

---

# Supplementary Materials

## AGMMU: A Comprehensive Agricultural Multimodal Understanding Benchmark

---

Aruna Gauba<sup>1,2,5\*</sup> Irene Pi<sup>1,3,5\*</sup> Yunze Man<sup>1,4,5†</sup> Ziqi Pang<sup>1,4,5†</sup>

Vikram S. Adve<sup>1,4,5</sup> Yu-Xiong Wang<sup>1,4,5</sup>

<sup>1</sup>University of Illinois Urbana-Champaign <sup>2</sup>Rice University <sup>3</sup>Carnegie Mellon University

<sup>4</sup>AIFARMS <sup>5</sup>Center for Digital Agriculture at UIUC

\*† Equal Contribution † Project Lead

1	<b>Contents</b>	
2	<b>A Evaluated Large Vision Language Models</b>	<b>1</b>
3	<b>B Dataset Curation Details</b>	<b>3</b>
4	B.1 Stage 1: Question Categorization . . . . .	3
5	B.2 Stage 2: Information Extraction . . . . .	6
6	B.3 Stage 3: Question Generation . . . . .	9
7	B.4 Final Stage: Human Verification . . . . .	10
8	<b>C More Evaluation, Implementation, and Design Choices</b>	<b>10</b>
9	<b>D More Dataset Visualization</b>	<b>11</b>
10	<b>E Limitations and Future Work</b>	<b>12</b>
11	<b>F Societal Impact</b>	<b>13</b>
12	<b>A Evaluated Large Vision Language Models</b>	

13 Our evaluation and analysis are conducted mainly on the group of models listed in Table 2 in the  
14 main paper. We have chosen models such that they cover most of the popular and best-performing  
15 methods used by recent multimodal understanding work. In this part, we discuss all the models we  
16 have used in our experiments and explain their evaluation details, the public checkpoints we have  
17 chosen, and display the prompts we used to adapt the model to our datasets.

18 During evaluation, we chose to follow the standard prompt provided by the authors whenever possi-  
19 ble for multiple-choice and short-answer questions. When the prompt is not provided for the model,  
20 we select a custom prompt that is created through several iterations of prompt engineering to select  
21 the one that produces the most effective results. The images are always included as the prefix.

22 **Proprietary Models.** We used three proprietary models in our evaluation: GPT-o4-mini [1], Gem-  
23 ini 1.5 Pro [9], and Claude 3 Haiku [10]. Below we note the model API version used for evaluation.

- 24 • GPT-o4-mini: May 13-15, 2025.
- 25 • Gemini 1.5 Pro: November 1-13, 2024.
- 26 • Claude 3 Haiku: November 13-14, 2024.

27 **Cambrian-1** [12]. Cambrian-1 is a recent state-of-the-art model that excels at visual-centric tasks.  
28 This model explores combinations of vision encoders, text and image integration techniques, and  
29 instruction tuning strategies. We use the official implementation and checkpoint<sup>1</sup> with a LLaMA3-  
30 8B-Instruct LLM backbone model in our evaluation.

31 **InternVL2** [11]. InternVL scales up the vision foundation model while aligning it with the back-  
32 bone LLM, and is trained on web-scale image-text data to achieve strong performance across a vari-  
33 ety of vision-centric tasks. We use the official implementation and checkpoint<sup>2</sup> with the InternViT-  
34 300M-448px vision backbone and Internlm2.5-7B-chat language backbone in our evaluation.

35 **LLaMA-3.2** [4]. LLaMA-3.2 is the first collection of multimodal large language model from the  
36 LLaMA family that was previously text-only. The integration of vision involves utilizing cross-  
37 attention layers and a pre-trained vision encoder that feeds directly into the text-processor. The  
38 model follows a commonly used training recipe that includes pretraining on noisy image-text pairs  
39 and then high-quality knowledge enhanced pairs. Notably, the language-model parameters were  
40 frozen during the training of alignment of image and text to retain strong text-only capabilities. We  
41 use the official implementation and checkpoint<sup>3</sup> that uses a LLaMA-3.1 text-only language backbone  
42 in our evaluation. When evaluating the model, we choose to use a custom prompt since no standard  
43 prompt is provided.

44 **LLaVA-NeXT** [8]. LLaVA-NeXT expands on LLaVA by using the same instruction tuning method  
45 to give the model the ability to process and reason about multi-images, multi-grames, and multi-  
46 views. We use the official implementation and checkpoint<sup>4</sup> with LLaMA-3-8B Instruct as the lan-  
47 guage backbone in our evaluation.

48 **LLaVA-OneVision** [5]. LLaVA-OneVision builds on LLaVA-NeXT with the capability to analyze  
49 single images, multi-images, and video scenarios. Most impressively, it allows for video under-  
50 standing through task transfer from images but this is not explored in our evaluation. We use the  
51 official implementation and checkpoint<sup>5</sup> that uses a base architecture consisting of SigLIP-SO400M-  
52 Patch14-384 and Qwen2-7b in our evaluation.

53 **LLaVA-1.5-7B / LLaVA-1.5-13B** [7]. LLaVA introduces the idea of instruction tuning a multi-  
54 modal model with GPT-4 generated instruction-following data for associated images. This gives it  
55 the ability to achieve impressive abilities to act as an instruction-following general agent. We use the  
56 official implementation and checkpoints<sup>6,7</sup> with a CLIP ViT-L/14 vision backbone and Vicuna1.5-7B  
57 / Vicuna1.5-13B in our evaluation.

58 **Qwen-VL-7B** [2]. Qwen-VL is a large vision language model that has the ability to perform vari-  
59 ous vision-language tasks including image captioning, visual grounding and more, not only limited  
60 to question answering. This model is multi-lingual in Chinese and English and was pre-trained using  
61 an interleaved image-text technique. We use the official implementation and checkpoint<sup>8</sup> that uses  
62 Qwen-7B as the language backbone and CLIP ViT bigG/14 as the vision encoder in our evaluation.

63 **VILA1.5-13B** [6]. VILA is trained using an enhanced pre-training method that involves inter-  
64 leaved visual language data. Additionally, during the supervised fine-tuning stage, the data includes

