimental approaches (*e.g.*, gene knockout, overexpression) to investigate gene roles in plant development; and (3) Variety Breeding & Trait Optimization, implementing breeding techniques (*e.g.*, hybridization, backcrossing) combined with agronomic trait selection for stable variety development.

for specialized benchmarks to evaluate the performance of LLMs in seed science.

### 3 SeedBench

This section outlines the design principles of SeedBench and details its 11 evaluation task types. SeedBench systematically evaluates LLMs in seed breeding by aligning with the typical workflow of breeding experts. It is structured into three primary categories: gene information retrieval, gene function and regulation, and variety breeding and agronomic traits. This taxonomy ensures comprehensive knowledge coverage and skill assessment.

#### 3.1 Breeding Process Overview

Breeding experts typically follow three key steps in seed selection (Figure 1; further details in Appendix A.1). These steps form the basis for SeedBench’s task categorization.

#### 3.2 Task Taxonomy

The tasks in SeedBench are categorized into three main areas, corresponding to the seed breeding workflow. These categories are further divided into ten subcategories in total, ensuring a systematic evaluation of LLM capabilities (Figure 2). The categorization follows the practice in seed science with the help from domain experts (Copeland and McDonald, 2012).

##### 3.2.1 Gene Information Retrieval

LLMs retrieve essential gene information linked to specific traits, including key genes in biological processes, gene sequences, functional descriptions, and expression patterns across environments and developmental stages. They also determine the intracellular localization of gene products, mapping

their distribution within the nucleus, cytoplasm, or membrane. These tasks corresponds to the initial step of the breeding process. Specific tasks include:

- Gene Basic Information Query
- Gene Expression Pattern Query
- Gene Product Cellular Localization Query

##### 3.2.2 Gene Function and Regulation

LLMs describe gene functions under specific experimental conditions using available data, analyze gene product regulation of downstream genes and pathways, and predict functions of uncharacterized genes. These tasks align with the second step of the breeding process. Specific tasks include:

- Gene Function Experimental Observation
- Gene Product Regulation of Downstream Genes Analysis
- Gene Function Prediction

##### 3.2.3 Variety Breeding and Agronomic Traits

LLMs gather information on breeding history, methods, objectives, and agronomic traits, including disease resistance, yield, and drought tolerance. They also propose suitable planting regions based on environmental factors. This aligns with the third step of the breeding process. Specific tasks include:

- Variety Breeding Process Query
- Variety Agronomic Trait Query
- Variety Cultivation and Technical Key Points Query
- Variety Suitable Planting Area Recommendation

#### 3.3 Benchmark Construction

SeedBench was developed through a structured three-step process: data collection, automatic question generation, and two-tier quality validation.

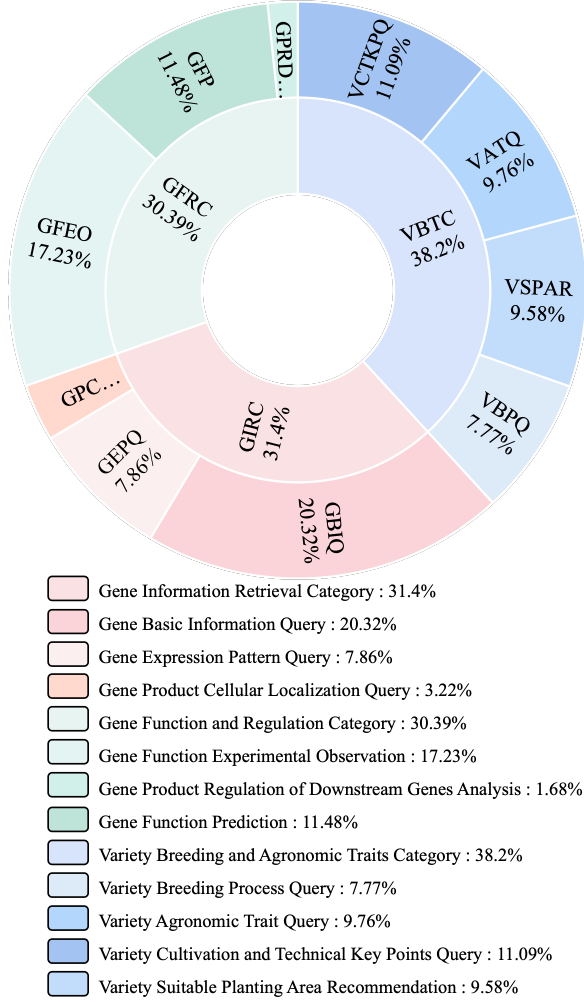


Figure 2: Benchmark Taxonomy Distribution. Three core breeding steps are further divided into ten expert-curated subcategories within SeedBench, which comprises a total of 2,264 questions. The percentages shown in the diagram represent the proportion of questions in each category relative to 2,264.

This methodology ensures a relevant, diverse, and scientifically rigorous evaluation of LLMs.

### 3.3.1 Data Collection

SeedBench is built on a comprehensive breeding knowledge base. We extracted 308,727 breeding-related publications in English and Chinese from publicly available sources to minimize language bias<sup>2</sup>. These papers were converted to Markdown using MinerU (Wang et al., 2024a) for consistency, primarily consisting of open-access academic works for credibility and reusability.

Data cleaning involved three steps: (1) heuristic filtering to remove noise, including corrupted

<sup>2</sup>See Appendix B.2 for detailed language composition. We also discuss the potential impact of linguistic differences on the model’s performance.

or irrelevant data; (2) deduplication using a local-sensitive hashing method; and (3) scoring corpus segments with the CCI3-HQ-Classifer (Wang et al., 2024b) to eliminate low-quality fragments. And we used the IndustryCorpus2 Classifier<sup>3</sup> to exclude content unrelated to seed breeding. These steps filtered out 86% of low-quality or duplicate data, yielding a corpus of 1.1 billion tokens.

From this corpus, domain experts selected 279 high-quality text segments, sourced from 113 documents, each averaging 300 words. These segments cover 10 predefined subcategories, ensuring both depth and breadth. Each segment is highly relevant to a specific subcategory, containing multiple knowledge points essential for model evaluation. To aid in automatic question generation, domain experts manually designed 293 reference questions. For illustration purpose, we consistently use rice to demonstrate the benchmark construction, empirical evaluation, and case studies across the main text and Appendix, given its global importance and representativeness in seed science.<sup>4</sup> The data collection and curation for maize, soy bean, wheat, and many other plants follow the same procedure.

### 3.3.2 Automatic Question Generation

Type ID	Question Type	Metric	Count (n)
<b>Q&amp;A</b>			
QA-1	Multiple Choice	Accuracy	199
QA-2	Multiple Answer	Macro-F1	186
QA-3	Fill-in-the-Blank	ROUGE-L	223
QA-4	Generation	ROUGE-L	241
<b>Summarization</b>			
SUM-1	Simple Summarization	ROUGE-L	224
SUM-2	Key Information Extraction	ROUGE-L	224
<b>Reading Comprehension</b>			
RC-1	Multiple Choice	Accuracy	112
RC-2	Multiple Answer	Macro-F1	107
RC-3	Fill-in-the-Blank	ROUGE-L	220
RC-4	Generation	ROUGE-L	239
RC-5	Subcategory Classification	Accuracy	278

Table 1: Benchmark Task Types. Each high-quality text segment systematically incorporates these 11 distinct task types to ensure diversity. ‘Count (n)’ indicates the number of questions for this particular task type after quality validation, from the total of 2,264 questions in SeedBench. The complete distribution statistics are provided in Appendix C.2.

The question generation phase begins with

<sup>3</sup>[https://huggingface.co/BAAI/IndustryCorpus2\\_Classifier](https://huggingface.co/BAAI/IndustryCorpus2_Classifier)

<sup>4</sup>Rice is the most widely consumed crop in the world, which feeds over 3.5 billion people in the world according to Food and Agriculture Organization of the United Nations and World Bank.

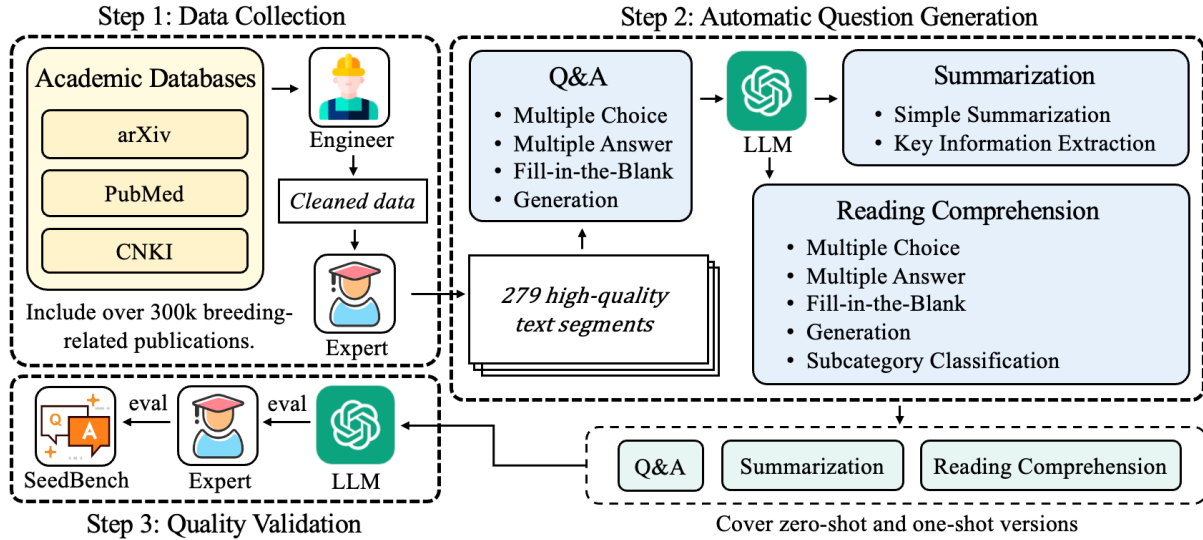


Figure 3: Benchmark Construction Pipeline. We developed SeedBench by extracting 308,727 breeding-related papers from English and Chinese sources and converting them into a unified Markdown format. The data underwent rigorous cleaning, ultimately yielding a 1.1-billion-token corpus. From this, experts curated 279 high-quality text segments, spanning 10 breeding subcategories, for generating LLM-based Q&A tasks. Validation included both automated and expert reviews, removing low-quality entries and ensuring relevance. SeedBench offers 2,264 refined questions across 11 task types, enabling fine-grained evaluation of LLMs in seed breeding.

expert-designed example questions across four Q&A types: multiple choice (QA-1), multiple answer (QA-2), fill-in-the-blank (QA-3), and generation (QA-4). This foundation extends to summarization (SUM-1, SUM-2) and reading comprehension (RC-1 to RC-5) tasks, detailed in Table 1.

Utilizing GPT-4’s natural language processing capabilities, unstructured text is converted into structured knowledge dictionaries through key information extraction (Appendix F.6), capturing genetic traits, phenotypic features, and practical applications. Based on these dictionaries, GPT-4 generates four Q&A task types:

- Multiple Choice Questions
- Multiple Answer Questions
- Fill-in-the-Blank Questions
- Generation Questions

Then, we generated 450 summarization questions (SUM-1, SUM-2), with two per text segment, to evaluate the model’s ability to summarize breeding literature without requiring prior domain knowledge. For reading comprehension tasks, the original text segments were provided as reference documents in the augmented context. Using GPT-4, we rephrased the questions to require retrieval, analysis, reasoning, and answering based on the given document, assessing the model’s capability in long-context breeding problems.

Additionally, we formulated 279 classification

questions, where the answer corresponds to the category of each text segment, testing the model’s ability to distinguish breeding areas. All questions are available in both zero-shot and one-shot settings. SeedBench thus serves as a multidimensional benchmark, covering 10 thematic subcategories and 11 task types for fine-grained breeding assessment. All prompt templates that we used are provided in Appendix F.

### 3.3.3 Quality Validation

Since SeedBench relies on GPT-4 for annotation, we implemented a two-stage validation process to ensure accuracy and reliability: automated machine screening and manual expert review.

**Automated Machine Screening.** GPT-4 first assessed coherence, logical consistency, and task adherence for each question, filtering out those with errors or contradictions. About 0.01% of the questions were excluded at this stage.

**Manual Expert Review.** Domain experts reviewed the remaining questions for relevance and alignment with expert perspectives, removing irrelevant or weakly contextualized ones (*e.g.*, “Is rice planted in Beijing or Shanghai?”) and eliminating about 20% of the initial set.

After validation, 2,264 high-quality questions were retained. Detailed case studies and question distribution are provided in Appendix C.2.

### 3.4 Evaluation

The evaluation framework consists of two stages: response standardization and task-specific scoring.

- Multiple Choice Tasks: Accuracy.
- Multiple Answer Tasks: Macro-F1.
- Fill-in-the-Blank: ROUGE-L F1 for segment-level comparison between model predictions and reference answers, averaged across segments.
- Generation Tasks: ROUGE-L F1 for sentence-level comparison of full generated responses against references. BERTScore included as an additional evaluation metric (Appendix H.7).

Formal mathematical definitions and post-processing methods are detailed in Appendix D.2. To account for prompt sensitivity in LLMs, Appendix G compares different prompt templates.

## 4 Experiment

### 4.1 Experimental Setup

**Models.** We compared the performance of 26 LLMs on the SeedBench, including 7 proprietary models, such as the GPT series (Achiam et al., 2023), Gemini series (Team et al., 2024), Claude-3.5-Sonnet<sup>5</sup>, and GLM-4-Plus<sup>6</sup>, as well as 16 open-source models, including the Qwen series (Yang et al., 2024a), DeepSeek-V3 (Liu et al., 2024), Llama series (Dubey et al., 2024), InternLM series (Cai et al., 2024), GLM-4 (GLM et al., 2024), and Mistral (Jiang et al., 2023). Furthermore, we evaluated 3 domain-specific models, represented by the PLLaMa series (Yang et al., 2024b) and Ak-sara<sup>7</sup>. This comprehensive comparison not only highlights the relative strengths and limitations of each model group but also provides key insights for future research and application development.

**Implementation Details.** We evaluated the performance of all models in both zero-shot and one-shot settings. In zero-shot inference, the model input includes only the task instructions and the query. In one-shot inference, the model input consists of the task instructions, an example query with its answer, followed by the actual query. The experiments were conducted using the OpenCompass<sup>8</sup> framework. For proprietary LLMs, we performed inference through their APIs. The evaluation of each model took approximately 1 hour. For open-

source LLMs, the evaluation was conducted on 8 NVIDIA A100 40GB GPUs, with an average completion time of 0.5 hours. The inference hyperparameters are detailed in Appendix D.1.

### 4.2 Performance Evaluation

Here we assess overall performance across breeding subcategories, as detailed in Table 2<sup>9</sup>. Among proprietary LLMs, GPT-4 achieves the highest average score on SeedBench (62.06), followed by GLM-4-plus (59.61). In contrast, open-source models show a different ranking: DeepSeek-V3 leads with an average score of 63.3, outperforming Qwen2.5-14B, which scores 54.3. Notably, DeepSeek-V3, despite being a recently released model with 671B parameters, surpasses GPT-4 on SeedBench. On the other hand, the three domain-specific LLMs perform relatively poorly, likely due to their limited conversational and instruction following capabilities. Interestingly, OpenAI o1, despite demonstrating strong reasoning abilities in mathematics and coding, scores lower than GPT-4 on SeedBench. This suggests that its reasoning strategy does not transfer effectively to breeding-related tasks. A similar trend is observed in Gemini-2.0-Flash Thinking and QwQ-32B, both of which exhibit explicit reasoning steps in their responses yet achieve only 34.24 and 33.55, respectively. A more detailed comparative analysis is provided in Appendix H, where we examine performance variations within the same model series, evaluate models of identical sizes, and assess the impact of task difficulty. And Appendix E provides a holistic error taxonomy and analysis of models' failures, summarizing eight primary causes of errors, such as Gene Name Confusion. These additional analyses collectively offer deeper insights into the factors influencing model performance on SeedBench.

### 4.3 Empirical Analysis

#### 4.3.1 Analysis on Reasoning Ability

Table 2 shows that LLMs with explicit "reasoning mode" do not consistently outperform those without specialized chain-of-thought mechanisms (RQ1). For instance, models designed for multi-step reasoning, such as Gemini-2.0-Flash (34.24) and QwQ-32B (33.55), achieve lower average scores than general-purpose LLMs such as GPT-4 (62.06) and the top-performing open-source model, DeepSeek-V3-671B (63.30). This discrepancy is

<sup>5</sup><https://www.anthropic.com/>

<sup>6</sup><https://open.bigmodel.cn/dev/api/normal-model/glm-4>

<sup>7</sup>[https://huggingface.co/cropinailab/aksara\\_v1](https://huggingface.co/cropinailab/aksara_v1)

<sup>8</sup><https://opencompass.org.cn/home>

<sup>9</sup>See Tables 10 11 12 in Appendix H.8 for complete results.