---

<sup>1</sup><https://github.com/cambrian-mlm/cambrian>

<sup>2</sup><https://huggingface.co/OpenGVLab/InternVL2-8B>

<sup>3</sup><https://huggingface.co/meta-llama/Llama-3.2-11B-Vision>

<sup>4</sup><https://huggingface.co/llava-hf/llama3-llava-next-8b-hf>

<sup>5</sup><https://huggingface.co/llava-hf/llava-onevision-qwen2-7b-ov-hf>

<sup>6</sup><https://huggingface.co/llava-hf/llava-1.5-7b-hf>

<sup>7</sup><https://huggingface.co/llava-hf/llava-1.5-13b-hf>

<sup>8</sup><https://huggingface.co/Qwen/Qwen-VL>

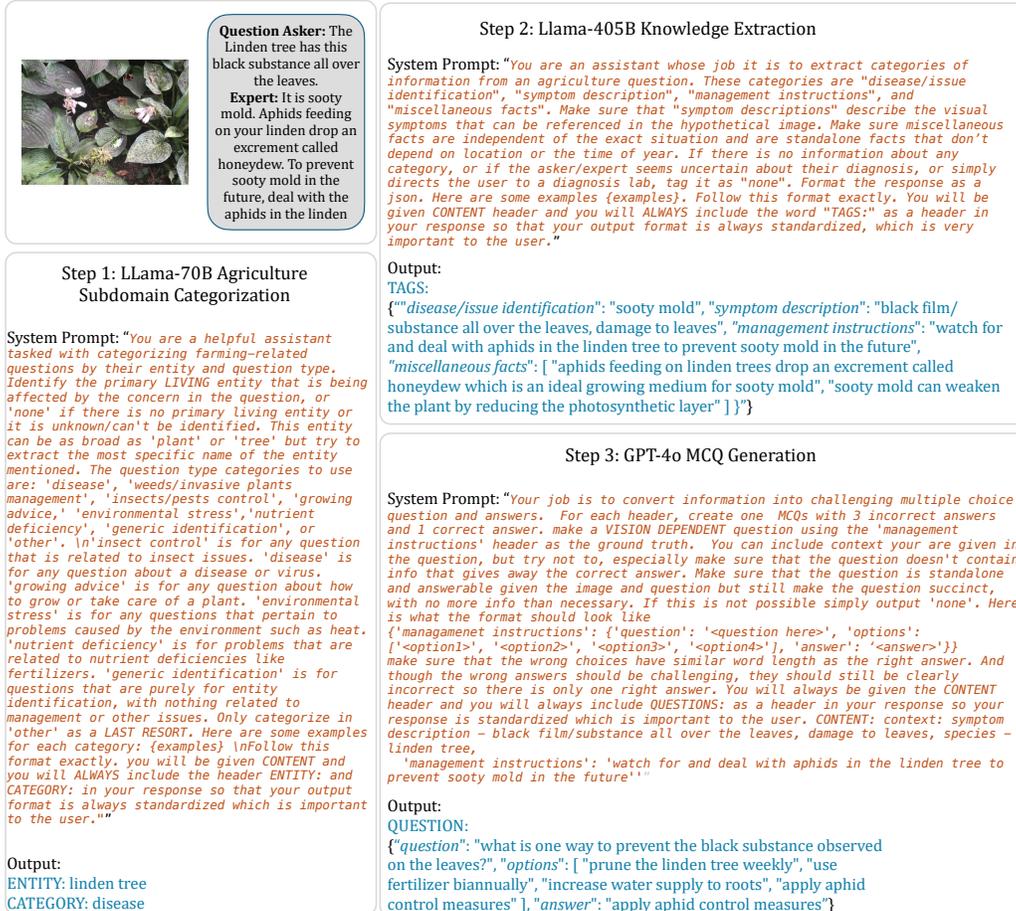


Figure A: Prompts used in different stages of our data curation pipeline.

65 text-only instruction data to help the model retain strong text-only capabilities. We use the official implementation and checkpoint<sup>9</sup> with a LLaMA3-8B LLM backbone and SigLIP-SO400M-Patch14-384 vision encoder in our evaluation.

## 68 B Dataset Curation Details

69 This section outlines the multi-stage curation pipeline of AGMMU and describes the prompts designed for each question type and subdomain.

### 71 B.1 Stage 1: Question Categorization

72 In the first step, we employ the Llama-70B model [4] to categorize questions into predefined agriculture subdomains while identifying the primary living entity affected by the query. Our systematically crafted prompt (Figure A) guides the model to extract the most specific living entity mentioned, such as “apple tree” or “honeybee,” or to assign “none” when the entity is unclear or absent.

76 The subdomains include *Disease*, *Weeds/Invasive Plant Management*, *Insect/Pest Control*, *Growing Advice*, *Environmental Stress*, *Nutrient Deficiency*, *Generic Identification*, and *Other*. Each subdomain is succinctly defined within the prompt, with illustrative examples provided in Figure B to address ambiguous or edge-case scenarios. The prompt enforces a standardized output format, ensuring consistency with the inclusion of “ENTITY:” and “CATEGORY:” headers.

81 To enhance robustness, the prompt includes examples of complex or overlapping cases, ensuring accurate classification even for questions that span multiple subdomains or lack explicit details. By

<sup>9</sup><https://github.com/NVLabs/VILA>

## Agriculture Domain + Species Extraction (step 1) Examples

CONTENT:  
Best ways to treat Azaleas and Mt. Laurel with a Lace infestation. #875156

Question Asker:  
My Azalea shrubs and Mt. Laurel are infested with Lace. I have sprayed them but I am not sure if I sprayed them adequately. 15ome of the shrubs are tall, and I will need a ladder to reach the top.

Is it too late to apply a liquid application around the base? How often should my shrubs be treated? Should I hire a private company?

any suggestions would be much appreciated.

Expert:  
Can you share a photo or two of what you are seeing? You can attach them directly to this reply. First off, stop spraying, and let us know what you are using. You wouldn't see damaged leaves recover, you would just see any new leaves look healthy. You can also burn leaves of plants (or affect non-target insects and plants when spray applications are made when the weather is hot or windy. Lace bugs tend to be worse in landscapes where azaleas are planted in full sun (which stresses them) and where pesticides are regularly used. In healthier landscapes with little or no pesticide use and an abundance of different plants, their populations are kept in check by beneficial insects. That is the ideal goal. Azalea can get lace bugs that are specific to azalea that you can learn about here: <link> there are also different lace bugs that are specific to Rhododendrons and Japanese Andromeda as well but we don't see them on Mountain Laurel. Let us see what your concerns are on the Mountain Laurel and we will assist. The most common problems that those shrubs have tend to be holes in the leaves (Shothole, which can look like insect chewing but is not, and is cosmetic and no chemical controls are recommended) and bark scale insects, which would look like white flocking along the limbs. Systemic soil drenches containing imidacloprid (a type of neonicotinoid) have been found to be damaging to pollinators. In 2016, the Maryland Pollinator Protection Act was passed, which prohibits homeowners from applying them. Only professional, licensed applicators may do so. For this reason, applications would only be a last resort, and there are many other, less toxic, more environmentally friendly ways to deal with many pests. Here is a page that explains more: <link>

ENTITY:  
azalea, mt. Laurel  
CATEGORY:  
insects/pests control

CONTENT:  
Large brown spots on bush bean leaves #873890

Question Asker:  
Hi,

I'm wondering what these brown patches on my green beans are and if there's anything I should do to stop/prevent further issues.

Thanks!  
Expert:  
This looks like abiotic damage, which means it was caused by environmental factors and not a pest or disease. In this case, it looks like sunscorch (also called sunscald or just "scorch"), which is essentially sunburn. Plants with reduced air circulation, such as being crowded or growing near a wall, solid fence, or near heat-reflective pavement or stones can be more vulnerable to scorch, but even well-spaced and unobstructed plants can still develop it. Beans can be among the more vulnerable veggies to scorch. Fortunately, mild scorch in beans generally does not affect yield. You can keep monitoring the plants for watering needs; feeling the soil a few inches down and watering if it becomes somewhat dry to the touch, but no other intervention is needed. Floating row cover and insect mesh netting can serve as a shade cloth of sorts (even if not needed for their pest-excluding or frost-shielding properties) if a full sun exposure is proving to be too much for certain plants, but we'd expect these plants will grow out of it well enough on their own. (Injured leaves cannot heal, but new foliage should emerge normally.)

Question Asker:  
Glad to hear this. Thanks

Expert:  
You're welcome!

ENTITY:  
green beans  
CATEGORY:  
environmental stress

CONTENT:  
Dead Grass #829812

Question Asker:  
Hi There,

Last summer was hard on my grass with most of it dying, particularly in full sun areas. I'm left with some dead patches but mostly bare dirt. I'm interested in doing a no-mow grass on my slope, regular on the flat yard, and am looking for recommendations on if I should sod/seed and what varieties and extra care steps (fertilizer, watering times, etc.) you might recommend. Thank you kindly!

Expert:  
No mow options can exceed city ordinances. Because city ordinances sometimes limit what can be planted on boulevards, you might want to check that first. It would also be good to get a soil test to see what plants are a good match to your soil. See: <https://soiltest.ci.ans.uminn.edu/>. The steepness of some of the area suggests that you also need erosion control for the area. The following websites offer some ideas that may help you decide. Whatever you chose, a deep rooted planting is better for this area than shallow rooted plants like grass. 1. For steep slopes see page 34. <link> 2. For landscape design see: <link> 3. For native prairie plants that require no fertilizer or watering see: <link> 4. For low growing ground covers see:<link> 5. Also see: <link> You could also take a trip to the Minnesota Arboretum in Chanhassen and see some examples of plants that may interest you.

ENTITY:  
grass  
CATEGORY:  
growing advice

CONTENT:  
What is this plant? #874057

Question Asker:  
I originally got this as a stray seedling with a peony plant I purchased at a local nursery. I potted it out of curiosity. It's grown into a lovely good sized plant. Can you tell me what this is? Thank you for your help.

Expert:  
Hello, happy to help. I suspect it may be a weed but I'd be happy to continue working with you to identify it. Could you send a photo of its flower and what month it bloomed when that happens? Though the question may look closed, when you add a reply, it will reopen and notify me. Thanks!

Question Asker:  
So far there hasn't been a hint of blossom or flower. Below are pics from just now. The largest leaf is now 7.5". I have not seen anything like this growing wild in our area. I live in rural Isanti county. Sandy soil country. Thank you so much for your help

Expert:  
Hello, It's burdock an invasive weed. The common burdock can be found everywhere in Minnesota but there are three varieties and all of them are invasive and should be eradicated. Here is information about all three types from Minnesota Wildflowers. You could cut one of the stems to see if it's hollow or not. If not hollow, it is the likely newer variety called Actium lappa. Good-luck!

Question Asker:  
Thank you so much for researching this for me. The leaves do look similar to the Great Burdock. Leaves on the other 2 are too pointy. I don't see any branchy stem coming up or flower buds. I got this seedling in April.Perhaps this matures late summer? Being as it is contained in a pot on my patio I will let it mature to see what it does. Should be interesting. Thank you again for naming my Mystery Plant and letting me know I shouldn't plant it in the garden!

ENTITY:  
burdock  
CATEGORY:  
weed/invasive plants management

CONTENT:  
Question about freezer jam #875023

Question Asker:  
Hello. I'd like to make both a strawberry and strawberry rhubarb freezer jam, however, for health reasons, I'd prefer to use raw honey in place of sugar.

I'm curious—can I substitute honey for sugar in any freezer jam recipe, and, if so, how much? Also wondered if you had any recipe suggestions in this vein.

Secondly, I have found recipes that already call for honey in lieu of sugar, if I was to use these or make my own substitution and use Suregel, is it safe to let the jam sit out at room temperature for the 24 hours required when using Suregel?

Thank you for your time.

Expert:  
Hi, As per the National Center for Home Food Preservation and USDA, Corn syrup and honey may be used to replace part of the sugar in recipes, but too much will mask the fruit flavor and alter the gel structure. Use tested recipes for replacing sugar with honey and corn syrup. (<links> if you are trying to reduce sugar, please know that honey is also pure sugar, just from a different source – so simply substituting this is not a solution to that challenge. There is information in the above link that does talk about making jams/jellies with reduced sugar – one option is using a "low-methoxyl pectin", which the brand name is Pamaona, another option you may want to try. You can substitute honey in Suregel products and it is safe to leave out for 24 hours when canned. I hope I have answered all of your questions, if not, please respond with further questions. Thank you,

ENTITY:  
strawberry  
CATEGORY:

Figure B: Examples included in prompt during the agriculture domain categorization (step 1).





Figure D: Prompts used for each subdomain during knowledge extraction (step 2).

83 embedding these clarifications, the design supports reliable categorization across diverse agricultural  
 84 contexts.

## 85 B.2 Stage 2: Information Extraction

86 In the second step, we design prompts to extract granular categories of information from agricultural  
 87 questions. These categories are tailored to the specific subdomain identified in Step 1, ensuring that  
 88 the extracted information is both relevant and actionable.

89 **Weeds/Invasive Plants Management.** For the “weeds/invasive plants management” subdomain,  
 90 the extraction focuses on: (1) *Image Description*, visual characteristics of the weed or invasive  
 91 plant, (2) *Management Instructions*, actionable strategies for control, and (3) *Miscellaneous Facts*,  
 92 contextual expert insights. The name of the weed itself is already extracted in Step 1. This ensures  
 93 that the emphasis remains on descriptions, actionable measures, and expert knowledge.

94 **Insects/Pests Control.** For this subdomain, the categories include: (1) *Insect/Pest*, identifying the  
 95 pest in focus, (2) *Image Description*, visual traits of the pest or evidence of damage, (3) *Management*  
 96 *Instructions*, guidance for mitigation, and (4) *Miscellaneous Facts*, contextual expert insights. The  
 97 primary plant affected, if exists, is identified in Step 1, thus this step concentrates on pest-specific  
 98 details, such as visual features or damage patterns, and the corresponding management strategies.

99 **Nutrient Deficiency, Disease, Environmental Stress.** For these subdomains, we group them due  
 100 to shared characteristics. The extracted categories are: (1) *Disease/Issue Identification*, specifying  
 101 the underlying cause, (2) *Symptom Description*, observable signs such as discoloration or stunted  
 102 growth, (3) *Management Instructions*, remediation or prevention strategies, and (4) *Miscellaneous*  
 103 *Facts*, contextual expert insights. These subdomains are defined by their symptomatic presentation,  
 104 the underlying conditions, and the need for targeted management interventions.

## Short Answer OEQ Grading Prompt

Your job is to grade student answers from the agriculture and biology domain. Your job is to look at a question, a gold target, and a predicted answer, and then assign a grade of either ['CORRECT', 'INCORRECT', 'NOT ATTEMPTED', 'PARTIALLY CORRECT'].

First, I will give examples of each grade, and then you will grade a new example. {examples}

Remember the following key points:

- a statement should be AT LEAST partially correct if the predicted answer is a subcategory of the gold target or the gold target is a subcategory of the predicted answer
- a statement is always partially correct if it has ANY overlap in content with the target

Grade the predicted answer of this new question as one of:

A: CORRECT  
B: INCORRECT  
C: NOT\_ATTEMPTED  
D: PARTIALLY CORRECT

Question: {question}  
Gold Target: {target}  
Predicted Answer: {predicted\_answer}

Just return the letters "A", "B", "C", or "D", with no text around it.

Figure E: Grading prompt for our LLM-as-judge on short-answer OEQ.