Models	Breeding Subcategories										Average
	C1	C2	C3	C4	C5	C6	C7	C8	C9	C10	
Proprietary LLMs											
Claude-3.5-Sonnet	48.77	57.72	66.02	57.54	47.82	49.36	57.47	60.11	58.06	58.89	55.45
Gemini-1.5-Pro	47.00	59.55	62.42	59.56	43.11	49.55	53.41	56.18	52.51	53.71	53.58
Gemini-2.0-Flash	33.67	27.37	53.04	32.07	25.87	44.41	33.57	36.77	31.78	31.70	34.24
GLM-4-Plus	52.72	59.62	70.62	60.11	50.60	56.75	65.02	64.17	61.70	62.90	59.61
GPT-4	59.59	60.55	76.32	61.16	56.34	59.35	63.67	64.74	60.65	67.66	62.06
GPT-4o mini	54.24	56.64	72.11	59.28	53.00	57.88	58.38	61.75	57.50	62.38	58.40
OpenAI o1-mini	49.16	55.58	59.37	54.77	44.43	50.73	54.57	55.36	54.91	54.19	53.25
Open-Source LLMs											
DeepSeek-V3-671B	56.03	62.42	74.81	63.17	55.23	58.84	68.23	69.04	66.46	68.48	63.30
GLM-4-Chat-9B	23.28	21.31	39.97	26.13	16.20	34.15	26.63	29.60	25.60	26.68	26.55
InternLM2-7B	27.55	21.14	39.64	28.57	15.16	36.12	28.74	30.80	27.32	29.22	28.71
InternLM2.5-7B	51.71	55.75	67.88	50.48	44.14	56.73	51.28	54.91	52.46	56.24	53.51
Llama3.1-8B	43.89	31.21	42.53	40.68	38.47	43.80	42.87	51.62	41.88	40.91	42.23
Llama3.1-70B	48.72	55.41	64.77	53.67	46.73	54.08	56.94	57.72	55.31	57.56	54.30
Llama3.3-70B	45.32	47.15	60.62	49.76	40.90	54.30	52.79	54.61	49.98	55.05	50.53
Mistral-v0.3-7B	42.61	38.28	57.02	40.41	29.97	44.22	36.31	43.98	39.92	43.51	41.59
Qwen2-0.5B	32.84	25.15	40.19	28.20	27.62	37.22	33.81	33.63	28.25	31.67	31.44
Qwen2-7B	44.21	40.41	63.00	47.36	35.37	52.30	45.61	48.73	44.88	46.89	46.51
Qwen2-57B	53.67	49.81	74.30	58.38	39.34	54.71	63.89	59.57	59.22	60.08	57.20
Qwen2-72B	51.16	58.10	74.07	59.72	51.58	57.76	58.85	61.63	56.69	59.11	57.62
Qwen2.5-7B	45.16	39.50	66.01	44.61	35.72	50.00	53.60	53.31	53.06	51.05	48.45
Qwen2.5-14B	50.91	50.73	68.62	52.15	47.14	54.54	57.02	62.05	54.37	54.15	54.21
Qwen2.5-72B	46.86	47.41	70.99	51.89	46.17	57.60	55.35	56.31	53.05	54.75	52.63
QwQ-32B	32.24	21.06	47.11	29.14	28.56	39.68	38.17	39.56	34.70	34.52	33.55
Domain Specific LLMs											
Aksara-v1-7B	36.72	36.69	48.32	35.41	24.26	36.83	31.17	34.64	31.15	34.14	35.04
PLLaMa-7B	17.85	13.69	17.99	16.81	11.66	21.67	14.34	17.36	12.39	16.11	16.46
PLLaMa-13B	15.10	14.18	28.41	18.83	13.96	23.28	18.53	17.37	14.15	18.51	17.57

Table 2: Evaluation of 26 LLMs on SeedBench. Performance (averaged across both zero-shot and one-shot configurations) is stratified by breeding subcategories, with open-source/domain-specific models evaluated through 3 repeated trials (mean scores reported). The scores represent averages across three different metrics for 11 task types. The columns delineate ten subcategories in breeding: (C1) Gene Basic Information Query, (C2) Gene Expression Pattern Query, (C3) Gene Product Cellular Localization Query, (C4) Gene Function Experimental Observation, (C5) Gene Product Regulation of Downstream Genes Analysis, (C6) Gene Function Prediction, (C7) Variety Breeding Process Query, (C8) Variety Agronomic Trait Query, (C9) Variety Cultivation and Technical Key Points Query, (C10) Variety Suitable Planting Area Recommendation. Top-3 performers per column are highlighted in red. Extended results (including **standard deviations**, **separate** zero-shot/one-shot scores, and task-type breakdowns) are provided in Appendix H.8.

evident in subcategories that primarily involve fact retrieval or straightforward inference, such as Gene Basic Information Query (C1) and Variety Cultivation and Technical Key Points Query (C9).

Furthermore, the results suggest potential drawbacks of verbose chain-of-thought reasoning in tasks that require single-step inferences. Lengthy reasoning chains can introduce unnecessary content, potentially reducing performance on precision-based metrics such as ROUGE. This effect is particularly evident in tasks like Gene Product Regulation of Downstream Genes Analysis (C5), where the best-performing models (*e.g.*, GPT-4 and DeepSeek-V3-671B) maintain conciseness while effectively capturing key information. Overall, these findings indicate that while reasoning-centric prompts may be advantageous in complex multi-step tasks (*e.g.*, coding or mathematical problem-

solving), seed breeding queries rely on direct knowledge retrieval or limited inference. Future research could focus on adjusting the reasoning strategy to match task complexity, dynamically adjusting reasoning chains based on task difficulty, rather than using complex reasoning for all tasks (Chen et al., 2024b).

#### 4.3.2 Impact of Domain-Specific Fine-Tuning

Contrary to expectations, domain-specific fine-tuned models (*e.g.*, Aksara-v1-7B, PLLaMa-7B, PLLaMa-13B) perform worse than general-purpose models on SeedBench. As shown in Table 2, these specialized models exhibit significantly lower overall scores (*e.g.*, 35.04 for Aksara-v1-7B and 17.57 for PLLaMa-13B). In comparison, mid-tier open-source LLMs such as Llama3.1-70B (54.30) and Qwen2.5-14B (54.21) outper-

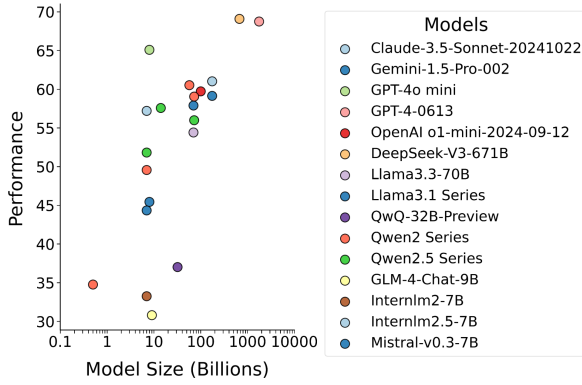


Figure 4: Performance vs. Model Size. We empirically validate scaling laws in seed breeding tasks, showing a logarithmic correlation between model size and average scores. The optimal model size for breeding tasks lies between 7B and 14B, balancing performance and computational efficiency.

form them, while top-performing general-purpose models like DeepSeek-V3-671B (63.30) and GPT-4 (62.06) achieve even higher scores. Notably, Aksara-v1-7B’s best subcategory score (48.32 in C3: Gene Product Cellular Localization Query) remains below many open-source models’ performance in similar tasks. We hypothesize that a key factor behind this underperformance is the deterioration of conversational and instruction-following capabilities, after fine-tuning on domain-specific data. In one-shot evaluations, PLLaMas frequently misinterpret prompts, treating in-context examples as direct queries. This shows that domain specialization may weaken general instruction adherence, which is essential for tasks requiring customized outputs and complex user interactions. Additionally, post-training on narrowly defined tasks or using data not validated by experts may further diminish models’ performance. Thus, we suggest that maintaining general-purpose abilities during fine-tuning, expanding the breadth of training data, and incorporating expert validation may effectively improve the performance of domain-specific fine-tuned models. Further investigations into the fine-tuning strategy and the scope and quality of the training corpus are required to fully answer **RQ2**.

#### 4.3.3 Impact of Model Size

As shown in Figure 4, models with 7B–14B parameters achieve the best trade-off between performance and computational efficiency (**RQ3**). Models in this range, such as Qwen2.5-14B and InternLM-2.5-7B, perform robustly with manageable resource requirements. In contrast, smaller

models (e.g., Qwen2-0.5B) underperform, and much larger models (e.g., Qwen2.5-72B) show diminishing returns. Additionally, the lack of performance improvement across scaled Qwen models (14B to 72B) suggests a distribution mismatch between our benchmark data and Qwen’s enlarged training corpus. We conjecture that the quality of training corpus is more important than model size for domain-specific tasks. In addition, we analyze performance differences between models of varying sizes, series, and subcategory difficulty in Appendices H.1 to H.6.

## 5 Discussion

The discrepancy between the seed scientists’ expectations and the reality of LLMs is evident. From the perspective of breeding experts, the immediate applicability of domain-specific fine-tuned models to seed science remains constrained. This may arise from factors such as training on narrowly defined tasks, reliance on data lacking expert validation, or catastrophic forgetting of general capabilities during fine-tuning. Conversely, while current general LLMs have shown good potential in text understanding and basic reasoning, they still fall short of meeting the deep and specialized requirements of actual breeding work. Several key gaps between seed scientists’ expectations and the capabilities of general LLMs are highlighted below.

**Domain Depth.** While LLMs cover a broad range of agricultural topics, their knowledge depth is insufficient for specialized breeding tasks. Complex issues like molecular breeding or trait introgression often require high-quality domain knowledge repositories, an area where current models are lacking. One potential solution is to integrate structured knowledge graphs mapping relationships across phenotypic, genomic, and environmental data.

**Multimodal Integration.** Breeding in practice heavily relies on sensory evaluations (such as detecting grain morphology, texture, and odor) and environmental data (such as climate and soil conditions). Current LLMs, primarily based on text input, cannot effectively integrate images, sensor readings, and field observations, limiting their performance in multimodal decision-making.

**Explainability and Risk Management.** The breeding decision cycle is costly and prolonged. Without transparent reasoning and risk evaluation mechanisms, misleading outputs could result in substantial losses. Producers must be able to ver-

ify and trace model conclusions to ensure safe and controllable implementation, *e.g.*, via safe RL.

## 6 Conclusion

In this study, we propose SeedBench, the first multi-task LLM benchmark tailored for seed science. The contributions of SeedBench to AI for Science are twofold. First, it demonstrates the complete benchmark construction process for seed science, a knowledge-intensive field. The knowledge behind can be transferred to other science domains, such as life science or physical science, to build a comprehensive and reliable scientific benchmark. Second, it evaluates the capabilities of LLMs on addressing seed breeding tasks. The results not only provide insightful empirical findings on the tasks on interest, but also pose future research directions on designing scientific LLM. By bridging the gap between the power of LLMs and the real-world scientific problem in seed science, we aim to make a foundational effort for successful implementation of LLMs for seed design in the future.

## Limitations

Given the highly specialized nature of breeding research, we primarily use the peer-reviewed scientific literature to ensure the accuracy and credibility of the data. Online sources often lack systematic review and professional validation. In the future, we will explore more reliable online databases and expert knowledge bases to further diversify our data sources.

There remains a gap between the expectations of the researchers and the actual capabilities of LLMs. Closing this gap calls for more domain-focused and in-depth professional datasets, as well as additional breeding-specific knowledge during model training. Another key direction is to incorporate sequential decision making and iterative learning into the models, so that they can adapt to planting cycles and experimental feedback. Additionally, developing LLMs that support multi-modal inputs—such as phenotypic, genomic, and environmental data—would be crucial for complex breeding scenarios. Finally, improving model interpretability and safety mechanisms will be essential for building trust when the model provides breeding recommendations.

## Ethical Considerations

This research adheres to the ethical principles outlined in the *ACL Code of Ethics*. We have taken steps to ensure that our work does not cause harm, particularly in the context of seed science, where the implications of our research could impact agricultural practices and food security.

We have carefully considered the potential risks of our approach, especially with respect to the use of LLMs in agricultural research. One of the main ethical concerns we addressed is the potential for bias in LLMs, which could affect seed breeding decisions or lead to misinterpretations in agricultural data. To mitigate this, we evaluated multiple LLMs, to better understand their limitations and potential biases, and have explicitly highlighted areas where further improvement is needed.

In addition, we ensured that all data used in our experiments were sourced responsibly, with due regard for privacy and intellectual property. The artifacts we used are governed under the CC-BY 4.0, Apache 2.0 and MIT licenses, which support open and ethical use of such resources. The creation of the SeedBench benchmark involved collaboration with domain experts to ensure that it accurately represents the complexities of seed breeding, while minimizing any unintended consequences. Moreover, we have taken care to ensure that the use of these artifacts aligns with the intended purpose of advancing research in seed breeding and AI applications.

Finally, we have considered the broader social impact of our work, recognizing that our research could influence farming practices, genetic resource management, and food security. We have outlined the necessary steps to avoid harm and ensure responsible application of our findings.

## Acknowledgments

This work was supported by Shanghai Artificial Intelligence Laboratory, the Yazhouwan National Laboratory Project (grant no. 2310CF01), and the Hainan Yazhou Bay Seed Laboratory Project (grant no. B21HJ0001). The authors thank anonymous reviewers for their constructive comments. The authors also thank Yazhouwan National Laboratory for data collection and validation of human experts. The authors also thank Songyang Zhang from Shanghai Artificial Intelligence Laboratory for academic discussion.

## References

- Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altschmidt, Sam Altman, Shyamal Anadkat, et al. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.
- Muhammad Arbab Arshad, Talukder Zaki Jubery, Tirtho Roy, Rim Nassiri, Asheesh K Singh, Arti Singh, Chinmay Hegde, Baskar Ganapathysubramanian, Aditya Balu, Adarsh Krishnamurthy, et al. 2024. Ageval: A benchmark for zero-shot and few-shot plant stress phenotyping with multimodal llms. *arXiv preprint arXiv:2407.19617*.
- Zheng Cai, Maosong Cao, Haojiong Chen, Kai Chen, Keyu Chen, Xin Chen, Xun Chen, Zehui Chen, Zhi Chen, Pei Chu, et al. 2024. Internlm2 technical report. *arXiv preprint arXiv:2403.17297*.
- Pengcheng Chen, Jin Ye, Guoan Wang, Yanjun Li, Zhongying Deng, Wei Li, Tianbin Li, Haodong Duan, Ziyang Huang, Yanzhou Su, et al. 2024a. Gmai-mmbench: A comprehensive multimodal evaluation benchmark towards general medical ai. *arXiv preprint arXiv:2408.03361*.
- Xingyu Chen, Jiahao Xu, Tian Liang, Zhiwei He, Jianhui Pang, Dian Yu, Linfeng Song, Qiuzhi Liu, Mengfei Zhou, Zhuosheng Zhang, et al. 2024b. Do not think that much for  $2+3=?$  on the overthinking of o1-like llms. *arXiv preprint arXiv:2412.21187*.
- Lawrence O Copeland and Miller F McDonald. 2012. *Principles of seed science and technology*. Springer Science & Business Media.
- Djavan De Clercq, Elias Nehring, Harry Mayne, and Adam Mahdi. 2024. Large language models can help boost food production, but be mindful of their risks. *Frontiers in Artificial Intelligence*, 7:1326153.
- Jacob Devlin. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, et al. 2024. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*.
- Lucy M. Egan, Rainer W. Hofmann, Warwick N. Stiller, and Valerio Hoyos-Villegas. 2024. Cultivating success: Bridging the gaps in plant breeding training in australia, canada, and new zealand. *Crop Science*, 64(4).
- Zhiwei Fei, Xiaoyu Shen, Dawei Zhu, Fengzhe Zhou, Zhuo Han, Alan Huang, Songyang Zhang, Kai Chen, Zhixin Yin, Zongwen Shen, Jidong Ge, and Vincent Ng. 2024. LawBench: Benchmarking legal knowledge of large language models. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 7933–7962, Miami, Florida, USA.
- Wensheng Gan, Zhenlian Qi, Jiayang Wu, and Jerry Chun-Wei Lin. 2023. Large language models in education: Vision and opportunities. In *2023 IEEE international conference on big data (BigData)*, pages 4776–4785. IEEE.
- Team GLM, Aohan Zeng, Bin Xu, Bowen Wang, Chenhui Zhang, Da Yin, Dan Zhang, Diego Rojas, Guanyu Feng, Hanlin Zhao, et al. 2024. Chatglm: A family of large language models from glm-130b to glm-4 all tools. *arXiv preprint arXiv:2406.12793*.
- Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shirong Ma, Peiyi Wang, Xiao Bi, et al. 2025. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. *arXiv preprint arXiv:2501.12948*.
- K Iversen, HWJ Cheng, K Helgason, and M LaFleur. 2021. Frontier technologies for smallholder farmers: Addressing information asymmetries and deficiencies. *Frontier Technology Issues*.
- Albert Q Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, et al. 2023. Mistral 7b. *arXiv preprint arXiv:2310.06825*.
- Josué Kpodo, Parisa Kordjamshidi, and A Pouyan Nejadhashemi. 2024. Agxqa: A benchmark for advanced agricultural extension question answering. *Computers and Electronics in Agriculture*, 225:109349.
- Matheus Thomas Kuska, Mirwaes Wahabzada, and Stefan Paulus. 2024. Ai for crop production—where can large language models (llms) provide substantial value? *Computers and Electronics in Agriculture*, 221:108924.
- Dongyang Li, Ruixue Ding, Qiang Zhang, Zheng Li, Boli Chen, Pengjun Xie, Yao Xu, Xin Li, Ning Guo, Fei Huang, et al. 2023a. Geoglu: A geographic language understanding evaluation benchmark. *arXiv preprint arXiv:2305.06545*.
- Yinheng Li, Shaofei Wang, Han Ding, and Hang Chen. 2023b. Large language models in finance: A survey. In *Proceedings of the fourth ACM international conference on AI in finance*, pages 374–382.
- Stephanie Lin, Jacob Hilton, and Owain Evans. 2022. TruthfulQA: Measuring how models mimic human falsehoods. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 3214–3252, Dublin, Ireland.
- Aixin Liu, Bei Feng, Bing Xue, Bingxuan Wang, Bochao Wu, Chengda Lu, Chenggang Zhao, Chengqi Deng, Chenyu Zhang, Chong Ruan, et al. 2024. Deepseek-v3 technical report. *arXiv preprint arXiv:2412.19437*.