Insect/Pest		Disease/Issue Identification	
<p><b>EXAMPLE 1:</b> Question: What insect is depicted in this image with distinctive yellow and black coloring? Gold Target: hornet Predicted Answer: Hornet nest Grade: CORRECT -This answer is correct because the correct insect is identified, so it doesn't matter that the habitat is mentioned too.</p> <p><b>EXAMPLE 2:</b> Question: What species of spider is shown in the image? Gold Target: cross orb weaver Predicted Answer: Orb weaver Grade: PARTIALLY CORRECT -This answer is partially correct because it gets the (family/orb weaver) correct, but not the species (cross orb weaver).</p> <p><b>EXAMPLE 3:</b> Question: What insect is indicated by the image? Gold Target: bagworms Predicted Answer: Bagworm moth Grade: CORRECT -This is correct because it correctly identifies the insect type.</p>	<p><b>EXAMPLE 4:</b> Question: What insect is indicated by this image? Gold Target: dragonfly beetle Predicted Answer: Carpenter ant Grade: INCORRECT -This answer is INCORRECT because they are referring to different species</p> <p><b>EXAMPLE 5:</b> Question: What insect is indicated by this image? Gold Target: Aphids Predicted Answer: I don't know Grade: NOT ATTEMPTED -This answer is NOT ATTEMPTED because it does not try to answer the question</p> <p><b>EXAMPLE 6:</b> Question: What insect is indicated by this image? Gold Target: spongy moth larva Predicted Answer: Caterpillar Grade: PARTIALLY CORRECT -This answer is PARTIALLY CORRECT because it gets the life stage correct (caterpillar = larvae), but is not as specific as the target.</p>	<p><b>EXAMPLE 1:</b> Question: What issue is indicated by the condition of the plants foliage in the image? Gold Target: burn on the foliage Predicted Answer: Chemical burn Grade: CORRECT -This answer is correct because they are both describing the same issue of a burn.</p> <p><b>EXAMPLE 2:</b> Question: What condition affects the plant in the image? Gold Target: heat stress Predicted Answer: Drought Grade: CORRECT -This answer is correct because they are both describing the same issue.</p> <p><b>EXAMPLE 3:</b> Question: What issue is affecting the plant as shown in the image? Gold Target: powdery mildew Predicted Answer: Fungal infection. Grade: PARTIALLY CORRECT -This answer is PARTIALLY CORRECT because it gets the general type of disease correct (fungal) but not the specific type of fungal (powdery mildew).</p>	<p><b>EXAMPLE 4:</b> Question: What issue is affecting the plant in the image? Gold Target: blossom end-rot Predicted Answer: I don't know Grade: NOT ATTEMPTED -This answer is NOT ATTEMPTED because it does not try to answer the question</p> <p><b>EXAMPLE 5:</b> Question: Which diseases affect the plant in the image? Gold Target: crown gall and powdery mildew Predicted Answer: Crown gall, scale insects. Grade: PARTIALLY CORRECT -This is partially correct because the prediction contains parts of the correct answer (crown gall) but is missing part (powdery mildew).</p> <p><b>EXAMPLE 6:</b> Question: What issue is affecting the plant in the image? Gold Target: drought stress Predicted Answer: Leaf scorch. Grade: PARTIALLY CORRECT -This is partially correct because the prediction, leaf scorch, is a type of drought stress, but the answer is more specific than the target.</p>
Species			
<p><b>EXAMPLE 1:</b> Question: What species is in the image? Gold Target: Potato plant Predicted Answer: I am not sure Grade: NOT ATTEMPTED -This answer is NOT ATTEMPTED because it does not try to answer the question.</p> <p><b>EXAMPLE 2:</b> Question: What species is shown in the image? Gold Target: Spider Predicted Answer: Wolf spider Grade: CORRECT -This answer is CORRECT because wolf spider is a type of spider.</p>	<p><b>EXAMPLE 3:</b> Question: What species is in the image? Gold Target: Kickxia sparis Predicted Answer: Yellow monkeyflower Grade: INCORRECT -This answer is INCORRECT because it identifies a different species than the target.</p> <p><b>EXAMPLE 4:</b> Question: What species is in the image? Gold Target: Banana squash Predicted Answer: Zucchini squash Grade: PARTIALLY CORRECT -This answer is PARTIALLY CORRECT because though it gets the genus r ight (squash), but gets the species wrong (Zucchini instead of Banana).</p>		

Figure F: Unique examples included for short-answer categories added to the grading prompt for our LLM-as-judge on short answer OEQ.

## Multi-Statement OEQ Grading Prompt

Your job is to grade student answers from the agriculture and biology domain. Your job is to look at a question, a gold target, and a predicted answer, and then assign grades to each statement in the response of 'correct', 'partially correct', 'incorrect', 'missing', 'irrelevant'.

- Correct is assigned to statements from the predicted answer that fully semantically map to a statement in the gold target.
- Partially correct is assigned to statements which partially semantically map to a statement in the gold target.
- Incorrect is assigned to statements from the predicted answer that directly semantically contradict a statement in the gold target.
- Missing is assigned to statements in the gold target which haven't been mapped within correct, partially correct, or incorrect.
- Irrelevant is assigned to statements in the predicted answer which neither directly contradict nor correspond in any way to statements in the gold target. EACH STATEMENT IN THE GOLD TARGET AND PREDICTED ANSWER SHOULD BE ASSIGNED TO EXACTLY ONE OF THESE CATEGORIES. Here are examples of correctly graded statements: {examples}

Remember the following key point:

- a statement is always partially correct if it has ANY overlap in content with the target.

Question: {question}  
Gold Target: {expected}  
Predicted Answer: {actual}

Follow the format of the examples exactly. Output only a json with no additional text.

Figure G: Prompt for categorizing statements in our LLM-as-judge on multi-sentence (long-answer) OEQ.