- Ummara Mumtaz, Awais Ahmed, and Summaya Mumtaz. 2024. Llm-healthcare: Current applications and challenges of large language models in various medical specialties. *Artificial Intelligence in Health*, 1(2):16–28.
- Alec Radford. 2018. Improving language understanding by generative pre-training.
- Gemini Team, Petko Georgiev, Ving Ian Lei, Ryan Burnell, Libin Bai, Anmol Gulati, Garrett Tanzer, Damien Vincent, Zhufeng Pan, Shibo Wang, et al. 2024. Gemini 1.5: Unlocking multimodal understanding across millions of tokens of context. *arXiv preprint arXiv:2403.05530*.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. *Advances in Neural Information Processing Systems*.
- Bin Wang, Chao Xu, Xiaomeng Zhao, Linke Ouyang, Fan Wu, Zhiyuan Zhao, Rui Xu, Kaiwen Liu, Yuan Qu, Fukai Shang, et al. 2024a. Mineru: An open-source solution for precise document content extraction. *arXiv preprint arXiv:2409.18839*.
- Liangdong Wang, Bo-Wen Zhang, Chengwei Wu, Hanyu Zhao, Xiaofeng Shi, Shuhao Gu, Jijie Li, Quanyue Ma, TengFei Pan, and Guang Liu. 2024b. Cci3.0-hq: a large-scale chinese dataset of high quality designed for pre-training large language models. *arXiv preprint arXiv:2410.18505*.
- Zeyu Wang. 2024. Causalbench: A comprehensive benchmark for evaluating causal reasoning capabilities of large language models. In *Proceedings of the 10th SIGHAN Workshop on Chinese Language Processing*, pages 143–151.
- Qianqian Xie, Weiguang Han, Zhengyu Chen, Ruoyu Xiang, Xiao Zhang, Yueru He, Mengxi Xiao, Dong Li, Yongfu Dai, Duanyu Feng, et al. 2024. The finben: An holistic financial benchmark for large language models. *arXiv preprint arXiv:2402.12659*.
- An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan Li, Dayiheng Liu, Fei Huang, Haoran Wei, et al. 2024a. Qwen2.5 technical report. *arXiv preprint arXiv:2412.15115*.
- Xianjun Yang, Junfeng Gao, Wenxin Xue, and Erik Alexandersson. 2024b. Pllama: An open-source large language model for plant science. *arXiv preprint arXiv:2401.01600*.
- Hang Zhang, Jiawei Sun, Renqi Chen, Wei Liu, Zhonghang Yuan, Xinzhe Zheng, Zhefan Wang, Zhiyuan Yang, Hang Yan, Han-Sen Zhong, et al. 2024. Empowering and assessing the utility of large language models in crop science. In *The Thirty-eight Conference on Neural Information Processing Systems Datasets and Benchmarks Track*, pages 52670–52722.
- Jeffrey Zhou, Tianjian Lu, Swaroop Mishra, Siddhartha Brahma, Sujoy Basu, Yi Luan, Denny Zhou, and Le Hou. 2023. Instruction-following evaluation for large language models. *arXiv preprint arXiv:2311.07911*.

# Appendix

## Table of Contents

---

<b>A Background Definitions</b>	<b>13</b>
A.1 Definition of Breeding Competencies . . . . .	13
A.2 Example Definitions of Breeding Terminology . . . . .	14
<b>B Source Data Collection</b>	<b>15</b>
B.1 Construction and Annotation of High-Quality Text Segments . . . . .	15
B.2 Language Composition and Impact . . . . .	16
<b>C Quality Verification</b>	<b>16</b>
C.1 Manual Quality Verification . . . . .	17
C.2 Descriptive Statistics . . . . .	18
<b>D Additional Experimental Setup</b>	<b>19</b>
D.1 Model Hyperparameters . . . . .	19
D.2 Evaluation Metrics . . . . .	19
<b>E Error Analysis of Large Language Model Outputs</b>	<b>21</b>
E.1 Errors in the First Task Type . . . . .	21
E.2 Errors in the Second Task Type . . . . .	27
E.3 Errors in the Third Task Type . . . . .	31
<b>F Prompt Templates</b>	<b>34</b>
F.1 Multiple-Choice Question Generation (Task QA-1) . . . . .	34
F.2 Multiple-Answer Question Generation (Task QA-2) . . . . .	35
F.3 Fill-in-the-Blank Question Generation (Task QA-3) . . . . .	36
F.4 Text-based Q&A Generation (Task QA-4) . . . . .	37
F.5 Naive Summarization (Task SUM-1) . . . . .	38
F.6 Key Information Extraction (Task SUM-2) . . . . .	39
F.7 Question Rewriting (Tasks RC-1, RC-2, RC-3, RC-4) . . . . .	40
<b>G Robustness Evaluation of Prompts</b>	<b>41</b>
<b>H Detailed Performance Comparison and Analysis</b>	<b>49</b>
H.1 Model Comparison on Different LLMs . . . . .	49
H.2 Comparison Between Different Model Sizes of the Same Model Series . . . . .	49
H.3 Top Performers in Each Subcategory . . . . .	49
H.4 Comparison of Models with Same Parameter Scale . . . . .	49
H.5 Task-Specific Performance . . . . .	50
H.6 Conclusion . . . . .	50
H.7 Assessment of the Generative Tasks using BERTScore . . . . .	51
H.8 Additional Quantitative Results . . . . .	52

---

## A Background Definitions

### A.1 Definition of Breeding Competencies

We define in Table 3 the core breeding competencies required to address complex, knowledge-intensive breeding tasks. The breeding process is divided into three steps: the acquisition and analysis of genetic information, the exploration of gene function and regulatory mechanisms, and variety selection alongside agronomic trait optimization. Each stage corresponds to specific knowledge and skill requirements.

Stage	Corresponding Task	Competency	Description
<b>Acquisition and Analysis of Genetic Information</b>	Basic Gene Information Retrieval	Ability to Retrieve Basic Gene Information	Is able to retrieve and integrate a gene's basic information (e.g., sequence, functional annotation, chromosomal location) based on user inputs such as gene names or identifiers.
	Gene Expression Pattern Retrieval	Ability to Analyze Gene Expression Patterns	Possesses the capacity to explore and interpret gene expression data under different tissues, developmental stages, or environmental conditions.
	Subcellular Localization of Gene Products	Ability to Determine Gene Product Localization and Characteristics	Uses existing protein or molecular data to infer the specific location of a gene product within the cell (e.g., nucleus, cytoplasm, cell membrane) and, by applying biological knowledge, deduces its potential functions (e.g., transcriptional regulation, signal transduction, material transport).
<b>Exploration of Gene Function and Regulatory Mechanisms</b>	Observations from Gene Function Experiments	Ability to Interpret Experimental Results on Gene Function	Accurately describes how a gene influences plant traits under specific conditions.
	Analysis of Gene Product Regulation on Downstream Genes	Ability to Analyze Gene Regulatory Networks and Downstream Genes	Relies on existing research or deductive reasoning to evaluate the regulatory effects of a target gene product on downstream genes or related pathways.
	Gene Function Prediction	Ability to Predict and Validate Unknown Gene Functions	Draws on known gene sequences, expression characteristics, and analogies with characterized genes to predict functions of genes that have not yet been extensively studied.
<b>Variety Selection and Agronomic Trait Optimization</b>	Retrieval of Variety Breeding Processes	Ability to Oversee Variety Breeding Processes	Collects and summarizes the breeding history, methods, and improvement goals for a specific variety.
	Querying Agronomic Traits of Varieties	Ability to Screen and Evaluate Agronomic Traits	Compares the agronomic traits (e.g., disease resistance, yield, drought tolerance) of a target variety or multiple varieties to assess their value in practical production, selecting the most suitable breeding materials or gene combinations.
	Querying Key Points of Variety Cultivation and Techniques	Ability to Plan Variety Cultivation Management and Key Technical Points	Gathers and analyzes crucial information on cultivation management (e.g., fertilization strategies, irrigation schedules, pest control) to optimize field trials or practical planting outcomes.
	Recommending Suitable Planting Regions for Varieties	Ability to Evaluate and Recommend Planting Regions	Takes into account environmental factors and agronomic traits to propose suitable planting regions for a given variety, while evaluating its potential for wider application.

Table 3: Core Breeding Competencies

## A.2 Example Definitions of Breeding Terminology

We define in Table 4 the specialized breeding terms mentioned in this paper, along with explanations to provide a unified reference for subsequent experimental design, task classification, and result interpretation.

Terminology	Definition
<b>Gene Information</b>	Refers to the fundamental data of a given gene, including its sequence, structure, functional descriptions, and expression patterns. This information is typically obtained from genomic databases or sequencing results, serving as the starting point for subsequent gene function research and variety improvement.
<b>Gene Expression Pattern</b>	Describes a gene’s expression levels and dynamic changes across various tissues, developmental stages, or environmental conditions. By analyzing this pattern, one can determine the gene’s potential impact on target traits (e.g., drought tolerance).
<b>Gene Regulatory Network</b>	A molecular regulatory system formed by interactions among gene products (e.g., transcription factors, enzymes). It determines how crops respond to environmental stimuli or developmental requirements and serves as an essential foundation for precision breeding analyses.
<b>Hybridization &amp; Backcrossing</b>	Involves breeding new offspring by crossing or pollinating different parental lines. Backcrossing is carried out with a desirable parent line over multiple generations to stabilize or strengthen specific traits. This approach is traditional yet effective, often combined with modern molecular marker techniques.
<b>Agronomic Trait</b>	Refers to key field performance characteristics of crops, such as plant height, number of tillers, yield, disease resistance, and maturity. These traits are crucial indicators for evaluating the economic value and potential for broader adoption of new varieties.
<b>Marker-Assisted Breeding</b>	Uses molecular markers (e.g., SNP, SSR) for genotypic screening to accelerate the breeding process and improve selection accuracy. This method is frequently employed during hybridization and backcrossing to rapidly pinpoint genes associated with targeted traits.
<b>LLM-Aided Breeding</b>	Leverages large language models (e.g., GPT) to analyze and reason over literature, databases, and experimental data, supporting gene information analysis, gene function prediction, and variety selection decisions. It is expected to shorten breeding cycles and enhance breeding efficiency.

Table 4: Breeding Terminology Definitions

## B Source Data Collection

### B.1 Construction and Annotation of High-Quality Text Segments

In Figure 5 below, Content is a snippet extracted by agricultural experts from selected scientific publications. Example Question is an illustrative question provided by the experts, demonstrating from which perspective or angle one may inquire about the snippet. Classification indicates the subfield of agriculture to which the snippet belongs, and Reference identifies the scientific publication from which the snippet originates.

<p><b>Content:</b> A pathway composed of five subunits of the heterotrimeric G proteins that regulate grain length in rice. The G<math>\beta</math> protein is essential for plant survival and growth. G<math>\alpha</math> provides a foundation for grain size expansion. Three G<math>\gamma</math> proteins, DEP1, GGC2 and GS3, antagonistically regulate grain size. DEP1 and GGC2, individually or in combination, increase grain length when in complex with G<math>\beta</math>. GS3, having no effect on grain size by itself, reduces grain length by competitively interacting with G<math>\beta</math>.</p> <p><b>Example Question:</b> what rice gene GS3 (AGIS_Os12g031280) is?</p> <p><b>Classification:</b> Basic Gene Information Query</p> <p><b>Reference:</b> A G-protein pathway determines grain size in rice</p>
<p><b>Content:</b> Field tests of these plant lines in Beijing in 2018 revealed that OsDREB1C overexpression led to increases in grain yield per plant of 45.1 to 67.6% and in yield per plot of 41.3 to 68.3% compared with wild-type (WT) plants (Fig. 1C). Conversely, OsDREB1C KO resulted in yield decreases (from 16.1 to 29.1% in yield per plant and 13.8 to 27.8% in yield per plot) compared with the WT (Fig. 1, D and E, and table S1). A detailed phenotypic analysis showed that the higher yield of the OsDREB1C-OE lines was mainly attributable to an enhanced grain number per panicle and an increased 1000-grain weight (Fig. 1F and fig. S3C), traits apparently resulting from increased secondary branch number and grain length, width, thickness, and density (fig. S3, A and B and D to K, respectively). The OsDREB1C-OE plants exhibited higher grain yield but reduced strawweight compared with WT plants (Fig. 1G), thus leading to an increased harvest index (the ratio of grain yield to aboveground biomass; Fig. 1H) and raising the possibility that OsDREB1C controls resource allocation between vegetative and reproductive tissues. The harvest index of OsDREB1C-OE plants was increased by 40.3 to 55.7%, whereas it was decreased by 22.4 to 33.7% in OsDREB1C-KO plants (table S1). In addition, key grain quality traits were enhanced in OsDREB1C-OE plants, suggesting that yield improvement does not entail a quality penalty (table S2). Leaves of OsDREB1C-OE plants contained higher levels of photosynthetic pigments (chlorophylls and carotenoids) compared with WT plants, whereas pigment levels were reduced in OsDREB1C-KO plants (fig. S7A). Analysis of leaf mesophyll cells revealed that both chloroplast number and size were increased in OsDREB1C-OE plants (fig. S7, Band C). Analysis of carbon and nitrogen distribution showed that the OsDREB1C-OE plants accumulated more carbon and nitrogen in the grains, but less in their mature leaves, without substantial alterations in the carbon-to-nitrogen ratio (Fig. 3H and fig. S11).</p> <p><b>Example Question:</b> What phenotypes do transgenic plants, rice plants knocking out OsDREB1C genes or plants overexpressing OsDREB1C genes exhibit?</p> <p><b>Classification:</b> Gene Function Experiment Observation</p> <p><b>Reference:</b> A transcriptional regulator that boosts grain yields and shortens the growth duration of rice</p>

Figure 5: An illustration of the Content, Example Question, Classification, and Reference fields.

## B.2 Language Composition and Impact

(1) The initial corpus of 308,727 articles comprises 63% English and 37% Chinese, with this imbalance reflecting the greater availability of English publications; (2) After cleaning, the 1.1 billion-token corpus consists of 75% English and 25% Chinese, with this shift due to the higher accuracy of the MinerU in processing English texts; (3) The final 279 segments used in SeedBench include 49% English and 51% Chinese, achieving balance through manual selection by breeding experts; (4) The 2,264 questions in SeedBench include 45% English and 55% Chinese.

Although LLMs exhibit strong cross-lingual capabilities and the linguistic differences in breeding questions do not alter the underlying scientific logic, we observed response drift when posing the same question in English and Chinese (with cleared histories). This phenomenon suggests potential issues in cross-lingual consistency, which deserves further research—especially when applying LLMs in specific domains where alignment across languages is critical:

- **Question1-EN:** *What effect did the overexpression of OsDREB1C have on the levels of photosynthetic pigments in the leaves of the plants?*  
A. Increased pigment levels  
B. Decreased pigment levels  
C. Pigment levels fluctuated unpredictably  
D. No significant change in pigment levels
- **Question1-CN:** *OsDREB1C 的过表达对植物叶片中光合色素水平有何影响?*  
A. 增加色素水平  
B. 减少色素水平  
C. 色素水平不可预测地波动  
D. 色素水平没有显著变化
- **Question2-EN:** *The expression profile of OsDT11 in different rice tissues was analyzed by \_\_\_\_\_.*
- **Question2-CN:** *OsDT11 在不同水稻组织中的表达谱通过\_\_\_\_\_分析。*

Table 5: Model responses to parallel English and Chinese domain-specific questions.

Model	Ans1-EN	Ans1-CN	Ans2-EN	Ans2-CN
DeepSeek-V3-671B	A	A	qRT-PCR	实时荧光定量PCR
GPT-4	A	A	qRT-PCR	转录组学
OpenAI o1-mini	A	A	qRT-PCR	Northern blot分析
Gemini-1.5-Pro	A	A	qRT-PCR	qRT-PCR

## C Quality Verification

The manual QA verification process involves a comprehensive assessment by experts to ensure the clarity, relevance, and rationality of questions. This includes evaluating whether the questions are designed to elicit accurate and valuable answers. Experts also check the correctness of the answers by consulting authoritative sources and applying professional knowledge to ensure they are error-free. Additionally, for questions that include multiple-answer options, it is essential to assess the rationality of these options. This assessment guarantees a clear distinction between them and ensures their close relevance to the correct answer. To enhance transparency regarding the expert correction process and the resolution of disagreements, we outline the following three points: (1) The composition of the expert panel, consisting of six Ph.D.-level experts in seed breeding; (2) An iterative review process, wherein each question is independently evaluated by at least two experts; (3) A disagreement resolution mechanism that addresses subjectivity in expert assessments, resolved by adopting the intersection of differing expert corrections. Additionally, we have open-sourced samples discarded during the expert correction process as “bad cases” on GitHub, offering readers further insight and reference.

## C.1 Manual Quality Verification

**Question:**

The results of regional trials can provide feedback to breeders, helping them understand the extent to which their breeding goals have been achieved and guiding the breeding direction for \_\_\_\_\_.

**Answer:**

the future



**Expert Review:**

This question does not assess any knowledge points related to **Seed Science and Breeding**.

**Question:**

Which disease is Ningxiangjing 9 (rice variety) resistant to?

A. Bacterial leaf blight of rice

B. Rice sheath blight

C. Rice stripe leaf blight

D. Rice blast

**Answer:**

C



**Expert Review:**

The question cannot be answered precisely and should be phrased as “**Which disease is Ningxiangjing 9 most resistant to?**”

**Question:**

What is the main focus of this study?

**Answer:**

The main focus of this study is to clone a dominant gene for resistance to BPH (Brown Planthopper) from the rice variety Ptb33.



**Expert Review:**

The question contains an unclear reference error that cannot be understood as there is no information about “**this study**” in the context.

**Question:**

In which region is Yuanfengzao (rice variety) suitable for planting?

- A. Pearl River Basin
- B. Yangtze River Basin
- C. Heilongjiang River Basin
- D. Yellow River Basin

**Answer:**

B

**Expert Review:**

The options provided are not appropriate as **there is little difference in the suitability of planting Yuanfengzao in the Yangtze River Basin, the Pearl River Basin, and the Yellow River Basin.** The QA should be transformed into:

**Question:**

In which region is Yuanfengzao suitable for planting?

- A. Songhua River Basin
- B. Liao River Basin
- C. Hai River Basin
- D. Yangtze River Basin

**Answer:**

D

## C.2 Descriptive Statistics

Table 6: Descriptive Statistics of Generated Questions and Answers

	C1	C2	C3	C4	C5	C6	C7	C8	C9	C10	total
QA-1	37 (1.63%)	9 (0.4%)	0 (0.0%)	23 (1.02%)	2 (0.09%)	33 (1.46%)	15 (0.66%)	32 (1.41%)	38 (1.68%)	11 (0.49%)	200 (8.83%)
QA-2	37 (1.63%)	10 (0.44%)	5 (0.22%)	44 (1.94%)	2 (0.09%)	26 (1.15%)	17 (0.75%)	11 (0.49%)	9 (0.4%)	26 (1.15%)	187 (8.26%)
QA-3	38 (1.68%)	32 (1.41%)	0 (0.0%)	34 (1.5%)	5 (0.22%)	13 (0.57%)	20 (0.88%)	23 (1.02%)	34 (1.5%)	25 (1.1%)	224 (9.89%)
QA-4	66 (2.92%)	12 (0.53%)	14 (0.62%)	47 (2.08%)	5 (0.22%)	34 (1.5%)	14 (0.62%)	16 (0.71%)	14 (0.62%)	20 (0.88%)	242 (10.69%)
SUM-1	39 (1.72%)	17 (0.75%)	6 (0.27%)	36 (1.59%)	4 (0.18%)	25 (1.1%)	20 (0.88%)	26 (1.15%)	28 (1.24%)	24 (1.06%)	225 (9.94%)
SUM-2	39 (1.72%)	17 (0.75%)	6 (0.27%)	36 (1.59%)	4 (0.18%)	25 (1.1%)	20 (0.88%)	26 (1.15%)	28 (1.24%)	24 (1.06%)	225 (9.94%)
RC-1	25 (1.1%)	7 (0.31%)	0 (0.0%)	15 (0.66%)	1 (0.04%)	17 (0.75%)	9 (0.4%)	16 (0.71%)	17 (0.75%)	6 (0.27%)	113 (4.99%)
RC-2	18 (0.8%)	4 (0.18%)	4 (0.18%)	32 (1.41%)	1 (0.04%)	15 (0.66%)	7 (0.31%)	6 (0.27%)	7 (0.31%)	14 (0.62%)	108 (4.77%)
RC-3	38 (1.68%)	32 (1.41%)	0 (0.0%)	33 (1.46%)	5 (0.22%)	13 (0.57%)	20 (0.88%)	23 (1.02%)	34 (1.5%)	23 (1.02%)	221 (9.76%)
RC-4	65 (2.87%)	12 (0.53%)	14 (0.62%)	47 (2.08%)	5 (0.22%)	33 (1.46%)	14 (0.62%)	16 (0.71%)	14 (0.62%)	20 (0.88%)	240 (10.6%)
RC-5	58 (2.56%)	26 (1.15%)	24 (1.06%)	43 (1.9%)	4 (0.18%)	26 (1.15%)	20 (0.88%)	26 (1.15%)	28 (1.24%)	24 (1.06%)	279 (12.32%)
<b>total</b>	460 (20.32%)	178 (7.86%)	73 (3.22%)	390 (17.23%)	38 (1.68%)	260 (11.48%)	176 (7.77%)	221 (9.76%)	251 (11.09%)	217 (9.58%)	2264 (100.0%)

## D Additional Experimental Setup

### D.1 Model Hyperparameters

We conduct evaluations on different large language models using opencompass as the primary tool. In order to ensure reproducibility and provide a reference for future research, the main hyperparameters used in our experiments, along with their meanings, are listed in Table 7:

Table 7: Main Hyperparameters for LLM Evaluation

Parameter	Meaning	Value
max_seq_len	The maximum context length (upper limit of input tokens)	7168
max_out_len	The maximum output length (upper limit of generated tokens)	2048
batch_size	The batch size (number of requests processed per generation)	80

For proprietary LLMs, the maximum generation tokens for each model were set to 2048, with a batch size of 80. The temperature was set to 0.7, and top- $p$  and top- $k$  were set to 0.8 and 10, respectively. For open-source LLMs, the maximum generation tokens were also set to 2048, with a batch size of 80. The temperature was set to 0.7, and top- $p$  and top- $k$  were set to 0.8 and 10, respectively.

### D.2 Evaluation Metrics

The evaluation process consists of two steps: the first step is answer extraction. After collecting the model’s responses, we first perform post-processing to extract the model’s replies. The second step is metric calculation. During evaluation, we have designed different evaluation methods for different question types:

**Single-choice Questions.** We use Accuracy as the evaluation metric to calculate the correctness of the model’s answers. The formula for Accuracy is:

$$\text{Accuracy} = \frac{1}{N_{\text{total}}} \sum_{i=1}^{N_{\text{total}}} \delta(y_i^{\text{pred}}, y_i^{\text{true}})$$

where  $N_{\text{total}}$  is the total number of questions,  $y_i^{\text{pred}}$  is the predicted answer for the  $i$ -th question by the model, and  $y_i^{\text{true}}$  is the true answer for the  $i$ -th question.  $\delta(a, b)$  is the Kronecker delta function, where  $\delta(a, b) = 1$  if  $a = b$ , and  $\delta(a, b) = 0$  otherwise.

**Multiple-choice Questions.** We use the **F1 score**, which balances the correct and incorrect answers from the model to provide a more comprehensive reflection of model performance. The formula for the F1 score is:

$$F1 = 2 \cdot \frac{P \cdot R}{P + R}$$

where precision ( $P$ ) is:

$$P = \frac{1}{N} \sum_{i=1}^N \frac{|Y_i^{\text{pred}} \cap Y_i^{\text{true}}|}{|Y_i^{\text{pred}}|}$$

and recall ( $R$ ) is:

$$R = \frac{1}{N} \sum_{i=1}^N \frac{|Y_i^{\text{pred}} \cap Y_i^{\text{true}}|}{|Y_i^{\text{true}}|}$$

where  $N$  is the total number of samples,  $Y_i^{\text{pred}}$  is the set of predicted answers for the  $i$ -th sample, and  $Y_i^{\text{true}}$  is the set of true answers for the  $i$ -th sample.

**Fill-in-the-Blank and Generation Questions.** We use ROUGE for evaluation and calculate the average F1 score of ROUGE-L. For fill-in-the-blank questions, since answers typically contain multiple blanks or paragraphs, we adopt a segmented evaluation method, comparing the similarity of each segment between the model’s predicted answer and the reference answer, and averaging the scores of multiple segments to more accurately reflect the model’s performance for each answer item. For generation questions, we use a sentence-level evaluation method, directly calculating the similarity between the model’s generated full answer and the reference answer. This method allows for a comprehensive evaluation of the overall fluency and semantic accuracy of the generated content.

The formula for ROUGE-L F1 is:

$$\text{ROUGE-L} = \frac{(1 + \beta^2) \cdot P \cdot R}{\beta^2 \cdot P + R}$$

where precision ( $P$ ) is:

$$P = \frac{1}{N} \sum_{i=1}^N \frac{\text{LCS}(X_i, Y_i)}{|X_i|}$$

and recall ( $R$ ) is:

$$R = \frac{1}{N} \sum_{i=1}^N \frac{\text{LCS}(X_i, Y_i)}{|Y_i|}$$

where  $N$  is the total number of samples,  $X_i$  and  $Y_i$  represent the predicted and reference answers for the  $i$ -th sample, and  $\text{LCS}(X_i, Y_i)$  is the length of the longest common subsequence between  $X_i$  and  $Y_i$ .  $|X_i|$  and  $|Y_i|$  are the lengths of  $X_i$  and  $Y_i$ , respectively.  $\beta$  is the balance factor between precision and recall, typically set to 1.

## E Error Analysis of Large Language Model Outputs

In this chapter, we thoroughly analyze the errors produced by large language models in different scenarios, according to the three task types defined above. By presenting typical cases and discussing both the causes and potential improvements, we provide a reference for optimizing the quality of LLM outputs in subsequent work.

### E.1 Errors in the First Task Type

The first task type primarily involves gene information retrieval capabilities. The model needs to demonstrate fundamental biological knowledge as well as accurate information retrieval. We observe several representative error cases, outlined below:

**Gene Name Confusion** When dealing with similar or homonymous gene names or gene identifiers, the model incorrectly provides the wrong gene sequences and functional descriptions. This mistake can lead to misalignment in subsequent gene function validation.

#### Example 1

**Example Question:**

What is the identifier of the OsAL50 gene in The Rice Annotation Project Database (RAP-DB)?

**Reference:**

Os01g0306800

**Answer:**

Os03g0559000

**Analysis:**

The model mistakenly identified "Os01g0306800" as "Os03g0559000," resulting in gene identifier confusion.

#### Example 2

**Example Question:**

GWD1 gene is the same gene as \_\_\_\_\_?

**Reference:**

LSE1 (line leaf starch excess 1)

**Answer:**

GWD2

**Analysis:**

The model incorrectly identified GWD1 and GWD2 as the same gene, leading to confusion in gene naming/relationships.

#### Example 3

**Example Question:**

What is the identifier of the GWD1 gene in The Rice Annotation Project Database (RAP-DB)?

**Reference:**

Os06g0498400

**Answer:**

GRS0000G0001999, GRS0000G0002000, GRS0000G0002001

**Analysis:**

The model returned a set of non-existent identifiers (which also do not conform to the RAP-DB format), causing gene identifier confusion.

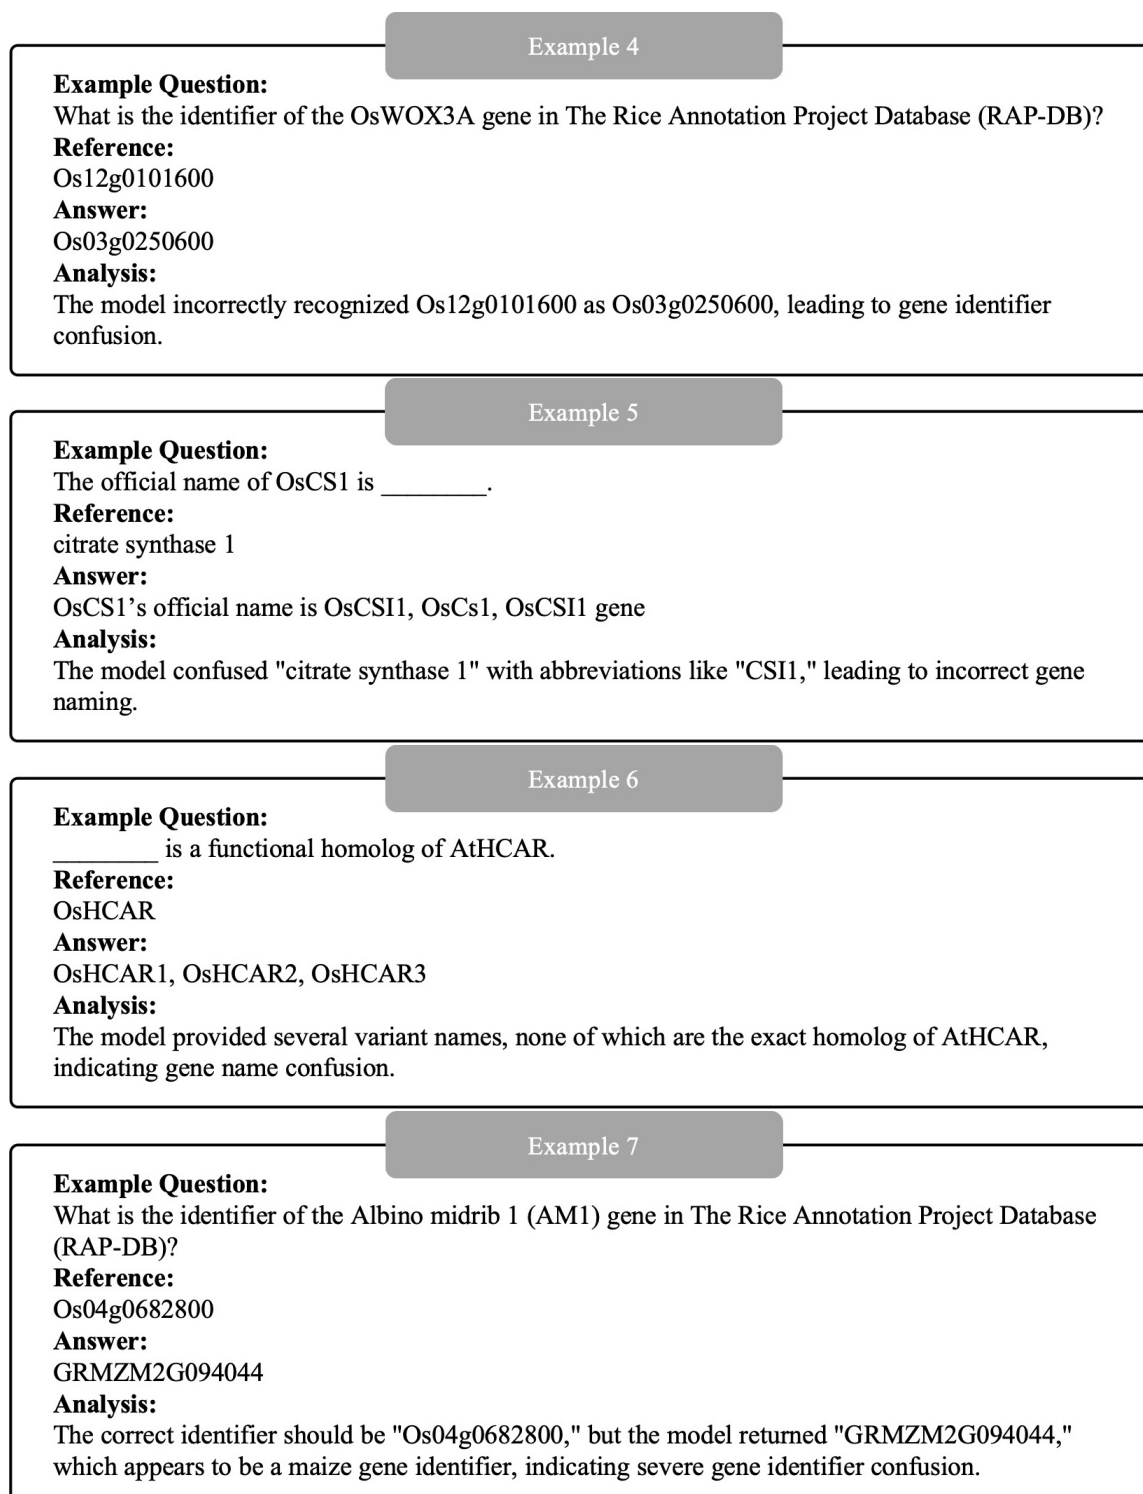


Figure 6: Example illustrating gene name confusion.

**Errors in Gene Sequence and Positional Information** The model provides incorrect answers regarding the physical location information of genes (chromosome number, start and end coordinates, sequence length, etc.), such as incorrect gene coordinates or sequence length discrepancies with actual databases. This may impact genome annotation and structural variation analysis, leading to experimental design biases.

#### Example 1

**Example Question:**

OsAL50 gene plays an important role in \_\_\_\_\_.

**Reference:**

Regulation of photosynthetic pigment metabolism, photosynthesis, and chloroplast formation

**Answer:**

scutellarein pathway, decarboxylation pathway

**Analysis:**

The model only answered "Rice," deviating from the actual functional or process description. The question expects an answer related to physiological metabolism or processes, but the model provided a more general species name.

#### Example 2

**Example Question:**

What protein does the OsAL50 gene encode?

**Reference:**

Endoribonuclease E-like protein

**Answer:**

Rice disease resistance-related protein

**Analysis:**

The model confused the specific protein type encoded by the gene, providing an entirely different functional protein description.

#### Example 3

**Example Question:**

What enzyme does the GWD1 gene encode?

**Reference:**

$\alpha$ -glucan water dikinase

**Answer:**

Glutamine synthetase

**Analysis:**

The model confused enzyme categories, incorrectly answering "glutamine synthetase" instead of " $\alpha$ -glucan water dikinase," leading to a serious functional misjudgment.

#### Example 4

**Example Question:**

NAL2 (narrow leaf2) and NAL3 are paralogs. What transcriptional activator do they encode?

**Reference:**

OsWOX3A

**Answer:**

NAC-domain-containing protein

**Analysis:**

The model mistakenly assumed that the gene encodes a NAC family transcription factor rather than a WOX family gene product, resulting in a functional description deviation.

#### Example 5

**Example Question:**

AtHCAR is a key enzyme in the \_\_\_\_\_ pathway for Chl degradation.

**Reference:**

PAO/phyllobilin

**Answer:**

OsCS1's official name is OsCS11, OsCs1, OsCS11 gene

**Analysis:**

The model confused the major pathways for chlorophyll degradation, providing irrelevant or rare metabolic pathway names.

#### Example 6

**Example Question:**

AtHCAR catalyzes the conversion of \_\_\_\_\_ to Chl a in vivo and in vitro.

**Reference:**

7-HMC a

**Answer:**

protochlorophyllide, Pchlde, ...

**Analysis:**

The model mistakenly confused the substrate "7-HMC a" with another chlorophyll precursor "Pchlde," leading to inaccurate functional descriptions.

#### Example 7

**Example Question:**

The athcar mutant shows a \_\_\_\_\_ phenotype during DIS.

**Reference:**

stay-green

**Answer:**

dwarfed, sterile flowering

**Analysis:**

The model incorrectly described the mutant phenotype, misidentifying "stay-green" as "dwarfed, sterile flowering," which are unrelated phenotypes.

Figure 7: Example illustrating cross-species gene information errors.

**Errors in Gene Function and Regulation** The model exhibits errors in describing gene functions, protein products, or regulatory mechanisms, such as misclassifying gene functions, misunderstanding their roles in signaling pathways, or incorrectly predicting protein products. These issues may mislead researchers' understanding of gene biological functions, potentially leading to incorrect downstream experiments and data analysis.

#### Example 1

**Example Question:**

On which chromosome is the OsAL50 gene located in rice? What is its specific position range?

**Reference:**

Chromosome 1

**Answer:**

Chromosome 4

**Analysis:**

The model made an error in answering the chromosome location, confusing the correct chromosome.

#### Example 2

**Example Question:**

What is the full sequence length of the OsAL50 gene?

**Reference:**

8917 bp

**Answer:**

1199

**Analysis:**

The model provided an incorrect sequence length number, which is significantly shorter than the actual length.

#### Example 3

**Example Question:**

On which chromosome strand is the GWD1 gene located?

**Reference:**

Chromosome 6

**Answer:**

5

**Analysis:**

The model provided an incorrect chromosome number.

#### Example 4

**Example Question:**

What is the full name of the OsUAM3 gene?

**Reference:**

UDP-ARABINOPYRANOSE MUTASE 3 gene

**Answer:**

The full name of the OsUAM3 gene is OsUAM3-like gene. Please note that this name may not be a widely recognized standard gene name, as OsUAM3-like suggests it might be a gene similar to OsUAM3 rather than a specifically recognized gene name. The exact name may need to be confirmed based on the latest scientific literature or databases.

**Analysis:**

In the actual database annotation, OsUAM3 corresponds to "UDP-ARABINOPYRANOSE MUTASE 3," but the model mistakenly assumed it to be an "OsUAM3-like" variant, indicating a lack of understanding of the full gene name.

#### Example 5

**Example Question:**

How is IPA1 regulated in vivo?

**Reference:**

IPA1 is regulated by microRNA (OsmiR156) in vivo.

**Answer:**

IPA1 regulation in vivo is mainly achieved through transcriptional control. Specifically, IPA1 expression is regulated by multiple transcription factors, including members of the MYB, bHLH, and WD40 families. These transcription factors regulate IPA1 expression by binding to specific cis-acting elements in the IPA1 gene promoter region. Additionally, environmental factors such as light and hormone levels also influence the activity of these transcription factors, thereby indirectly regulating IPA1 expression. In some cases, IPA1 expression may also be regulated by microRNA.