Symptom/Visual Description Examples	Management Instructions Examples
<p><b>EXAMPLE 1:</b> Gold Target: ['Top of the tree is bare of needles'] Predicted Answer: The image features a tall tree with a mix of green foliage and some bare branches. The background is a clear blue sky. The tree appears to have healthy green leaves at the lower part, while the upper branches are more sparse and dry. Scoring: {'correct': {'upper branches are more sparse and dry': 'top of the tree is bare of needles'}, 'incorrect': [], 'partially correct': [], 'missing': [], 'irrelevant': ['healthy green leaves at the lower part', 'tall tree with a mix of green foliage and some bare branches']} <b>EXAMPLE 2:</b> Gold Target: ['Leaves turning brown', 'leaves dying', 'new growth is dying'] Predicted Answer: The image features a variety of leaves in different colors and conditions. There are green leaves, some with yellowing edges, and a few that are brown and wilted. The overall composition shows a mix of healthy and unhealthy foliage, indicating a range of plant health. The background consists of more green leaves, creating a dense, leafy appearance. Scoring: {'correct': {'brown and wilted leaves': 'leaves turning brown'}, 'incorrect': [], 'partially correct': ['unhealthy foliage': 'leaves dying'], 'missing': ['new growth is dying'], 'irrelevant': ['The background consists of more green leaves, creating a dense, leafy appearance.', 'There are green leaves, some with yellowing edges']}] <b>EXAMPLE 3:</b> Gold Target: ['Brown spots on leaves'] Predicted Answer: The image features several tall green plants with elongated, fleshy leaves. The leaves show signs of damage, including brown spots and edges, indicating possible disease or environmental stress. The plants have small clusters of yellowish flowers at the top. In the background, there is a maroon door and a welcome mat. Scoring: {'correct': {'leaves show signs of damage, including brown spots and edges': 'brown spots on leaves'}, 'incorrect': [], 'partially correct': [], 'missing': [], 'irrelevant': ['tall green plants with elongated, fleshy leaves', 'small clusters of yellowish flowers at the top']}]</p>	<p><b>EXAMPLE 1:</b> Gold Target: ['The entire plant (including the root/rhizome system) by digging around the plant's base'] Predicted Answer: To manage nutcase, consider the following strategies: 1. **Cultural Control**: Improve drainage and reduce soil compaction to make the environment less favorable for nutcase growth. Regularly mow or cut the plants before they seed. 2. **Mechanical Control**: Hand-pulling can be effective, but ensure you remove the entire root system to prevent regrowth. This may need to be repeated multiple times. 3. **Chemical Control**: Use pre-emergent herbicides in early spring and post-emergent herbicides specifically labeled for nutcase. Follow all application guidelines carefully. 4. **Mulching**: Apply a thick layer of mulch to suppress growth and prevent light from reaching the nutcase. 5. **Monitoring and Follow-up**: Regularly check for new growth and address it promptly to prevent establishment. Combining these strategies will yield the best results in managing nutcase effectively. Scoring: {'correct': {'mechanical control': 'hand-pulling can be effective, but ensure you remove the entire root system to prevent regrowth. this may need to be repeated multiple times.'}, 'incorrect': {'chemical control': 'use pre-emergent herbicides in early spring and post-emergent herbicides specifically labeled for nutcase. follow all application guidelines carefully.'}, 'missing': [], 'irrelevant': {'cultural control': 'improve drainage and reduce soil compaction to make the environment less favorable for nutcase growth. regularly mow or cut the plants before they seed.', 'mulching': 'apply a thick layer of mulch to suppress growth and prevent light from reaching the nutcase.', 'monitoring and follow-up': 'regularly check for new growth and address it promptly to prevent establishment.'}] <b>EXAMPLE 2:</b> Gold Target: ['Deep water every 2 to 3 weeks through dry summer months', 'irrigate the entire area underneath the canopy, plus 1 to 2 feet beyond', 'moisten the soil to at least 12 inches deep', 'determine irrigation runtime by testing soil moisture penetration'] Predicted Answer: To ensure consistent watering, especially during dry periods, as elevated planters can dry out quickly. Deep water every 2 to 3 weeks through dry summer months. Irrigate the entire area underneath the canopy, plus 1 to 2 feet beyond. Moisturize the soil to at least 12 inches deep. Determine irrigation runtime by testing soil moisture penetration. Scoring: {'correct': [], 'incorrect': [], 'partially correct': ['ensure consistent watering, especially during dry periods, as elevated planters can dry out quickly.', 'deep water every 2 to 3 weeks through dry summer months.', 'irrigate the entire area underneath the canopy, plus 1 to 2 feet beyond.', 'moisten the soil to at least 12 inches deep.', 'determine irrigation runtime by testing soil moisture penetration.'], 'missing': [], 'irrelevant': ['consider mulching around the base to retain moisture and improve soil health.', 'monitor for pests or diseases.', 'consult an arborist if issues persist.']} <b>EXAMPLE 3:</b> Gold Target: ['The plants have deep roots and should be dug up', 'or you can use chemical control'] Predicted Answer: To manage Clematis (leaves/honeyvine/bluevine), consider the following strategies: 1. **Mechanical Removal**: Cut and pull the vines, ensuring to remove as much of the root system as possible to prevent regrowth. 2. **Herbicide Application**: Use a systemic herbicide containing glyphosate or triclopyr, applying it to freshly cut stems or foliage during the growing season for effective control. 3. **Monitoring and Follow-up**: Regularly check the area for new growth and reapply management strategies as needed. 4. **Preventive Measures**: Promote native vegetation to outcompete the invasive species and reduce its spread. Always follow local guidelines and regulations when using herbicides. Scoring: {'correct': {'mechanical removal': 'cut and pull the vines, ensuring to remove as much of the root system as possible to prevent regrowth.'}, 'partially correct': [], 'missing': [], 'irrelevant': {'herbicide application': 'use a systemic herbicide containing glyphosate or triclopyr, applying it to freshly cut stems or foliage during the growing season for effective control.', 'or you can use chemical control.'}, 'incorrect': [], 'partially correct': [], 'missing': [], 'irrelevant': {'monitoring and follow-up': 'regularly check the area for new growth and reapply management strategies as needed.', 'preventive measures': 'promote native vegetation to outcompete the invasive species and reduce its spread.'}] <b>EXAMPLE 4:</b> Gold Target: ['Easily pulled by hand', 'remove them before they get to seed'] Predicted Answer: The recommended management strategy for garlic mustard includes: 1. **Manual Removal**: Pull out the plants by hand, ensuring to remove the entire root system, especially before they set seed in late spring. 2. **Regular Monitoring**: Check the area regularly for new growth and remove any emerging plants. 3. **Mulching**: Apply mulch to suppress further growth and prevent seed germination. 4. **Herbicide Application**: If necessary, consider using a targeted herbicide in the fall when native plants are dormant, following local guidelines. Always ensure to dispose of removed plants properly to prevent establishment. Scoring: {'correct': {'manual removal': 'pulling to remove the entire root system, especially before they set seed in late spring.', 'remove them before they get to seed.'}, 'missing': [], 'irrelevant': {'regular monitoring': 'check the area regularly for new growth and remove any emerging plants.', 'herbicide application': 'if necessary, consider using a targeted herbicide in the fall when native plants are dormant, following local guidelines.', 'mulching': 'apply mulch to suppress further growth and prevent seed germination.'}]</p>

Figure H: Unique examples included for multi-statement categories added to the grading prompt for our LLM-as-judge on multi-sentence (long-answer) OEQ.

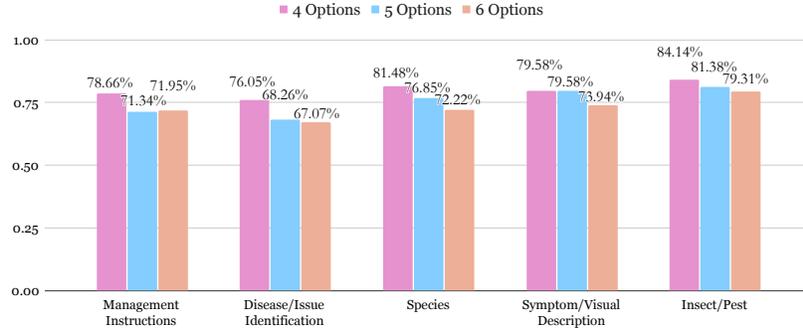


Figure I: **GPT-4o accuracy with increasing MCQ options.** Model performance on MCQs across different categories, comparing accuracy scores when varying the number of answer options (4, 5, and 6). We observe a 5-10% difference in accuracy across categories between the 4-option and 6-option configurations, with performance generally decreasing as the number of options increases.

105 **Growing Advice.** For this subdomain, the variability in question structure necessitates tailored  
 106 extractions: (1) *Succinct Question*, a concise reformulation of the user query, (2) *Succinct Answer*,  
 107 a precise response to the query, (3) *Image Description*, any relevant visual details, and (4) *Miscella-*  
 108 *neous Facts*, contextual expert insights.

109 Importantly, besides distinguishing the extraction types, we also put different examples of pre-made  
 110 knowledge extraction into the prompt, see Figure C. Prompts given to the model for each subdomain  
 111 can be seen in Figure D.

112 The *Miscellaneous Facts* category is extracted across all subdomains but is not directly used in  
 113 subsequent steps. Instead, it captures standalone expert information that can contextualize a user’s  
 114 issue.

115 To optimize extraction accuracy, we distinguish between “Symptom Description” (used for nutri-  
 116 ent deficiency, disease, and environmental stress) and “Image Description” (used for weeds/invasive  
 117 plants and insects/pests). While these serve a similar purpose—capturing observable or visual de-  
 118 tails—they are unified under the term “Symptom/Visual Description” in subsequent steps to main-  
 119 tain consistency.