**Analysis:**

According to literature, IPA1 is most classically and definitively regulated by miRNA156 in rice. However, the model provided a general "transcription factor group" regulation mechanism, which, although biologically plausible, deviates from the widely recognized primary regulatory pathway, making it an off-topic answer or a case of key information omission.

Figure 8: Examples of gene function and regulatory errors.

## E.2 Errors in the Second Task Type

The second task type concerns gene function and regulatory mechanisms, requiring more advanced logical reasoning and deeper domain knowledge. The model is expected to demonstrate “causal understanding” or “upstream/downstream interactions” thinking in gene function validation and regulatory network analysis. Similarly, we observe several representative error types, listed below:

**Knowledge Errors** Because the model lacks thorough understanding of the background or domain-specific knowledge, its answers can be incorrect. In Table X (not shown), we list erroneous examples where the model misinterprets domain-specific terminology or experimental data, resulting in inaccurate responses.

### Example 1

**Example Question:**

GL10 overexpression results in a significant increase in grain length and \_\_\_\_\_ compared to NIL-gl10.

**Reference:**

grain weight

**Answer:**

grain width

**Analysis:**

The model confused "grain weight" with "grain width," indicating an inaccurate understanding of specific experimental results, which constitutes a clear professional knowledge error.

### Example 2

**Example Question:**

To which protein family does the OsJRL45 gene belong?

**Reference:**

Jacalin-related lectin (JRL) family

**Answer:**

Photoreceptor family

**Analysis:**

The model incorrectly classified the OsJRL45 gene into the "Photoreceptor family." It confused the Jacalin-related lectin family with other protein families that may contain keywords such as "light" or "pigment," demonstrating a significant background knowledge bias.

### Example 3

**Example Question:**

What method was used in the research on LOC\_Os06g03240 (Bph32 gene)?

**Reference:**

The method used was the transformation of a susceptible indica variety (Kasalath) with the cDNA sequence of the Ptb33 Bph32 gene.

**Answer:**

CRISPR/Cas9 gene-editing technology

**Analysis:**

The core method described in the literature was the introduction of the "cDNA sequence of the Bph32 gene" into a susceptible indica rice variety to verify its function. However, the model answered with CRISPR/Cas9 editing, which contradicts known facts, indicating a lack of understanding of literature details or professional background.

#### Example 4

**Example Question:**

In which database is OsJRL45 recorded under the gene identifier T2T?

**Reference:**

RiceSuperPIRdb

**Answer:**

T2T, OsJRL45

**Analysis:**

The question asks for the database where the gene identifier T2T is recorded, but the model only repeated "T2T, OsJRL45" without providing the correct database name. This represents a clear knowledge error or a failure to distinguish between databases and identifiers.

#### Example 5

**Example Question:**

Which protein domain does the OsJRL45 gene-encoded protein contain?

**Reference:**

jacalin-like lectin

**Answer:**

RNA-binding domain

**Analysis:**

The model confused the "jacalin-like lectin" domain with the "RNA-binding domain." This is a typical domain knowledge error, indicating that the model failed to correctly recognize or understand the protein domain annotation related to this gene.

Figure 9: Example illustrating knowledge errors in gene function and regulation.

**Insufficient Answer Precision** When the question requires detailed specifics, such as particular experimental methods, tools, or procedures, the model fails to provide precise information and instead offers overly generic descriptions. As a result, the answer does not meet the exact requirements of the question, indicating over-generalization.

#### Example 1

**Example Question:**

What were the findings regarding the effect of elevated DEP1 accumulation on grain length?

**Reference:**

Elevated DEP1 accumulation increased the grain length by 6.85 – 9.58% with a normal plant stature.

**Answer:**

Elevated DEP1 accumulation was found to significantly increase grain length in the studied plants. Specifically, the grains from plants with increased DEP1 levels were approximately 10-15% longer compared to the control group. This effect was consistent across different genotypes and environmental conditions, suggesting a robust and reliable outcome. However, further research is needed to understand the underlying molecular mechanisms and potential impacts on grain yield and quality.

**Analysis:**

While the model's response is directionally correct, the specific values (10-15%) do not match the actual literature (6.85-9.58%). Additionally, the model adds extraneous, generalized information that was not required by the question, failing to provide the exact percentage increase in grain length.

#### Example 2

**Example Question:**

According to Gene Ontology (GO) analysis, what are the predicted functions of the APX9 gene?

**Reference:**

Examples include the top five GO annotations (e.g., response to reactive oxygen species, catalytic activity, peroxidase activity, etc.).

**Answer:**

The predicted functions of the APX9 gene in GO analysis mainly include antioxidant activity, hydrogen peroxide metabolism, and antioxidant defense systems.

**Analysis:**

The question requires specific GO functional entries or IDs, but the model only provides broad descriptions related to "antioxidant activity" without matching the corresponding GO terms, making the answer imprecise.

#### Example 3

**Example Question:**

The significant reductions in C30, C32, and C34 waxes in wsl4-1 leaves were \_\_\_\_\_ in the complementation line.

**Reference:**

rescued

**Answer:**

restored

**Analysis:**

In the context of genetics and functional validation, "rescued" specifically means "restoring the original phenotype through exogenous gene expression or treatment." While "restored" is a similar term, it is not equivalent in this precise scientific context, making this a vocabulary accuracy issue.

Figure 10: Example illustrating lack of precision in model answers.

**Limited Reasoning Ability** The model does not incorporate the given information to perform logical analysis and contextual interpretation, thereby producing responses that do not meet practical needs. It either fails to link the contextual clues correctly or ignores implicit details in the question.

#### Example 1

**Example Question:**

What was the effect of overexpressing RFT1 in the study?

**Reference:**

Overexpression of RFT1 resulted in the direct formation of spikelets from most of the transgenic calli.

**Answer:**

Overexpressing RFT1 in the study led to earlier flowering in plants, which is beneficial for breeding programs in temperate regions where early flowering can help avoid summer heat stress. Additionally, it enhanced flowering under long-day conditions, indicating its potential role in regulating flowering time under different photoperiods.

**Analysis:**

The most significant and specific experimental observation was "the direct formation of spikelets from transgenic calli." However, the model omitted this core experimental result, instead providing a general discussion of "earlier flowering," indicating that it failed to deeply analyze or match the contextual information in the question.

#### Example 2

**Example Question:**

What is the name of the gene discussed in the text?

**Reference:**

The name of the gene discussed in the text is OsCHR709.

**Answer:**

Unable to answer because the text does not provide it.

**Analysis:**

The question explicitly refers to "OsCHR709" as the target gene; the model should have been able to answer by connecting contextual information. However, it claimed it "could not answer," demonstrating insufficient contextual information extraction and reasoning ability.

Figure 11: Example illustrating limited reasoning ability.

### E.3 Errors in the Third Task Type

The third task type encompasses a wide range of needs related to variety selection and agronomic traits, requiring the model to integrate gene-level knowledge with real-world agricultural management experience. We observe several representative error cases, listed below:

**Insufficient Semantic Understanding** The model does not fully grasp the core descriptions or context in the question, leading to incorrect interpretations of critical information. Consequently, the final output deviates from the reference solution.

#### Example 1

**Example Question:**

How is gene editing technology performed?

**Reference:**

Gene editing is a process that modifies specific targets in an organism's genome using gene editing technologies. It enables precise and efficient gene insertion, deletion, or replacement, thereby altering genetic information and phenotypic traits.

**Answer:**

Gene editing technology typically follows these steps:

1. Target Gene Selection: Identify the target gene to be edited, usually based on an understanding of specific traits or functions.
  2. Guide Sequence Design: Use specific sequences (such as guide RNA in the CRISPR-Cas9 system) to precisely locate the target gene.
  3. Delivery of Editing Tools: Introduce the editing tools (such as Cas9 enzyme) and guide sequences into target cells. This can be achieved through various methods, including direct injection, viral vectors, or electroporation.
  4. Gene Editing: The editing tool makes cuts at the target gene site, triggering the cell's repair mechanisms to insert, delete, or replace the gene.
  5. Screening and Validation: Molecular biology techniques are used to screen successfully edited cells and verify whether the editing occurred as expected.
- This technology can be applied to improving disease resistance, stress tolerance, yield, or nutritional value in crops.

**Analysis:**

Both the prompt and the reference answer specify that "the response should only include the answer to the question, without additional content." However, the model provides a structured step-by-step explanation and additional applications of the technology, failing to strictly follow the instruction of "only providing the answer itself without unnecessary elaboration." This is a typical example of inadequate comprehension of the prompt instructions.

#### Example 2

**Example Question:**

The paternal parent of Koshihikari is \_\_\_\_\_, and the maternal parent is \_\_\_\_\_.

**Reference:**

Norin 1, Norin 22.

**Answer:**

Dek, Dongnong 288.

**Analysis:**

The documented pedigree of Koshihikari is well established, yet the model provides completely unrelated parent names, showing a fundamental misunderstanding of rice variety lineage.

Figure 12: Example illustrating insufficient semantic understanding in variety selection tasks.

**Knowledge Errors** Similar to the other task types, this category also shows frequent knowledge-related mistakes. Lacking the necessary domain background or relevant information leads the model to give incorrect answers.

#### Example 1

**Example Question:**

During the breeding process of Longjing 13, in 1995, Longhua 91-340 was used as the maternal parent, and \_\_\_\_\_ was used as the paternal parent for sexual hybridization.

**Reference:**

Kongyu 139.

**Answer:**

Hei 2499.

**Analysis:**

The official breeding records clearly state that “Kongyu 139” was used as the paternal parent. The model’s incorrect response of “Hei 2499” indicates a significant knowledge gap in breeding lineage.

#### Example 2

**Example Question:**

The paternal parent of Koshihikari is \_\_\_\_\_, and the maternal parent is \_\_\_\_\_.

**Reference:**

Norin 1, Norin 22.

**Answer:**

Dek, Dongnong 288.

**Analysis:**

The documented pedigree of Koshihikari is well established, yet the model provides completely unrelated parent names, showing a fundamental misunderstanding of rice variety lineage.

#### Example 3

**Example Question:**

Yuxiangyouzhan was developed through the hybridization of \_\_\_\_\_ and \_\_\_\_\_, followed by multiple generations of directional selection to create a photoperiod-sensitive indica conventional rice variety.

**Reference:**

TY36, IR100.

**Answer:**

Early indica rice, late indica rice.

**Analysis:**

The model’s response is overly generic, providing broad rice categories rather than the specific parent varieties “TY36” and “IR100,” highlighting a deficiency in precise knowledge of hybrid breeding lineages.

#### Example 4

**Example Question:**

Longjing 13 is a rice variety bred by \_\_\_\_\_.

**Reference:**

Heilongjiang Academy of Agricultural Sciences, Rice Research Institute.

**Answer:**

Liaoning Academy of Agricultural Sciences, Rice Research Institute.

**Analysis:**

The correct breeding institution for this variety is the “Heilongjiang Academy of Agricultural Sciences, Rice Research Institute,” but the model incorrectly states “Liaoning Academy of Agricultural Sciences, Rice Research Institute,” demonstrating a factual knowledge error or confusion of institutions.

#### Example 5

**Example Question:**

The breeding technology of Longjing 13 is a comprehensive breeding technology that combines conventional techniques with \_\_\_\_\_.

**Reference:**

Anther culture technology.

**Answer:**

Hybridization technology.

**Analysis:**

The model mistakenly identifies “hybridization technology” instead of the correct “anther culture technology,” revealing a misunderstanding of the breeding process. This is a classic example of a domain-specific knowledge error.


Figure 13: Example illustrating knowledge errors in variety selection and agronomic traits.

## F Prompt Templates

### F.1 Multiple-Choice Question Generation (Task QA-1)

**Prompt for Task QA-1**

**Topic Description:**  
Instruct the model to generate multiple-choice questions.



**User:**  
Please generate independent multiple-choice questions based on the provided context. Each multiple-choice question should contain four options, with one correct answer and three distractors. Generate {num\_of\_questions} questions for each context.

Please note:


1. If the context is in Chinese, the questions should also be in Chinese. If the context is in English, the questions should be in English.
2. The correct answer for each question should strictly follow the specified format, with no extra output.
3. Each question should be complete and independent, without relying on the answers to other questions.

{example\_1}

{example\_2}

Contextual materials:  
{context}

Below are the {num\_of\_questions} multiple-choice questions generated based on the given description:



**Example:**  
Question: The field test conducted in Beijing in 2018 on plant lines found that overexpression of OsDREB1C led to what effect on grain yield per plant?  
A. Decrease by 45.1 to 67.6%  
B. Increase by 45.1 to 67.6%  
C. No significant change  
D. Decrease by 16.1 to 29.1%  
Answer: B

Question: In the same field test, knockout (KO) of OsDREB1C resulted in what effect on yield per plot?  
A. Increase by 13.8 to 27.8%  
B. Increase by 41.3 to 68.3%  
C. No significant change  
D. Decrease by 13.8 to 27.8%  
Answer: D  
...

Figure 14: Illustration for multiple-choice question generation.

## F.2 Multiple-Answer Question Generation (Task QA-2)

### Prompt for Task QA-2

#### Topic Description:

Instruct the model to generate multiple-answer questions.



#### User:

Please generate independent multiple-answer questions based on the context, with each question containing four options, some of which are correct answers while others are distractors. Generate {num\_of\_questions} multiple-choice questions for each piece of context.

Please note:

1. If the context is in Chinese, the questions should also be in Chinese. If the context is in English, the questions should be in English.
2. The correct answer for each question should strictly follow the specified format, with no extra output.
3. Each question should be complete and independent, without relying on the answers to other questions.

{example\_1}

{example\_2}

Contextual materials:

{context}

Below are the {num\_of\_questions} multiple-answer questions generated based on the given description:



#### Example:

Question: What are the observed characteristics of the Bg1-D mutant compared to the wild type (WT)?

- A. Increased grain size
- B. Decreased plant height
- C. Longer and wider leaves
- D. Smaller panicles

Answer: A, C

Question: What changes were observed in the cells of the Bg1-D mutant?

- A. Increased area of parenchyma cells
- B. Decreased number of parenchyma cells
- C. Down-regulation of genes associated with cell cycle and cell expansion
- D. Longer epidermis cells of both palea and

Answer: A, D

...

Figure 15: Illustration for multiple-answer question generation.

### F.3 Fill-in-the-Blank Question Generation (Task QA-3)

#### Prompt for Task QA-3

##### Topic Description:

Instruct the model to generate fill-in-the-blank questions.



##### User:

Please generate independent fill-in-the-blank questions based on the provided context. Fill-in-the-blank questions typically require filling in the gaps according to the requirements of the question. Generate {num\_of\_questions} fill-in-the-blank questions for each contextual material.

Please note:

1. If the context is in Chinese, the questions should also be in Chinese. If the context is in English, the questions should be in English.
2. The correct answer for each question should strictly follow the specified format, with no extra output.
3. Each question should be complete and independent, without relying on the answers to other questions.

{example\_1}

{example\_2}

Contextual materials:

{context}

Below are the {num\_of\_questions} fill-in-the-blank questions generated based on the given description:



##### Example:

Question: AFD1 transcripts were found in all examined \_\_\_\_\_ and organs.

Answer: tissues

Question: AFD1 exhibited more abundant expressions in \_\_\_\_\_, young panicles, lemmas, paleae, and lodicules than those in other tissues or organs examined.

Answer: internodes

...

Figure 16: Illustration for fill-in-the-blank question generation.

## F.4 Text-based Q&A Generation (Task QA-4)

### Prompt for Task QA-4

**Topic Description:**

Instruct the model to generate text generation questions.

**User:**

Please generate independent knowledge-based questions and answers based on the provided context. Both the input and output of the questions and answers should be in text format. Generate {num\_of\_questions} questions and answers for each context.

Please note:

1. If the context is in Chinese, the questions should also be in Chinese. If the context is in English, the questions should be in English.
2. The correct answer for each question should strictly follow the specified format, with no extra output.
3. Each question should be complete and independent, without relying on the answers to other questions.

{example\_1}

{example\_2}

Contextual materials:

{context}

Below are the {num\_of\_questions} text generation questions generated based on the given description:

**Example:**

Question: What role does the  $G\alpha$  protein play in the heterotrimeric G proteins pathway?

Answer: The  $G\alpha$  protein provides a foundation for grain size expansion.

Question: How do the three  $G\gamma$  proteins DEP1, GGC2 and GS3 regulate grain size?

Answer: DEP1 and GGC2 increase grain length when in complex with  $G\beta$ , either individually or in combination. GS3 reduces grain length by competitively interacting with  $G\beta$ , but has no effect on grain size by itself.

...

Figure 17: Illustration for text-based Q&A generation.

## F.5 Naive Summarization (Task SUM-1)

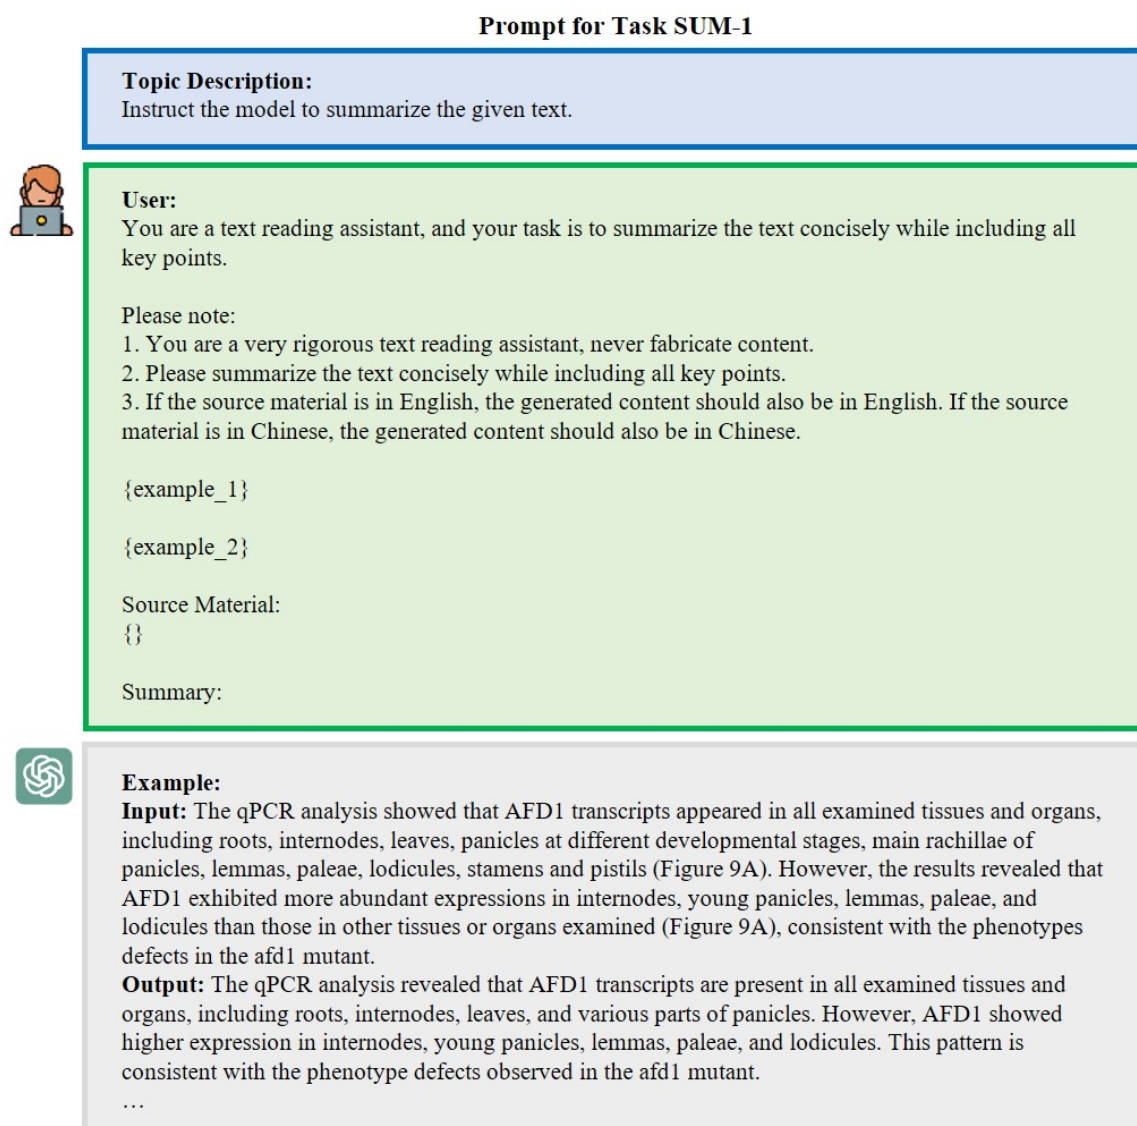



Figure 18: Illustration for naive summarization.


## F.6 Key Information Extraction (Task SUM-2)

**Prompt for Task SUM-2**

**Topic Description:**  
Instruct the model to extract key information.



**User:**  
Please extract the key information from the following description.  
---  
Task: Extract key information.  
Goal: The extracted content should be as comprehensive as possible and cover the given text.  
Method: Analyze the text, identify and extract key information points from the text.  
Language restriction: If the description text is in Chinese, the extracted key information should also be in Chinese. If the description text is in English, the extracted key information should also be in English.  
---  
{example\_1}  
{example\_2}  
Description:  
{}  
  
Extracted key information:



**Example:**  
Research topic: qPCR analysis of AFD1 transcripts  
Findings:  
- AFD1 transcripts were found in all examined tissues and organs, including roots, internodes, leaves, panicles at different developmental stages, main rachillae of panicles, lemmas, paleae, lodicules, stamens, and pistils.  
- AFD1 exhibited more abundant expressions in internodes, young panicles, lemmas, paleae, and lodicules than those in other tissues or organs examined.  
- The results are consistent with the phenotypes defects in the *afd1* mutant.  
...

Figure 19: Illustration for key information extraction.

## F.7 Question Rewriting (Tasks RC-1, RC-2, RC-3, RC-4)

### Prompt for Task RC-1, RC-2, RC-3, RC-4

#### Topic Description:

Instruct the model to rephrase the question.



#### User:

You are a language expert responsible for rephrasing specific questions. The requirements are as follows:

1. Change the wording and sentence structure without altering the original meaning.
2. Maintain the core meaning of the original sentence without introducing new concepts or distorting the original intent.
3. Ensure fluency and accuracy in the language.
4. The original questions can be multiple-choice, single-choice, fill-in-the-blank, or open-ended; when rephrasing, the original question type must be preserved.
5. Do not change the language of the question; if the original question is in Chinese, the rephrased question should also be in Chinese, and if the original question is in English, the rephrased question should be in English as well.

{example\_1}

{example\_2}

{original question}

Rephrased question:



#### Example:

**Input:** The field test conducted in Beijing in 2018 on plant lines found that overexpression of OsDREB1C led to what effect on grain yield per plant?

**Output:** What was the impact on grain yield per plant as a result of overexpressing OsDREB1C in the 2018 Beijing field test on plant lines?

...

Figure 20: Illustration for question rewriting across multiple subtasks.

## G Robustness Evaluation of Prompts

This section analyzes the sensitivity of model outputs to different prompt templates and evaluates how various prompt styles impact model performance. Specifically, we modify the wording and style of prompts and compare how different styles affect the consistency and stability of model-generated results.

### Example 1

**Question:**

a pathway composed of five subunits of the heterotrimeric G proteins that regulate grain length in rice. The G $\beta$  protein is essential for plant survival and growth. G $\alpha$  provides a foundation for grain size expansion. Three G $\gamma$  proteins, DEP1, GGC2 and GS3, antagonistically regulate grain size. DEP1 and GGC2, individually or in combination, increase grain length when in complex with G $\beta$ . GS3, having no effect on grain size by itself, reduces grain length by competitively interacting with G $\beta$ . How does GS3 influence the length of the grain?

**Reference:**

GS3 influences the length of the grain by competitively interacting with G $\beta$ , thereby reducing grain length.

**Instruction A:**

As a researcher in the field of agricultural bioengineering, please provide a professional response to the following agricultural question.

Important notes: Your response must be precise and limited strictly to the answer itself. The question is as follows:

**Answer A:**

GS3 influences the length of the grain by competitively interacting with G $\beta$ , thereby reducing grain length.

**Score A:**

54.9

**Instruction B:**

This is a question about agriculture. As an agricultural university student, you should answer according to the following rules:

The input and output of the question are both in text format. Note: Your answer should include only the response to the question, without any additional content.

Question:

**Answer B:**

GS3 reduces grain length by competitively interacting with G $\beta$ .

**Score B:**

68.7

## Example 2

### Question:

OsSPL14 is localized to the nucleus (Fig. 1g), consistent with a role as a transcription factor. RNA in situ hybridization revealed that OsSPL14 was predominantly expressed in the shoot apex at both the vegetative (Fig. 2a) and reproductive (Fig. 2b) stages. It was also highly expressed in the promordia of primary and secondary branches (Fig. 2c). Expression patterns of OsSPL14 and OsmiR156 in various rice organs revealed by real-time PCR and miRNA gel blot analyses showed that OsSPL14 was highly expressed in the culm and shoot apex, which is complementary with the expression pattern of OsmiR156 in vivo (Fig. 2e,f). Consistently, overexpression of OsmiR156 resulted in a substantial decrease in OsSPL14 transcripts (Supplementary Fig. 7a), whereas the interruption of OsmiR156 (overexpression of MIM156) led to a marked increase in OsSPL14 transcripts (Supplementary Fig. 7b). These results indicated that OsSPL14 was regulated by the OsmiR156-directed cleavage in vivo. In which part of the plant is OsSPL14 mainly expressed?

- A. In the flower apex
- B. In the leaf apex
- C. In the root apex
- D. In the shoot apex

### Reference:

D

### Instruction A:

Now, please act as an agricultural university student and complete this multiple-choice question. The task requires selecting all correct answers from the four options: A, B, C, and D.

Answering rules: Your answer should only include the correct option letters. If multiple answers are correct, separate them with a comma. Do not add any extra text or spaces. For example, if A and B are correct, write: A, B. Please begin answering:

### Answer A:

D

### Score A:

100

### Instruction B:

This is a multiple-choice question about agriculture. As an agricultural university student, you should answer according to the following rules: Select one or more correct options from A, B, C, and D, and respond accordingly. Note that your answer should include only the correct option letters, without extra spaces or other characters. When there are multiple correct options, separate them with a comma. For example: Correct answer: A, B Strictly follow this format and provide no additional output.  
Question:

### Answer B:

D, B, A

### Score B:

50

### Example 3

**Question:**

How is OsSPL14 regulated?

- A. By the OsmiR159-directed cleavage in vivo.
- B. By the OsmiR157-directed cleavage in vivo.
- C. By the OsmiR156-directed cleavage in vivo.
- D. By the OsmiR158-directed cleavage in vivo.

**Reference:**

C

**Instruction A:**

Below is a multiple-choice question on agricultural knowledge. Answer it from the perspective of an agricultural student.

Carefully read the question and select all correct options from A, B, C, and D.

Answer format: Only list the correct option letters, with multiple answers separated by commas. Standard format: Correct answer: A, B.

**Answer A:**

A

**Score A:**

0

**Instruction B:**

This is a multiple-choice question about agriculture. As an agricultural university student, you should answer according to the following rules: Select one or more correct options from A, B, C, and D, and respond accordingly. Note that your answer should include only the correct option letters, without extra spaces or other characters. When there are multiple correct options, separate them with a comma. For example: Correct answer: A, B Strictly follow this format and provide no additional output. Question:

**Answer B:**

A, C

**Score B:**

50

Example 4

**Question:**

What changes were observed in the cells of the Bg1-D mutant?

- A. Increased area of parenchyma cells
- B. Decreased number of parenchyma cells
- C. Down-regulation of genes associated with cell cycle and cell expansion
- D. Longer epidermis cells of both palea and lemma.

**Reference:**

D, A

**Instruction A:**

Here is a multiple-choice question in the field of agriculture. As an agricultural student, please select all the appropriate answers from A, B, C, and D.

Note: Your answer should be concise and clear, with multiple correct options separated by commas. Example format: Correct answer: A, B. Please follow this format when answering the question.

**Answer A:**

C

**Score A:**

0

**Instruction B:**

This is a multiple-choice question about agriculture. As an agricultural university student, you should answer according to the following rules: Select one or more correct options from A, B, C, and D, and respond accordingly. Note that your answer should include only the correct option letters, without extra spaces or other characters. When there are multiple correct options, separate them with a comma. For example: Correct answer: A, B Strictly follow this format and provide no additional output. Question:

**Answer B:**

A, C

**Score B:**

50

### Example 5

**Question:**

Where is OsSPL14 predominantly expressed?

- A. In the flower apex
- B. In the leaf apex
- C. In the root apex
- D. In the shoot apex.

**Reference:**

D

**Instruction A:**

Let's solve this multiple-choice question in the field of agriculture together. Your role is that of an agricultural university student, and you need to select all the correct answers from options A, B, C, and D. Remember, your answer should only include the corresponding letters. If there are multiple correct answers, separate them with a comma. Reference format: Correct answer: A, B

**Answer A:**

A, B, D

**Score A:**

50

**Instruction B:**

This is a multiple-choice question about agriculture. As an agricultural university student, you should answer according to the following rules: Select one or more correct options from A, B, C, and D, and respond accordingly. Note that your answer should include only the correct option letters, without extra spaces or other characters. When there are multiple correct options, separate them with a comma. For example: Correct answer: A, B Strictly follow this format and provide no additional output. Question:

**Answer B:**

A, B, D

**Score B:**

50

### Example 6

**Question:**

What growth changes were observed in one-week-old seedlings of the Bg1-D mutant compared with WT?

- A. Decreased length of leaf and root
- B. No significant changes in leaf and root length
- C. Decreased plant height
- D. Increased length of leaf and root.

**Reference:**

D

**Instruction A:**

As an agricultural student, carefully read the question and answer this multiple-choice question. Select all the correct answers from the given options A, B, C, and D. If multiple answers are correct, separate them with a comma and list only the letter options. Example format: Correct answer: A, B Please strictly follow this answer format.

**Answer A:**

A, C

**Score A:**

0

**Instruction B:**

This is a multiple-choice question about agriculture. As an agricultural university student, you should answer according to the following rules: Select one or more correct options from A, B, C, and D, and respond accordingly. Note that your answer should include only the correct option letters, without extra spaces or other characters. When there are multiple correct options, separate them with a comma. For example: Correct answer: A, B Strictly follow this format and provide no additional output. Question:

**Answer B:**

A, C

**Score B:**

0

### Example 7

**Question:**

OsSPL14 is localized to the nucleus (Fig. 1g), consistent with a role as a transcription factor. RNA in situ hybridization revealed that OsSPL14 was predominantly expressed in the shoot apex at both the vegetative (Fig. 2a) and reproductive (Fig. 2b) stages. It was also highly expressed in the promordia of primary and secondary branches (Fig. 2c). Expression patterns of OsSPL14 and OsmiR156 in various rice organs revealed by real-time PCR and miRNA gel blot analyses showed that OsSPL14 was highly expressed in the culm and shoot apex, which is complementary with the expression pattern of OsmiR156 in vivo (Fig. 2e,f). Consistently, overexpression of OsmiR156 resulted in a substantial decrease in OsSPL14 transcripts (Supplementary Fig. 7a), whereas the interruption of OsmiR156 (overexpression of MIM156) led to a marked increase in OsSPL14 transcripts (Supplementary Fig. 7b). These results indicated that OsSPL14 was regulated by the OsmiR156-directed cleavage in vivo. What function does OsSPL14 serve?

- A. It acts as a termination factor.
- B. It acts as a translation factor.
- C. It acts as a replication factor.
- D. It acts as a transcription factor.

**Reference:**

D

**Instruction A:**

Here is a multiple-choice test on agricultural knowledge. As an agricultural student, please read the question carefully and select all the correct answers from A, B, C, and D. Answer requirements: Only write the letter(s) of the correct option(s). If multiple answers are correct, separate them with a comma. For example, if A and B are correct, write: A, B. Keep your answer concise and do not add any explanations. The question is:

**Answer A:**

D

**Score A:**

100

**Instruction B:**

This is a multiple-choice question about agriculture. As an agricultural university student, you should answer according to the following rules: Select one or more correct options from A, B, C, and D, and respond accordingly. Note that your answer should include only the correct option letters, without extra spaces or other characters. When there are multiple correct options, separate them with a comma. For example: Correct answer: A, B Strictly follow this format and provide no additional output. Question:

**Answer B:**

D

**Score B:**

100

## Example 8

### Question:

Among them, Bg1-D has the most obviously increased grain size (Fig. 1A). Compared with WT (Nipponbare, *Oryza sativa* L. ssp. japonica), the 1,000-grain weight of Bg1-D has increased about 33.8%, with the grain length and grain width increased about 15.2% and 17.0%, respectively (Fig. 1B). Observation of spikelet hull by scanning electron microscopy showed that the epidermis cells of both palea and lemma in Bg1-D are much longer than that of the WT (Fig. S1). Careful examination of the hull cross-section before flowering revealed significant increases of both number and area of the parenchyma cells in the Bg1-D mutant (Fig. 1C and D). Consistent with this, a number of genes associated with cell cycle and cell expansion were up-regulated in the panicles of the Bg1-D mutant (Fig. S2). Despite the greatly enlarged grain size, Bg1-D also exhibits enhanced growth of other tissues at both the vegetative and reproductive stages (Fig. 1E and Fig. S3). One-week-old seedlings of the Bg1-D mutant have obviously increased length of leaf and root compared with WT (Fig. S3 A–D). When plants enter reproductive stage, the leaves of Bg1-D are longer and wider than those of the WT plants, and the Bg1-D also has increased plant height (Fig. 1E, Fig. S3E, and Table S1). In addition, the Bg1-D panicles are much larger than those of the WT (Fig. S3F and G). We overexpressed BG1 cDNA under the control of the rice ACTIN1 promoter in WT and found that the BG1-overexpressing plants (BG1-OE) showed obviously increased grain size and other phenotypes, such as enhanced plant height, longer leaves, and larger panicles, which resemble the Bg1-D mutant phenotypes (Fig. 2C and Fig. S6A and B). Importantly, severity of the phenotypes was apparently correlated with the BG1 expression level; that is, plants with higher BG1 expression have bigger grain size as well as increased 1,000-grain weight (Fig. 2D–G), demonstrating that activation of BG1 is the cause for the Bg1-D mutant phenotypes. What were the observed outcomes when the BG1 cDNA was overexpressed, regulated by the rice ACTIN1 promoter in WT?

- A. Enhanced plant height
- B. Larger panicles
- C. Decreased grain size
- D. Longer leaves.

### Reference:

A, B, D

### Instruction A:

Below is a multiple-choice test in the field of agriculture. As an agricultural student, identify all the correct answers from the four options: A, B, C, and D. Note: Your answer must be concise—only include the correct option letters. If there are multiple correct answers, separate them with a comma. For example, if A and B are correct, write: A, B. Please follow this format and do not add any explanations. The question is:

### Answer A:

A, D, B

### Score A:

100

### Instruction B:

This is a multiple-choice question about agriculture. As an agricultural university student, you should answer according to the following rules: Select one or more correct options from A, B, C, and D, and respond accordingly. Note that your answer should include only the correct option letters, without extra spaces or other characters. When there are multiple correct options, separate them with a comma. For example: Correct answer: A, B Strictly follow this format and provide no additional output. Question:

### Answer B:

A, D, B

### Score B:

100

Figure 21: Illustration of robustness evaluation for different prompt styles.

## H Detailed Performance Comparison and Analysis

To analyze the experiment results from the provided table, we can examine several aspects, including the performance of different models, the comparison between different sizes of the same series, the variation in performance across tasks, and the relationship between model size and performance.

### H.1 Model Comparison on Different LLMs

**Proprietary Models:** Among the proprietary models, GPT-4 consistently performs well across most of the tasks, often achieving top or second place in many breeding subcategories (C1, C3, C5, C6). The model's average score is notably high at 62.06, which is the highest among proprietary models.

**Open-Source Models:** Models such as DeepSeek-V3-671B and Qwen2-57B exhibit strong performance in certain categories (e.g., Qwen2-57B performs exceptionally well in C3 and C7).

### H.2 Comparison Between Different Model Sizes of the Same Model Series

**Qwen Models:** The different model sizes of the Qwen models show varying performance trends. While Qwen2-7B, a mid-tier model, significantly outperforms Qwen2-0.5B, it falls short of Qwen2-72B, the largest model in the series. However, contrary to expectations, increasing model size does not guarantee progressive performance gains. For instance, Qwen2.5-72B fails to demonstrate clear improvements over Qwen2.5-14B in many subcategories, suggesting that parameter count alone may not dictate performance in this domain. Notably, Qwen2.5-14B excels in several subcategories (e.g., C6, C8), achieving top-tier rankings, whereas Qwen2.5-72B shows inconsistent superiority—exhibiting only marginal gains in C3 and C6. This pattern implies diminishing returns for larger models in specific subcategories, highlighting the importance of training data and strategies beyond mere scale.