### 120 B.3 Stage 3: Question Generation

121 In the final step, the extracted agricultural facts are transformed into evaluative question-answer  
 122 (QA) pairs, comprising multiple-choice questions (MCQs) and open-ended questions (OEs) gen-  
 123 erated using GPT-4o. To enhance relevance, we exclude two knowledge types: (1) *Growing Advice*,  
 124 as image content often lacks direct correlation with the user’s issue, and (2) *Miscellaneous Facts*,  
 125 since these provide general context but do not directly relate to the user’s image. This refinement  
 126 narrows the scope to five key knowledge types for downstream processing, including *Disease/Issue*  
 127 *Identification*, *Symptom/Visual Description*, *Management Instructions*, *Insect/Pest*, and *Species*.

128 To ensure clarity and relevance, we employ a standardized prompt structure (see Figure A tailored  
 129 to each knowledge type. While the core structure remains consistent, the phrasing explicitly refer-  
 130 ences the specific knowledge type being addressed. This targeted design allows the prompts to focus  
 131 on generating well-contextualized and relevant questions. For added precision, the prompts incor-  
 132 porate contextual details where applicable: (1) For *species-related questions*, only symptom/visual  
 133 description information is referenced, ensuring the focus remains on observable traits, and (2) for  
 134 *symptom/visual-related questions*, species information is used to provide context, helping to ground  
 135 the questions in specific agricultural scenarios.

136 This contextualization ensures that the generated questions integrate both user-provided informa-  
 137 tion and extracted context seamlessly. The result is a set of comprehensive and “fair” evaluative  
 138 questions, designed to effectively assess multimodal agricultural understanding.

## 139 B.4 Final Stage: Human Verification

140 To guarantee the quality of the evaluation questions, we implemented a human verification process  
141 that validates faithfulness, certainty, quality, and MCQ feasibility. The data was distributed through  
142 an HTML file containing AGMMU questions and answers, original user questions, expert answers,  
143 and corresponding images. Each annotator was given a corresponding Excel file where the user  
144 just has to mark false (uncheck the box) for each condition not met per question. To further assist  
145 the annotator, we provided a few complex examples of questions that meet and do not meet the  
146 requirements, functioning as in-context examples. After collecting these data, only the completely  
147 unproblematic ones (all boxes remain checked) were kept.

148 **Faithfulness:** *Do you think the question, ground-truth, and context extract faithful information*  
149 *from the original farmer question?* Our questions are directly based on the original questions and  
150 this step functions as a sanity check ensuring the quality of our dataset. The annotator needs to read  
151 through the question and the original conversations between the user and the expert.

152 **Certainty:** *Is the expert certain about the answer?* As our ground truth answers are extracted from  
153 the expert answers, we only want to include those that are very certain. A higher certainty from the  
154 expert means that it is more likely to be correct. We observe that the behaviors of the annotator are  
155 to read the responses from the expert and look for keywords like “may,” “not sure,” “you have to go  
156 to a lab for further inspection.”

157 **Quality:** *Are the images suitable for answering the questions?* (Images are not in low-resolution,  
158 blurs, pure blackness, etc.) *Are the image/symptom descriptions visible in the presented images?* As  
159 our benchmark attempts to evaluate the visual understanding of the models, our human verification  
160 removes the questions that do not depend on the images and those with broken images. For example,  
161 it is not fair for our question to ask about the fruit of the plant when the submitted photos only capture  
162 the leaves of the tree or the image is blurry.

163 **Feasibility:** *Are all of the wrong choices wrong?* The incorrect choices were generated with GPT-  
164 4o, so we need to check to ensure there are no multiple correct answers or an answer that overlaps  
165 with the correct answer and remove. For example, the wrong choice might be the common name of  
166 a species displayed by the scientific name in the ground truth.

## 167 C More Evaluation, Implementation, and Design Choices

168 **LLM-as-judge.** To perform evaluation on few-word and multi-statement OEQ responses, we im-  
169 plemented the LLM-as-judge methodology using GPT-4.1. Our prompts for few-word responses  
170 (Figure E, F) and multi-statement responses (Figure G, H) contain several in-context examples  
171 based on the question category to guide the LLM to correctly categorize the answer as “correct,”  
172 “incorrect,” “partially correct,” and “irrelevant.”

173 **Number of MCQ options.** To determine the optimal number of answer choices for our MCQs,  
174 we conducted an ablation study comparing GPT-4o’s accuracy when presented with four, five, and  
175 six options. For efficiency, we conducted this experiment on a subset of 821 questions, generating 5  
176 wrong answers with GPT-4o. We randomly choose 3, and 4 wrong answers, for the four-choice and  
177 five-choice experiment, respectively, and take all choices for the six-choice experiment. While this  
178 limited subset may not capture the full variability of the dataset, it provides sufficient evidence to  
179 inform our design decisions. Due to the risk of process of elimination with MCQs, we believe that  
180 OEQs more accurately capture model performance.

181 The results, shown in Figure I, indicate that accuracy decreases as the number of answer choices  
182 increases. Specifically, we observed a 5-10% reduction in accuracy between the 4-option and 6-  
183 option configurations. This trend suggests that the model might rely on a process of elimination  
184 when selecting an answer, making it more challenging to identify the correct response as the number  
185 of options increases. While the decrease in accuracy is not overly significant, we think it justifies  
186 our choice to use four options for MCQs.

187 **Implementation of AGBASE fine-tuning.** We fine-tune the LLaVA-v1.5-7B model using a LoRA-  
188 based setup. The training is performed with a learning rate of  $2e-4$ , without weight decay, and a  
189 cosine learning rate schedule with a 3% warm-up ratio. We use a per-device batch size of 16 with

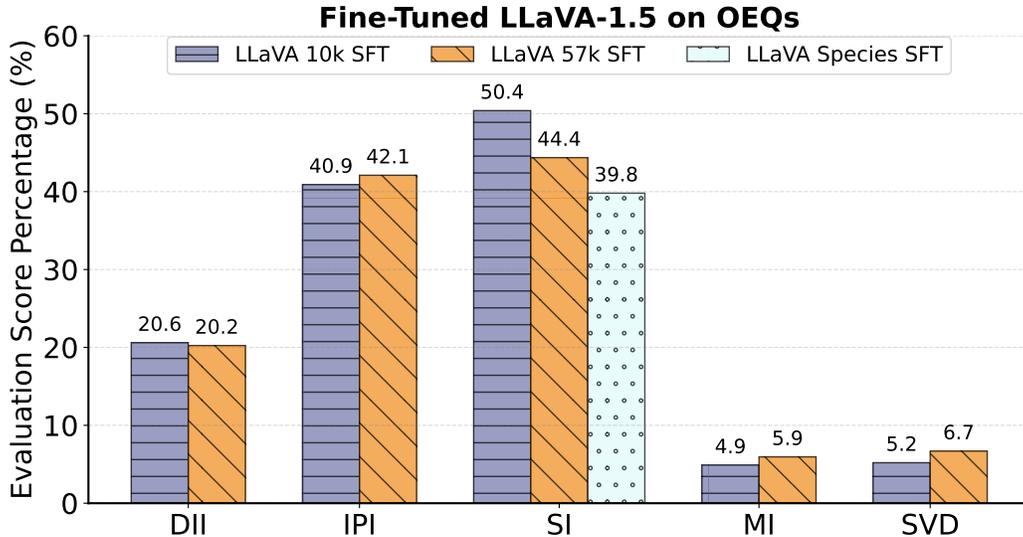


Figure J: Evaluation scores across five error categories for three fine-tuning setups using the AG-BASE dataset. Models are fine-tuned on: (1) **LLaVA 10k SFT**—a mix of AGBASE and 10k LLaVA-Instruct samples, (2) **LLaVA 57k SFT**—a 50-50 blend of AGBASE and LLaVA’s original 57k SFT samples, and (3) **LLaVA Species SFT**—a specialized set focused on species identification with contextual augmentation.