### H.3 Top Performers in Each Subcategory

- C1:** (Gene Basic Information Query: GPT-4, 59.59)
- C2:** (Gene Expression Pattern Query: DeepSeek-V3, 62.42)
- C3:** (Gene Product Cellular Localization Query: GPT-4, 76.32)
- C4:** (Gene Function Experimental Observation: DeepSeek-V3, 63.17)
- C5:** (Gene Product Regulation of Downstream Genes Analysis: GPT-4, 56.34)
- C6:** (Gene Function Prediction: GPT-4, 59.35)
- C7:** (Variety Breeding Process Query: DeepSeek-V3, 68.23)
- C8:** (Variety Agronomic Trait Query: DeepSeek-V3, 69.04)
- C9:** (Variety Cultivation and Technical Key Points Query: DeepSeek-V3, 66.46)
- C10:** (Variety Suitable Planting Area Recommendation: DeepSeek-V3, 68.48)

**Consistent Top Performers:** DeepSeek-V3-671B is a standout performer, securing top-3 positions in all ten subcategories and ranking first in C2 (62.42), C4 (63.17), C7 (68.23), C8 (69.04), C9 (66.46), and C10 (68.48). Its dominance in moderately hard to medium tasks (e.g., C7, C8, C10) underscores its strength in breeding-related applications.

**Other Consistent Top Performers:** GPT-4 also demonstrates exceptional consistency, appearing in the top-3 for nine out of ten subcategories (C1, C2, C3, C4, C5, C6, C8, C9, C10) and ranking first in C1 (59.59), C3 (76.32), C5 (56.34), and C6 (59.35). Its strong performance across both simpler tasks (e.g., C3) and more complex ones (e.g., C5) highlights its versatility and robustness. GLM-4-Plus frequently appears in the top-3 for six subcategories (C2, C4, C7, C8, C9, C10), with notable scores such as 65.02 in C7 and 64.17 in C8, making it a reliable performer, particularly in medium-difficulty tasks. Additionally, Qwen2-57B and Qwen2-72B occasionally appear in the top-3 (e.g., C3, C7 for Qwen2-57B; C5 for Qwen2-72B), but their consistency is less pronounced compared to GPT-4 and DeepSeek-V3-671B.

### H.4 Comparison of Models with Same Parameter Scale

Parameter Level	Best Model (Avg.)	Second Best (Avg.)	Advantaged Domains
7B	InternLM2.5-7B (53.51)	Qwen2.5-7B (48.45)	C3, C9
70B	Qwen2-72B (57.62)	Llama3.1-70B (54.30)	C2, C5
large-scale	DeepSeek-V3-671B (63.30)	GPT-4 (62.06)	C5, C10

H.5 Task-Specific Performance

Task Difficulty Spectrum (Based on average scores across all models):

Difficulty Level	Subcategory	Avg. Score
Most Difficult	C5	37
Moderately Hard	C1, C2, C4, C9	42-45
Medium	C6, C7, C8,C10	46-49
Easiest	C3	58

Model Performance on Specific Subcategories (C1–C10)

The performance of large language models (LLMs) across the ten breeding subcategories (C1–C10) reveals distinct patterns in task difficulty and model capability. By classifying tasks based on average model performance, we observe that larger models excel in complex tasks, while smaller models can achieve competitive results in simpler ones. The following analysis aligns with the provided difficulty classification, adjusted to reflect observed performance trends.

**Most Difficult Task (C5, Avg. Score 37):** Gene Product Regulation of Downstream Genes Analysis (C5) is the most complex, with a wide performance gap. Larger models like GPT-4 (56.34) and DeepSeek-V3-671B (55.23) lead, but smaller models, such as Qwen2-0.5B (27.62), GLM-4-Chat-9B (16.20), and PLLaMa-7B (11.66), score significantly lower, aligning with the 37 average for mid-to-low-tier models. This task demands advanced reasoning and knowledge integration, favoring larger models.

**Easiest Task (C3, Avg. Score 58):** Gene Product Cellular Localization Query (C3) is the least challenging, with top models like GPT-4 (76.32) and DeepSeek-V3-671B (74.81) achieving high scores. Even smaller models, such as InternLM2.5-7B (67.88) and Qwen2.5-7B (66.01), perform well, indicating that C3 requires less computational capacity and is accessible to models with fewer parameters. The average score of 58 reflects mid-tier model performance, though top performers significantly exceed this.

**Additional Insights:** Proprietary models (e.g., GPT-4) and large open-source models (e.g., DeepSeek-V3-671B) dominate across all categories, particularly in difficult tasks like C5. Domain-specific models, such as Aksara-v1-7B (24.26 in C5) and PLLaMa-13B (13.96 in C5), underperform, suggesting limited generalization. The consistent presence of GPT-4 and DeepSeek-V3-671B in top-3 rankings underscores the advantage of model scale in complex tasks, while smaller models remain viable for easier tasks.

In summary, the difficulty classification highlights that smaller models can perform competitively in easier tasks (e.g., C3), but larger models are essential for moderately hard to difficult tasks (e.g., C5, C1, C2). This emphasizes the importance of selecting appropriately scaled models to match task complexity in breeding-related applications.

H.6 Conclusion

**Model Choice:** DeepSeek-V3-671B emerges as the top performer overall and should be considered the best model for most tasks. However, models like GPT-4 and GLM-4-Plus demonstrate competitive performance and could be preferable for certain tasks.

**Task-Level Analysis:** For more difficult tasks such as C1 and C5, larger models like GPT-4 tend to excel, while Qwen2-57B performs strongly in easier categories such as C3.

**Insights from Scaling Law:** Conventional scaling laws suggest that larger models, such as Qwen2.5-72B, should consistently outperform smaller models like Qwen2.5-7B and Qwen2.5-14B due to their greater model sizes and capacity. However, the experimental results challenge this expectation, as Qwen2.5-7B often performs comparably to or even surpasses its larger counterparts in several subcategories (e.g., 53.06 in C9 vs. 53.05 for Qwen2.5-72B). This discrepancy indicates that model size alone does not guarantee superior performance. Instead, factors such as training data distribution, task-specific optimization, and the nature of the tasks play critical roles in determining model effectiveness.

The findings suggest that larger models may not be fully optimized for the specific task set evaluated here, potentially due to mismatches between their training data and the demands of breeding-related subcategories. For certain tasks, smaller models like Qwen2.5-7B may be better suited, particularly when tasks align closely with their training or require less complex reasoning. These results challenge the

universal applicability of scaling laws across all task types, highlighting the need for careful consideration of training data characteristics and task design when selecting model size for optimal performance.

### H.7 Assessment of the Generative Tasks using BERTScore

We have conducted an assessment of the generative tasks (QA-4, SUM-1, SUM-2, RC-3, RC-4) using BERTScore, and calculated the Pearson correlation coefficient between BERTScore and ROUGE-L. The results are as follows:

Table 8: Pearson correlation between BERTScore and ROUGE-L on generative tasks.

Task	Pearson Correlation (BERTScore vs. ROUGE-L)
QA-4	0.7937
SUM-1	0.6770
SUM-2	0.5811
RC-3	0.7515
RC-4	0.7338

The following table reports the BERTScore (%) for five generative task subsets in SeedBench.

Table 9: BERTScore (%) across models on five generative tasks.

Model	QA-4	SUM-1	SUM-2	RC-3	RC-4
Claude-3.5-Sonnet	48.43	50.74	51.83	43.28	48.47
Gemini-1.5-Pro	72.39	74.89	73.89	95.59	79.48
Gemini-2.0-Flash	64.60	74.84	70.85	61.71	72.00
GLM-4-Plus	79.70	83.42	80.19	95.03	84.34
GPT-4o mini	81.10	83.50	82.79	93.98	86.01
GPT-4	79.66	87.45	85.95	96.02	88.27
OpenAI o1-mini	77.36	75.18	69.91	94.61	81.97
DeepSeek-V3	81.22	83.04	81.55	94.25	85.89
GLM-4-Chat-9B	56.10	69.45	74.72	51.17	60.93
InternLM2-7B	57.49	72.76	74.04	53.83	61.34
InternLM2.5-7B	76.97	84.07	81.63	91.01	85.72
Llama3.1-8B	74.91	82.11	77.87	83.28	81.18
Llama3.1-70B	77.97	83.84	82.62	91.12	82.98
Llama3.3-70B	76.35	81.89	79.08	87.91	82.52
Mistral-v0.3-7B	75.13	81.98	84.54	75.94	83.66
Qwen2-0.5B	75.27	78.37	76.46	64.88	79.64
Qwen2-7B	77.21	78.71	78.85	84.85	82.60
Qwen2-57B	80.31	82.82	83.93	92.00	85.83
Qwen2-72B	78.75	80.23	85.11	94.51	84.53
Qwen2.5-7B	77.66	81.19	82.19	82.82	83.35
Qwen2.5-14B	75.30	75.38	62.66	83.46	82.80
Qwen2.5-72B	79.50	84.38	83.50	93.16	83.14
QwQ-32B	70.16	70.43	69.57	65.62	73.06
Aksara-v1-7B	71.48	77.60	80.33	70.86	79.36
PLLaMa-7B	69.53	64.44	63.62	62.41	69.60
PLLaMa-13B	61.73	62.46	59.66	56.95	66.64

## H.8 Additional Quantitative Results

Models	Breeding Subcategories										Average
	C1	C2	C3	C4	C5	C6	C7	C8	C9	C10	
Proprietary LLMs											
Claude-3.5-Sonnet	48.77	57.72	66.02	57.54	47.82	49.36	57.47	60.11	58.06	58.89	55.45
Gemini-1.5-Pro	47.00	59.55	62.42	59.56	43.11	49.55	53.41	56.18	52.51	53.71	53.58
Gemini-2.0-Flash	33.67	27.37	53.04	32.07	25.87	44.41	33.57	36.77	31.78	31.70	34.24
GLM-4-Plus	52.72	59.62	70.62	60.11	50.60	56.75	65.02	64.17	61.70	62.90	59.61
GPT-4	59.59	60.55	76.32	61.16	56.34	59.35	63.67	64.74	60.65	67.66	62.06
GPT-4o mini	54.24	56.64	72.11	59.28	53.00	57.88	58.38	61.75	57.50	62.38	58.40
OpenAI o1-mini	49.16	55.58	59.37	54.77	44.43	50.73	54.57	55.36	54.91	54.19	53.25
Open-Source LLMs											
DeepSeek-V3-671B	56.03	62.42	74.81	63.17	55.23	58.84	68.23	69.04	66.46	68.48	63.30
GLM-4-Chat-9B	23.28	21.31	39.97	26.13	16.20	34.15	26.63	29.60	25.60	26.68	26.55
InternLM2-7B	(±0.11)	(±0.12)	(±0.43)	(±0.12)	(±0.01)	(±0.12)	(±0.08)	(±0.01)	(±0.02)	(±0.05)	(±0.03)
	27.55	21.14	39.64	28.57	15.16	36.12	28.74	30.80	27.32	29.22	28.71
InternLM2.5-7B	(±0.01)	(±0.01)	(±0.00)	(±0.00)	(±0.00)	(±0.03)	(±0.02)	(±0.04)	(±0.01)	(±0.02)	(±0.01)
	51.71	55.75	67.88	50.48	44.14	56.73	51.28	54.91	52.46	56.24	53.51
Llama3.1-8B	(±0.01)	(±0.03)	(±0.02)	(±0.14)	(±0.01)	(±0.01)	(±0.02)	(±0.03)	(±0.01)	(±0.01)	(±0.03)
	43.89	31.21	42.53	40.68	38.47	43.80	42.87	51.62	41.88	40.91	42.23
Llama3.1-70B	(±0.35)	(±1.07)	(±0.30)	(±0.16)	(±0.38)	(±0.03)	(±0.05)	(±0.05)	(±0.51)	(±0.14)	(±0.23)
	48.72	55.41	64.77	53.67	46.73	54.08	56.94	57.72	55.31	57.56	54.30
Llama3.3-70B	(±0.07)	(±0.01)	(±0.00)	(±0.29)	(±0.00)	(±0.01)	(±0.12)	(±0.09)	(±0.11)	(±0.56)	(±0.13)
	45.32	47.15	60.62	49.76	40.90	54.30	52.79	54.61	49.98	55.05	50.53
Mistral-v0.3-7B	(±0.00)	(±0.06)	(±0.00)	(±0.00)	(±0.00)	(±0.09)	(±0.06)	(±0.00)	(±0.02)	(±0.00)	(±0.01)
	42.61	38.28	57.02	40.41	29.97	44.22	36.31	43.98	39.92	43.51	41.59
Qwen2-0.5B	(±0.32)	(±0.19)	(±0.14)	(±0.15)	(±0.21)	(±0.06)	(±0.39)	(±0.23)	(±0.10)	(±0.05)	(±0.06)
	32.84	25.15	40.19	28.20	27.62	37.22	33.81	33.63	28.25	31.67	31.44
Qwen2-7B	(±0.26)	(±0.64)	(±0.64)	(±0.24)	(±0.49)	(±0.19)	(±0.58)	(±0.21)	(±0.33)	(±0.22)	(±0.13)
	44.21	40.41	63.00	47.36	35.37	52.30	45.61	48.73	44.88	46.89	46.51
Qwen2-57B	(±0.03)	(±0.10)	(±0.11)	(±0.10)	(±0.12)	(±0.27)	(±0.06)	(±0.14)	(±0.13)	(±0.17)	(±0.02)
	53.67	49.81	74.30	58.38	39.34	54.71	63.89	59.57	59.22	60.08	57.20
Qwen2-72B	(±0.21)	(±0.05)	(±0.37)	(±0.04)	(±0.40)	(±0.16)	(±0.37)	(±0.04)	(±0.01)	(±0.07)	(±0.01)
	51.16	58.10	74.07	59.72	51.58	57.76	58.85	61.63	56.69	59.11	57.62
Qwen2.5-7B	(±1.70)	(±4.07)	(±0.04)	(±3.44)	(±0.42)	(±0.75)	(±2.35)	(±4.89)	(±2.16)	(±3.33)	(±2.60)
	45.16	39.50	66.01	44.61	35.72	50.00	53.60	53.31	53.06	51.05	48.45
Qwen2.5-14B	(±0.43)	(±0.19)	(±0.36)	(±0.19)	(±0.51)	(±0.24)	(±0.29)	(±0.21)	(±0.20)	(±0.31)	(±0.10)
	50.91	50.73	68.62	52.15	47.14	54.54	57.02	62.05	54.37	54.15	54.21
Qwen2.5-72B	(±0.00)	(±0.00)	(±0.00)	(±0.06)	(±0.00)	(±0.00)	(±0.03)	(±0.00)	(±0.12)	(±0.00)	(±0.02)
	46.86	47.41	70.99	51.89	46.17	57.60	55.35	56.31	53.05	54.75	52.63
QwQ-32B	(±0.01)	(±0.01)	(±0.00)	(±0.00)	(±0.00)	(±0.13)	(±0.00)	(±0.01)	(±0.01)	(±0.01)	(±0.02)
	32.24	21.06	47.11	29.14	28.56	39.68	38.17	39.56	34.70	34.52	33.55
	(±0.00)	(±0.00)	(±0.00)	(±0.00)	(±0.00)	(±0.00)	(±0.00)	(±0.00)	(±0.00)	(±0.00)	(±0.00)
Domain Specific LLMs											
Aksara-v1-7B	36.72	36.69	48.32	35.41	24.26	36.83	31.17	34.64	31.15	34.14	35.04
	(±0.18)	(±0.10)	(±0.27)	(±0.06)	(±0.02)	(±0.22)	(±0.03)	(±0.02)	(±0.01)	(±0.05)	(±0.05)
PLLaMa-7B	17.85	13.69	17.99	16.81	11.66	21.67	14.34	17.36	12.39	16.11	16.46
	(±0.03)	(±0.07)	(±0.44)	(±0.14)	(±0.20)	(±0.34)	(±0.21)	(±0.03)	(±0.17)	(±0.22)	(±0.01)
PLLaMa-13B	15.10	14.18	28.41	18.83	13.96	23.28	18.53	17.37	14.15	18.51	17.57
	(±0.04)	(±0.03)	(±0.36)	(±0.03)	(±0.01)	(±0.07)	(±0.12)	(±0.21)	(±0.08)	(±0.03)	(±0.03)

Table 10: Evaluation of 26 LLMs on SeedBench. Performance is stratified by breeding subcategories, with open-source/domain-specific models evaluated through 3 repeated trials (mean scores reported). The scores represent averages across three different metrics for 11 task types. The columns delineate ten subcategories in breeding: (C1) Gene Basic Information Query, (C2) Gene Expression Pattern Query, (C3) Gene Product Cellular Localization Query, (C4) Gene Function Experimental Observation, (C5) Gene Product Regulation of Downstream Genes Analysis, (C6) Gene Function Prediction, (C7) Variety Breeding Process Query, (C8) Variety Agronomic Trait Query, (C9) Variety Cultivation and Technical Key Points Query, (C10) Variety Suitable Planting Area Recommendation. Top-3 performers per column are highlighted in red.

Models	Question Type											Average
	QA-1	QA-2	QA-3	QA-4	SUM-1	SUM-2	RC-1	RC-2	RC-3	RC-4	RC-5	
Proprietary LLMs												
Claude-3.5-Sonnet	57.50	74.68	25.84	21.82	33.97	44.94	98.23	97.67	85.97	44.58	88.17	61.22
Gemini-1.5-Pro	56.50	73.26	27.37	16.96	22.99	28.88	99.12	97.01	86.75	43.18	84.95	57.91
Gemini-2.0-Flash	62.00	58.59	1.83	11.04	31.01	27.76	97.35	65.25	7.37	25.74	95.34	43.93
GLM-4-Plus	64.50	75.40	25.82	32.51	48.53	51.53	99.12	97.02	84.39	49.31	85.30	64.85
GPT-4o mini	57.50	72.33	17.88	44.47	49.51	60.66	97.35	95.38	85.30	57.73	84.23	65.67
GPT-4	60.50	73.87	21.35	36.07	58.73	62.89	100.00	96.44	87.86	62.29	86.74	67.88
OpenAI o1-mini	57.50	73.81	22.25	27.93	38.40	37.80	100.00	96.17	83.46	36.40	82.80	59.68
DeepSeek-V3	72.50	79.84	29.29	40.63	48.06	54.67	100.00	97.22	87.89	55.19	86.74	68.37
Open-Source LLMs												
GLM-4-Chat-9B	53.00 (±0.00)	62.41 (±0.10)	0.95 (±0.00)	1.57 (±0.01)	5.69 (±0.03)	13.38 (±0.02)	98.53 (±0.51)	68.51 (±0.29)	6.26 (±0.01)	2.29 (±0.01)	93.55 (±0.36)	36.92
InternLM2-7B	56.50 (±0.00)	66.93 (±0.17)	1.85 (±0.00)	2.14 (±0.00)	8.66 (±0.07)	13.88 (±0.03)	98.23 (±0.00)	85.88 (±0.00)	14.31 (±0.02)	6.51 (±0.02)	96.77 (±0.00)	41.06
InternLM2.5-7B	58.50 (±0.00)	65.86 (±0.00)	12.99 (±0.05)	32.69 (±0.07)	50.98 (±0.10)	53.64 (±0.02)	99.12 (±0.00)	94.48 (±0.00)	73.39 (±0.00)	55.41 (±0.01)	70.37 (±0.21)	60.67
Llama3.1-8B	48.00 (±0.00)	64.33 (±0.30)	10.60 (±0.24)	24.57 (±0.12)	48.11 (±0.14)	42.62 (±0.09)	94.69 (±0.00)	81.05 (±0.09)	54.31 (±0.76)	45.63 (±0.47)	74.67 (±0.21)	53.51
Llama3.1-70B	56.00 (±0.71)	73.73 (±0.07)	19.53 (±0.09)	34.65 (±0.00)	52.09 (±0.00)	54.83 (±0.02)	99.12 (±0.00)	96.62 (±0.00)	74.76 (±0.98)	51.28 (±0.01)	83.15 (±0.01)	63.25
Llama3.3-70B	58.50 (±0.00)	71.50 (±0.00)	18.65 (±0.00)	26.55 (±0.01)	47.79 (±0.13)	53.03 (±0.01)	99.12 (±0.00)	93.88 (±0.00)	64.92 (±0.09)	49.96 (±0.02)	80.29 (±0.00)	60.38
Mistral-v0.3-7B	39.50 (±0.00)	58.14 (±0.19)	5.91 (±0.18)	30.01 (±0.42)	45.02 (±0.05)	62.03 (±0.03)	83.19 (±0.00)	76.43 (±0.00)	38.51 (±0.22)	50.31 (±0.53)	73.12 (±0.00)	51.11
Qwen2-0.5B	40.50 (±0.00)	62.20 (±0.29)	2.82 (±0.07)	31.21 (±0.47)	44.72 (±0.20)	53.73 (±0.34)	55.46 (±1.35)	65.89 (±0.61)	15.98 (±0.08)	42.04 (±0.76)	50.18 (±0.72)	42.25
Qwen2-7B	57.00 (±0.00)	69.44 (±0.09)	12.72 (±0.17)	32.42 (±0.08)	31.18 (±0.00)	40.30 (±0.08)	97.35 (±0.00)	82.36 (±0.17)	57.83 (±0.20)	46.96 (±0.28)	84.59 (±0.00)	55.65
Qwen2-57B	56.00 (±0.00)	74.79 (±0.03)	21.24 (±0.11)	39.56 (±0.23)	46.04 (±0.03)	60.27 (±0.64)	99.12 (±0.00)	95.95 (±0.00)	76.55 (±0.09)	56.66 (±0.25)	80.83 (±0.25)	64.27
Qwen2-72B	59.50 (±0.00)	75.98 (±0.00)	19.55 (±0.00)	31.62 (±0.00)	31.08 (±0.00)	63.09 (±0.03)	99.12 (±0.00)	94.24 (±0.00)	72.20 (±0.00)	51.58 (±0.00)	89.96 (±0.00)	62.54
Qwen2.5-7B	57.00 (±1.00)	71.19 (±0.24)	17.46 (±0.10)	32.47 (±0.09)	43.96 (±0.01)	54.70 (±0.07)	96.46 (±0.00)	88.33 (±0.07)	54.47 (±0.00)	47.97 (±0.02)	79.21 (±0.00)	58.47
Qwen2.5-14B	57.00 (±0.00)	72.49 (±0.00)	17.43 (±0.00)	22.66 (±0.00)	54.88 (±0.08)	56.20 (±0.00)	99.12 (±0.00)	93.36 (±0.00)	67.72 (±0.25)	51.57 (±0.00)	83.51 (±0.00)	61.45
Qwen2.5-72B	70.50 (±0.00)	73.71 (±0.00)	17.86 (±0.00)	29.84 (±0.01)	51.33 (±0.00)	59.52 (±0.03)	100.00 (±0.00)	87.79 (±0.00)	57.07 (±0.00)	48.41 (±0.00)	83.51 (±0.00)	61.77
QwQ-32B	61.50 (±0.00)	58.81 (±0.00)	3.00 (±0.00)	14.54 (±0.00)	15.60 (±0.00)	26.93 (±0.00)	91.15 (±0.00)	64.17 (±0.00)	19.23 (±0.00)	19.09 (±0.00)	91.04 (±0.00)	42.28
Domain Specific LLMs												
Aksara-v1-7B	34.167 (±0.289)	53.4 (±0.139)	4.28 (±0.061)	17.263 (±0.025)	26.503 (±0.006)	38.32 (±0.00)	87.61 (±0.00)	72.89 (±0.00)	30.037 (±0.075)	35.397 (±0.396)	66.67 (±0.00)	42.41
PLLaMa-7B	6.5 (±0.5)	35.47 (±0.485)	2.613 (±0.032)	20.223 (±0.326)	13.753 (±0.126)	20.577 (±0.046)	6.19 (±0.00)	35.03 (±0.00)	12.76 (±0.00)	24.583 (±0.167)	4.78 (±0.208)	16.59
PLLaMa-13B	24.167 (±0.289)	50.433 (±0.133)	1.013 (±0.04)	18.17 (±0.171)	13.37 (±0.00)	14.84 (±0.00)	38.94 (±0.00)	57.37 (±0.00)	6.95 (±0.00)	18.007 (±0.243)	40.98 (±0.416)	25.84

Table 11: Evaluation results (zero-shot) on different question type. Top-3 per column are highlighted.

Models	Question Type											Average
	QA-1	QA-2	QA-3	QA-4	SUM-1	SUM-2	RC-1	RC-2	RC-3	RC-4	RC-5	
Proprietary LLMs												
Claude-3.5-Sonnet	53.50	75.54	27.92	23.99	32.36	42.88	98.23	97.67	85.68	45.40	86.02	60.84
Gemini-1.5-Pro	55.50	73.49	27.13	27.03	28.98	35.65	99.12	97.45	86.98	46.08	86.74	60.38
Gemini-2.0-Flash	64.50	59.30	1.73	12.99	31.99	36.43	97.35	65.91	6.25	26.63	93.19	45.12
GLM-4-Plus	67.00	76.46	26.47	46.13	51.35	52.73	99.12	96.62	85.99	59.14	80.29	67.39
GPT-4o mini	55.00	71.89	19.10	41.59	50.21	57.59	97.35	95.00	86.00	53.76	82.08	64.51
GPT-4	57.50	73.24	23.43	46.19	63.93	68.07	100.00	96.86	88.97	64.53	83.15	69.63
OpenAI o1-mini	61.50	74.78	19.82	35.75	32.89	25.91	99.12	96.84	83.31	47.12	80.65	59.79
DeepSeek-V3	71.50	79.16	32.28	46.59	50.16	58.52	100.00	97.00	86.85	60.59	85.30	69.81
Open-Source LLMs												
GLM-4-Chat-9B	55.50	59.31	0.51	1.24	4.46	11.12	97.35	64.35	7.67	1.92	91.16	35.87
	(±0.00)	(±0.00)	(±0.00)	(±0.01)	(±0.01)	(±0.00)	(±0.00)	(±0.00)	(±0.00)	(±0.00)	(±0.41)	
InternLM2-7B	50.50	59.30	0.16	1.55	7.82	12.76	94.69	75.12	6.35	2.98	88.17	36.31
	(±0.00)	(±0.03)	(±0.00)	(±0.01)	(±0.00)	(±0.00)	(±0.00)	(±0.00)	(±0.00)	(±0.01)	(±0.97)	
InternLM2.5-7B	58.50	69.42	14.74	35.05	53.63	56.07	99.12	95.18	75.04	57.03	68.22	62.00
	(±0.00)	(±0.00)	(±0.00)	(±0.12)	(±0.00)	(±0.00)	(±0.00)	(±0.00)	(±0.00)	(±0.11)	(±0.06)	
Llama3.1-8B	35.50	49.10	7.97	32.90	49.22	46.62	86.43	75.39	54.04	41.27	68.10	49.68
	(±0.00)	(±0.00)	(±0.01)	(±0.63)	(±0.07)	(±0.01)	(±0.51)	(±0.38)	(±0.00)	(±0.01)	(±0.62)	
Llama3.1-70B	55.00	70.01	19.60	33.17	52.24	59.04	99.12	96.02	74.49	51.42	80.65	62.80
	(±0.00)	(±0.06)	(±0.00)	(±0.07)	(±0.00)	(±0.06)	(±0.00)	(±0.00)	(±0.33)	(±0.00)	(±0.33)	
Llama3.3-70B	54.75	73.55	17.35	30.61	50.09	45.12	99.12	93.52	63.22	45.62	79.21	59.29
	(±0.35)	(±0.13)	(±0.00)	(±0.06)	(±0.00)	(±0.01)	(±0.00)	(±0.00)	(±0.00)	(±0.06)	(±0.22)	
Mistral-v0.3-7B	39.50	60.85	3.82	31.09	45.74	62.27	79.06	75.79	35.65	48.81	74.55	50.65
	(±0.50)	(±0.13)	(±0.01)	(±0.12)	(±0.00)	(±0.00)	(±0.51)	(±0.00)	(±0.30)	(±0.10)	(±0.75)	
Qwen2-0.5B	44.50	62.94	2.28	33.36	36.05	38.12	67.26	68.93	12.10	40.83	25.09	39.22
	(±0.00)	(±0.31)	(±0.25)	(±0.65)	(±0.38)	(±0.44)	(±0.00)	(±0.42)	(±0.98)	(±0.83)	(±0.46)	
Qwen2-7B	52.17	65.62	14.13	34.50	43.35	52.33	88.50	78.53	52.03	48.77	85.42	55.94
	(±0.58)	(±0.79)	(±0.02)	(±0.18)	(±0.00)	(±0.00)	(±0.00)	(±0.08)	(±0.24)	(±0.01)	(±0.31)	
Qwen2-57B	58.00	76.42	21.17	43.49	49.60	60.53	97.35	95.28	76.86	56.87	77.42	64.82
	(±0.00)	(±0.00)	(±0.09)	(±0.36)	(±0.01)	(±0.12)	(±0.00)	(±0.00)	(±0.16)	(±0.16)	(±0.63)	
Qwen2-72B	65.00	77.73	24.51	41.70	54.36	66.48	99.12	96.13	78.54	57.85	84.41	67.80
	(±0.00)	(±0.13)	(±2.57)	(±0.96)	(±0.07)	(±0.01)	(±0.00)	(±1.57)	(±7.40)	(±0.69)	(±0.57)	
Qwen2.5-7B	57.667	70.32	16.867	38.27	43.897	59.597	92.92	80.093	52.003	52.46	75.987	58.19
	(±0.00)	(±1.11)	(±0.62)	(±0.29)	(±0.18)	(±0.40)	(±0.00)	(±0.00)	(±0.30)	(±0.33)	(±0.39)	
Qwen2.5-14B	63.5	72.57	18.79	38.4	56.24	62.14	99.12	94.39	69.877	56.83	82.08	64.90
	(±0.00)	(±0.00)	(±0.00)	(±0.00)	(±0.00)	(±0.00)	(±0.00)	(±0.00)	(±0.335)	(±0.00)	(±0.00)	
Qwen2.5-72B	67.50	74.79	21.39	40.425	55.86	62.01	99.12	89.11	55.995	51.38	82.8	63.67
	(±0.00)	(±0.00)	(±0.00)	(±0.078)	(±0.00)	(±0.00)	(±0.00)	(±0.00)	(±0.007)	(±0.028)	(±0.00)	
QwQ-32B	63.50	60.26	6.81	25.35	13.12	20.35	93.81	65.07	22.72	27.84	75.99	43.17
	(±0.00)	(±0.00)	(±0.00)	(±0.00)	(±0.00)	(±0.00)	(±0.00)	(±0.00)	(±0.00)	(±0.00)	(±0.00)	
Domain Specific LLMs												
Aksara-v1-7B	41.667	55.627	2.653	23.35	38.7	57.14	84.07	70.88	27.037	40.087	71.33	50.65
	(±0.289)	(±0.182)	(±0.072)	(±0.157)	(±0.00)	(±0.00)	(±0.00)	(±0.00)	(±0.167)	(±0.389)	(±0.00)	
PLLaMa-7B	11.167	37.22	6.437	25.203	30.42	27.71	5.31	28.38	20.27	24.723	29.75	22.42
	(±0.577)	(±0.859)	(±0.202)	(±0.117)	(±0.00)	(±0.00)	(±0.00)	(±0.00)	(±0.208)	(±0.067)	(±0.00)	
PLLaMa-13B	24.0	52.697	0.507	7.167	12.34	13.73	30.09	50.51	3.81	15.55	32.26	22.06
	(±0.00)	(±0.029)	(±0.006)	(±0.086)	(±0.00)	(±0.00)	(±0.00)	(±0.00)	(±0.035)	(±0.00)	(±0.00)	

Table 12: Evaluation results (one-shot) on different question type. Top-3 per column are highlighted.

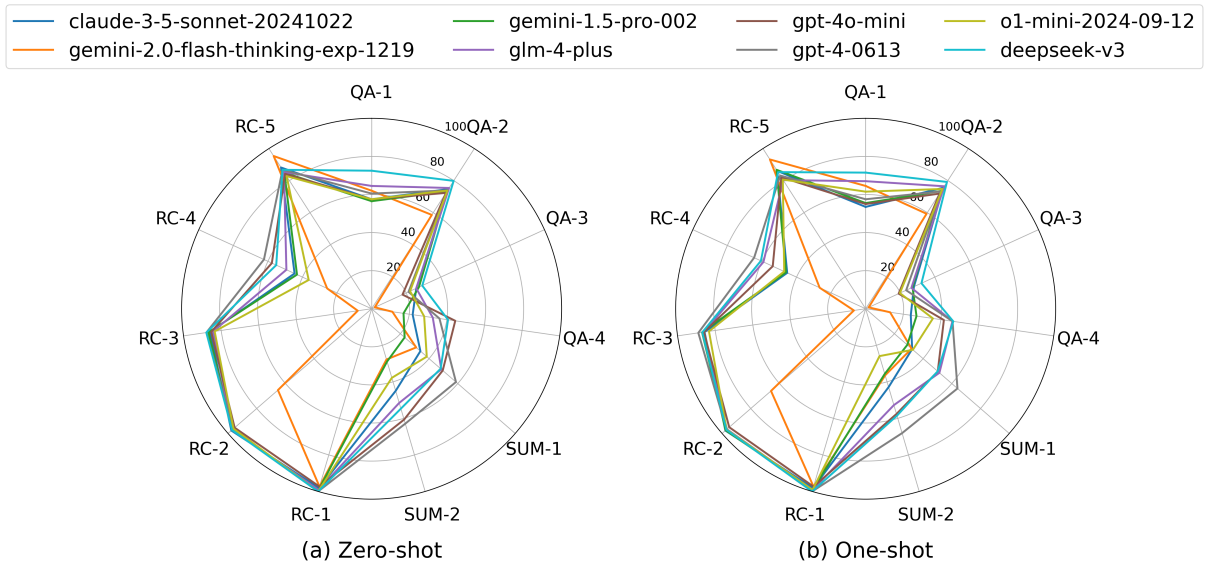


Figure 22: Evaluation of Proprietary LLMs on SeedBench. Performance is stratified by task-type.

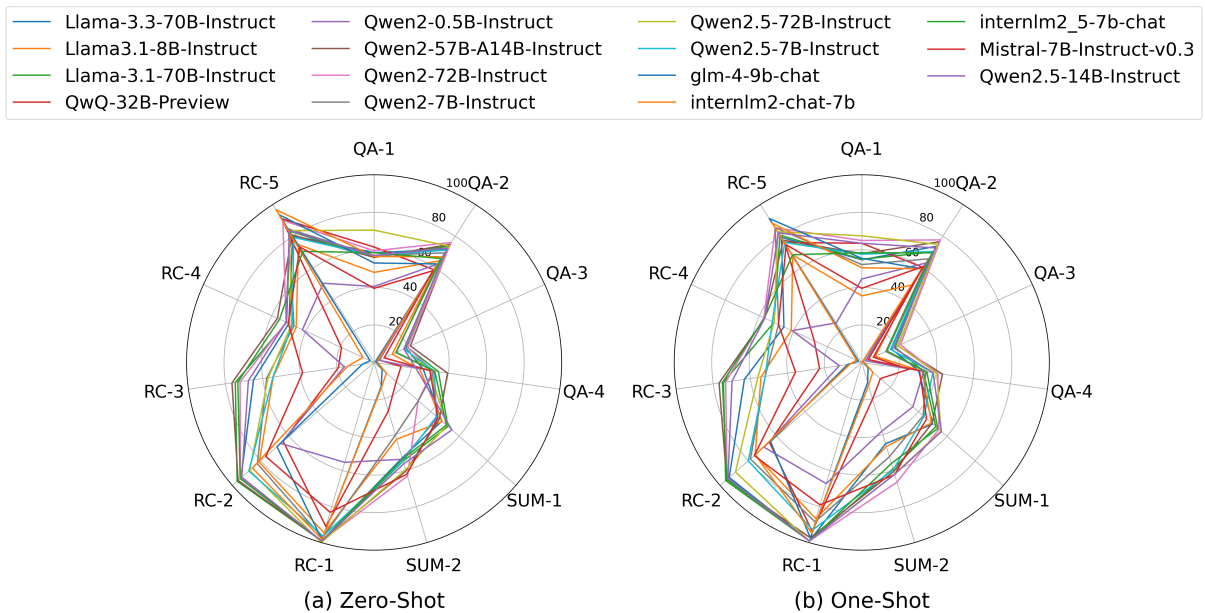


Figure 23: Evaluation of Open-Source LLMs on SeedBench. Performance is stratified by task-type.