190 gradient accumulation steps set to 2, resulting in an effective batch size of 64. The model is trained  
 191 over 2 epochs using 2 NVIDIA A6000 GPUs.

192 Dataset preparation involved curating structured multi-turn conversations from a horticultural FAQ  
 193 knowledge base, paired with user-uploaded images. From an initial pool of 367,331 QA-image  
 194 pairs, we filtered out questions that had a species value in [*tree, bee, shrub, weed, wasp, plant, insect,*  
 195 *grass, none, moth, beetle, snake, caterpillar, spider, ant, mushroom, fungus*], because we observe  
 196 that questions with these common non-species species extractions often contain vague or uncertain  
 197 examples. This gives us a high-quality dataset of 57,079 samples. Considering the influence of  
 198 data mixture for training VLMs, we conduct three fine-tuning experiments. (1) The first experiment  
 199 involves fine-tuning on a combination of our domain-specific dataset, AGBASE, and 10,000 samples  
 200 from LLaVA’s original instruction-tuning dataset, LLaVA-Instruct-150K. (2) The second experiment  
 201 employs a 50-50 mixture of AGBASE by using 57,079 samples from LLaVA’s original SFT set [7]  
 202 (3) The third experiment focuses solely on species identification and consists of 18,109 QA pairs  
 203 constructed by prepending the full original user queries to 33,777 generic identification samples,  
 204 allowing us to test the effect of user context on classification accuracy.

205 In Figure J, we find that the LLaVA 10k SFT model achieves a slightly higher overall accuracy (0.25)  
 206 compared to the LLaVA 57k SFT model (0.24), suggesting that a smaller, well-curated dataset mixed  
 207 with domain-specific data may be more effective than a larger, more generic one for knowledge-  
 208 intensive domain fine-tuning. Additionally, the LLaVA Species SFT model, which includes added  
 209 user query context for species identification, performs worse than the other models in the species  
 210 category, indicating that this additional context provides limited benefit for classification accuracy.

## 211 D More Dataset Visualization

212 In Figure K, we demonstrate more samples in AGMMU with questions and **multiple choice an-**  
 213 **swers.**

214 In Figure L, we demonstrate more samples in AGMMU with **open-ended questions** and responses.  
 215 We especially emphasize the long-form responses required from the model for symptom description  
 216 and management instructions, normally containing multiple facts.

Insect Identification				
	<p>Q: what type of pest is indicated by the yellow speckles on these rose leaves?</p> <p>A: corn earworm B: cabbage worm C: apple maggot D: <b>rose sawfly</b></p>	<p>Q: what insect is responsible for the damage seen on the oregano leaves in the image?</p> <p>A: red spider mite B: <b>four-lined plant bug</b> C: green stink bug D: brown marmorated stink bug</p>	<p>Q: which insect or pest is likely causing the sooty mold on the leaves in the image?</p> <p>A: weevils B: caterpillars C: <b>scales</b> D: beetles</p>	<p>Q: what type of insect is depicted in the image, often mistaken for a speck of dirt due to its size?</p> <p>A: <b>springtails</b> B: weevils C: larvae D: termites</p>
Plant Identification				
	<p>Q: what species of plant is shown in the image?</p> <p>A: prairie aster B: himalayan lily C: <b>mariposa lily</b> D: bluebell flower</p>	<p>Q: what species does the plant with these leaves belong to?</p> <p>A: maple tree B: oak tree C: <b>tree of heaven</b> D: pine tree</p>	<p>Q: what species of plant is shown in the image?</p> <p>A: arrowroot (maranta arundinacea) B: toothwort (cardamine diphylla) C: feverfew (tanacetum parthenium) D: <b>gingerwort (zingiber officinale)</b></p>	<p>Q: what tree species is showing sap oozing out of its trunk?</p> <p>A: <b>peach tree</b> B: plum tree C: cherry tree D: apple tree</p>
Management Instruction				
	<p>Q: what disease is characterized by orange spots on the leaves of a pear tree?</p> <p>A: pear scab B: <b>pacific coast pear rust</b> C: fire blight D: apple scab"</p>	<p>Q: what condition is affecting the cherry laurel leaves shown in the image?</p> <p>A: root rot B: insect infestation C: <b>winter burn</b> D: leaf blight</p>	<p>Q: what is a recommended practice to address potential issues?</p> <p>A: increase nitrogen application B: cover soil with plastic mulch C: <b>apply a copper-based fungicide</b> D: add synthetic fertilizer</p>	<p>Q: based on the appearance of the plant, what is suggested to prevent the spread of the infection?</p> <p>A: use nitrogen fertilizers B: increase water intake C: add more multi around plant D: <b>apply myclobutanil fungicide</b></p>
Disease Categorization				
	<p>Q: based on the bark image showing small exit holes, what is the recommended action for managing an infested tree?</p> <p>A: remove and discard leaves B: prune dead branches C: apply chemical spray D: <b>remove infested tree</b></p>	<p>Q: what is the recommended treatment for the grass shown in the image?</p> <p>A: mow every other day B: <b>use roundup on the entire area</b> C: use organic mulch D: apply fertilizer monthly</p>	<p>Q: based on the visible damage to the tree in the image, what is the most likely cause?</p> <p>A: <b>back rub damage</b> B: insect infestation C: fungal infection D: wind damage</p>	<p>Q: what condition is likely affecting the oak leaves shown in the image?</p> <p>A: anthracnose B: <b>oak leaf blister</b> C: powdery mildew D: leaf spot</p>
Symptom Description				
	<p>Q: based on the image, what symptom affects the holly tree?</p> <p>A: <b>mostly brown leaves</b> B: yellowing bark C: falling branches D: wilting flowers</p>	<p>Q: what unusual feature is present on the rudbeckia flower?</p> <p>A: <b>round growth without seeds</b> B: leaf discoloration C: stunted stem growth D: wilted petals</p>	<p>Q: what symptom is shown on the trunk of the tree?</p> <p>A: fungal rings B: dark spots C: moss growth D: <b>huge split</b></p>	<p>Q: what does the image reveal about the condition of the grass?</p> <p>A: <b>brown turf area</b> B: lush green patch C: overgrown weeds D: sandy texture</p>

Figure K: Additional visualization of samples in AGMMU. Ground truth selections of each question are highlighted in yellow.

## 217 E Limitations and Future Work

218 While our work makes unique contributions to agricultural benchmark development and VLM evalu-  
219 ation through knowledge-intensive tasks, we acknowledge several limitations and identify promising  
220 directions for future research in this section.

221 **Advanced Utilization of Training Data.** Although our curated dataset, AGBASE, has proven  
222 significant effectiveness for fine-tuning VLMs [7] as shown in Section 4 and Figure 6, its potential  
223 extends beyond our current usage. As a comprehensive knowledge repository, the dataset presents  
224 opportunities for knowledge retrieval and augmented generation (RAG) approaches [3]. In partic-  
225 ular, the development of vision-centric multimodal RAG systems remains an under-explored yet

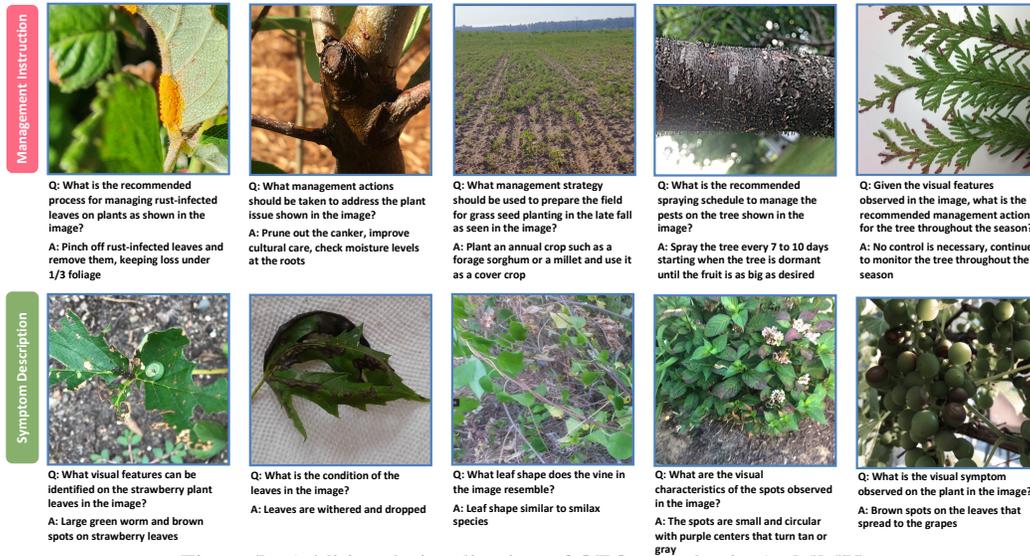


Figure L: Additional visualization of OEQ samples in AGMMU.

226 promising direction. This alternative could enable more effective knowledge extraction and utilization  
 227 from our dataset, potentially improving model performance on agricultural understanding tasks.  
 228 We leave the exploration of these advanced techniques for future work.

229 **Expanded Model Coverage and Evaluation Protocols.** While our current study encompasses  
 230 several state-of-the-art and most commonly used VLMs for zero-shot evaluation and fine-tuning  
 231 analysis, we acknowledge that they represent only a subset of available multimodal architectures  
 232 and methodologies. To enhance the robustness and generalizability of our findings, we plan to in-  
 233 corporate a broader spectrum of VLMs. Additionally, we plan to conduct more extensive ablation  
 234 studies and comparative analyses across different model scales and architectures. This comprehen-  
 235 sive evaluation will provide deeper insights into the relative strengths and limitations of various  
 236 approaches in agricultural understanding tasks.

## 237 F Societal Impact

238 We anticipate no direct negative societal impact of our work. Our dataset is ethically designed,  
 239 respecting the privacy of Extension.org users by removing personal identifying information such as  
 240 name, gender, username, and location. Additionally, we have verified to the best of our ability to  
 241 ensure the removal of images that contain human faces. During dataset curation, we put in great  
 242 effort to eliminate bias by creating a dataset representative of the original Extension.org questions  
 243 as well as a balanced dataset across all question types.

244 **Positive Impact:** We hope that the creation and release of this challenging vision-knowledge inten-  
 245 sive dataset can support active research in this domain. Our comprehensive dataset is adapted from  
 246 real-world conversations between users and experts, creating samples that are more representative  
 247 of questions and images one may ask. This enables more accurate responses as demonstrated by  
 248 our fine-tuning experiments. This dataset can be used to support the development of an agricultural  
 249 vision language model that can provide users with instant assistance on various topics like insect/  
 250 pest identification, disease categorization, and most importantly, management instructions. When  
 251 properly used, these models have the potential to assist sustainability goals, prevent yield loss, and  
 252 improve resource use.

## 253 References

254 [1] Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman,  
 255 Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. GPT-4 technical report.  
 256 *arXiv preprint arXiv:2303.08774*, 2023. 2

- 257 [2] Jinze Bai, Shuai Bai, Shusheng Yang, Shijie Wang, Sinan Tan, Peng Wang, Junyang Lin, Chang Zhou, and  
258 Jingren Zhou. Qwen-vl: A versatile vision-language model for understanding, localization, text reading,  
259 and beyond. *arXiv preprint arXiv:2308.12966*, 2023. [2](#)
- 260 [3] Angels Balaguer, Vinamra Benara, Renato Luiz de Freitas Cunha, Todd Hendry, Daniel Holstein, Jennifer  
261 Marsman, Nick Mecklenburg, Sara Malvar, Leonardo O Nunes, Rafael Padilha, et al. Rag vs fine-tuning:  
262 Pipelines, tradeoffs, and a case study on agriculture. *arXiv preprint arXiv:2401.08406*, 2024. [12](#)
- 263 [4] Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Let-  
264 man, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, et al. The llama 3 herd of models. *arXiv*  
265 *preprint arXiv:2407.21783*, 2024. [2](#), [3](#)
- 266 [5] Bo Li, Yuanhan Zhang, Dong Guo, Renrui Zhang, Feng Li, Hao Zhang, Kaichen Zhang, Yanwei Li, Ziwei  
267 Liu, and Chunyuan Li. LLaVA-OneVision: Easy visual task transfer. *arXiv preprint arXiv:2408.03326*,  
268 2024. [2](#)
- 269 [6] Ji Lin, Hongxu Yin, Wei Ping, Pavlo Molchanov, Mohammad Shoeybi, and Song Han. VILA: On pre-  
270 training for visual language models. In *CVPR*, 2024. [2](#)
- 271 [7] Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. Visual instruction tuning. In *NeurIPS*, 2023.  
272 [2](#), [11](#), [12](#)
- 273 [8] Haotian Liu, Chunyuan Li, Yuheng Li, Bo Li, Yuanhan Zhang, Sheng Shen, and Yong Jae Lee. LLaVA-  
274 NeXT: Improved reasoning, OCR, and world knowledge, 2024. [2](#)
- 275 [9] Machel Reid, Nikolay Savinov, Denis Teplyashin, Dmitry Lepikhin, Timothy Lillicrap, Jean-baptiste  
276 Alayrac, Radu Soricut, Angeliki Lazaridou, Orhan Firat, Julian Schrittwieser, et al. Gemini 1.5: Unlock-  
277 ing multimodal understanding across millions of tokens of context. *arXiv preprint arXiv:2403.05530*,  
278 2024. [2](#)
- 279 [10] Anthropic Team. Introducing the next generation of claude, 2024. [2](#)
- 280 [11] OpenGVLab Team. Internvl2: Better than the best—expanding performance boundaries of open-source  
281 multimodal models with the progressive scaling strategy, 2024. [2](#)
- 282 [12] Shengbang Tong, Ellis Brown, Penghao Wu, Sanghyun Woo, Manoj Middepogu, Sai Charitha Akula,  
283 Jihan Yang, Shusheng Yang, Adithya Iyer, Xichen Pan, Austin Wang, Rob Fergus, Yann LeCun, and  
284 Saining Xie. Cambrian-1: A fully open, vision-centric exploration of multimodal llms. *arXiv preprint*  
285 *arXiv:2406.16860*, 2024. [2](#)

286