

# GEOBench-VLM: Benchmarking Vision-Language Models for Geospatial Tasks

## Supplementary Material

This supplementary material includes the dataset table (S1), dataset verification details (S2), Additional comparisons (S3) and quantitative results to illustrate multiple cases for assessing model responses (S4). It also provides results on multispectral images (S5) along with a comparison between bi-temporal and multi-temporal approaches (S6). Additionally, a geographical analysis is included (S7) to observe the span of data coverage across locations, followed by a detailed description of the word cloud (S8).

### S1. Datasets

The datasets we use in our evaluation cover a wide range of geospatial tasks, showing the variety and depth of challenges in geospatial analysis. As shown in Table A1, these datasets include tasks like scene understanding, spatial relation, instance counting, temporal understanding, referring expression segmentation, and working with non-optical data. This diversity enables us to create versatile question-answer pairs tailored to each specific task. The inclusion of datasets from recent years ensures that our evaluation tackles recent challenges and uses up-to-date information.

These datasets also offer a rich variety of annotation types, sensor data, and spatial resolutions, reflecting the diverse nature of geospatial data. The annotation types range from class labels and bounding boxes to semantic and instance masks, giving different levels of detail for model evaluation. The sensor data includes RGB images, Multispectral Imaging (MSI), and Synthetic Aperture Radar (SAR), with resolutions from fine to coarse scales. This heterogeneity allows us to test models under different imaging conditions and resolutions, fostering robustness and generalizability. For example, datasets like FAIR1M [21], DIOR [5], and DOTA [24] provide high-resolution RGB images with bounding box annotations, which are critical for tasks like object detection and understanding spatial relationships in complex scenes. Temporal understanding datasets such as fMoW [10], xBD[9], PASTIS[16], FPCD[22], and GVLM[26] are crucial for tracking changes over time, helping in tasks such as disaster assessment and monitoring urban development. Non-optical datasets like So2Sat [28] and QuakeSet [15] introduce SAR data, expanding our analysis to situations where optical imagery isn't available due to weather or lighting conditions. Scene Understanding datasets like AiRound [13] and RESICS45 [4] offer class annotations that help categorize large-scale scenes, essential for land use and land cover classification.

### S2. Dataset Verification Details

The benchmark was developed using a structured pipeline with semi-automated curation and manual verification. Two annotators manually annotated spatial relationships, and two others handled validation. All tasks were based on referenced datasets and cross-validated by all four. GPT assisted in drafting QA pairs, but all samples were manually verified. For captioning, we provided spatial and annotated inputs to generate fine-grained descriptions. In total, 750 human hours were spent on data curation, manual spatial-relation annotation, validation, and refinement, covering issues such as ambiguity, hallucinations, and distractor quality.

### S3. Additional Results

We include two additional results. First, we report the (Fig. A1 left) performance of Claude-Sonnet 3.7 on classification tasks. The model shows strong performance in several categories, including Scene Classification (0.76), Spatial Relationship (0.66), and Water Bodies Counting (0.55), with an overall classification average of 0.2957 across 15 tasks. Second, to provide a more informative measure of performance on counting tasks, we compute mean absolute error (MAE) across all models (Fig. A1 right). LLaVA-OneVision and GPT-4o show lower MAEs (63.6 and 69.5 respectively), reflecting more precise quantitative estimation than what is captured by categorical accuracy alone. Including MAE offers a finer-grained understanding of model behavior, especially in tasks involving numerical reasoning.

### S4. Qualitative Results

The images in Fig. A2 show patterns in how the models performed on geospatial tasks relevant to scene understanding. Models perform well in identifying scenes with distinct features, such as “interchange”, where most models succeeded except Ferret [25], RS-LLaVA [3], and Sphinx [11], respectively. In the third image, all models except Ferret correctly identified the “stadium”, demonstrating notable contextual understanding. For the fourth image, only a few models correctly identified “mixed cereal” crops, with failures attributed to the ambiguous nature of crop patterns. The first image in the second row shows dense greenery, indicating a moist environment, with the fire risk correctly classified as “low”. The second image in the second row benefits from clear context, aiding classification as a water treatment facility. In contrast, the third image in the same row lacks context, making it prone to misclassification. This compar-

Name	Task	Annotation Type	Sensor (Res)	Year
AiRound[13]	Scene Understanding, Object Classification	Class	RGB, Sentinel-2 (10m)	2020
RESICS45[4]			RGB	2017
PatternNet[27]			RGB	2018
MtSCCD[12]			RGB (1m)	2024
FireRisk[18]			RGB (1m)	2023
FGSCR[8]			RGB	2021
FAIR1M[21]	Spatial Relation Classification, Referring Expression Detection, Captioning	Bounding Box	RGB (0.3–0.8m)	2021
DIOR[5]			RGB	2020
DOTA[24]			RGB (0.1–1m)	2021
Forest Damage[1]	Counting	Bounding Box	RGB	2021
Deforestation[6]			RGB	2024
COWC[14]			RGB (15 cm)	2016
NASA Marine Debris[17]			RGB (3m)	2024
The RarePlanes Dataset[19]			RGB (0.3m)	2020
fMoW[10]	Temporal Understanding	Class	RGB (1m)	2018
xBD[9]		Bounding Box, Instance Mask, Class	RGB (0.8m)	2019
PASTIS[16]		Semantic Mask	MSI (10m)	2021
FPCD[22]		Semantic Mask	RGB (1m)	2022
GVLM[26]		Class	RGB (0.6m)	2023
DeepGlobe Land Cover[7]	Referring Expression Segmentation	Semantic Mask	RGB (0.5m)	2018
GeoNRW[2]			RGB (1m)	2021
So2Sat[28]	Non-Optical	Class	SAR, MSI (10m)	2020
QuakeSet[15]		Number	SAR (10m)	2024

Table A1. Comprehensive overview of geospatial datasets utilized for evaluating Vision-Language Models (VLMs) across diverse tasks, including Scene Understanding, Spatial Relation, Object Classification, Spatial Relation Classification, Referring Expression Detection, Captioning, Temporal Understanding, Referring Expression Segmentation, and Non-Optical tasks. The datasets are categorized by annotation types (e.g., class, bounding box, semantic mask) and sensor types (e.g., RGB imagery, Multispectral Imaging (MSI), Synthetic Aperture Radar (SAR)), highlighting their versatility for a wide range of geospatial applications.

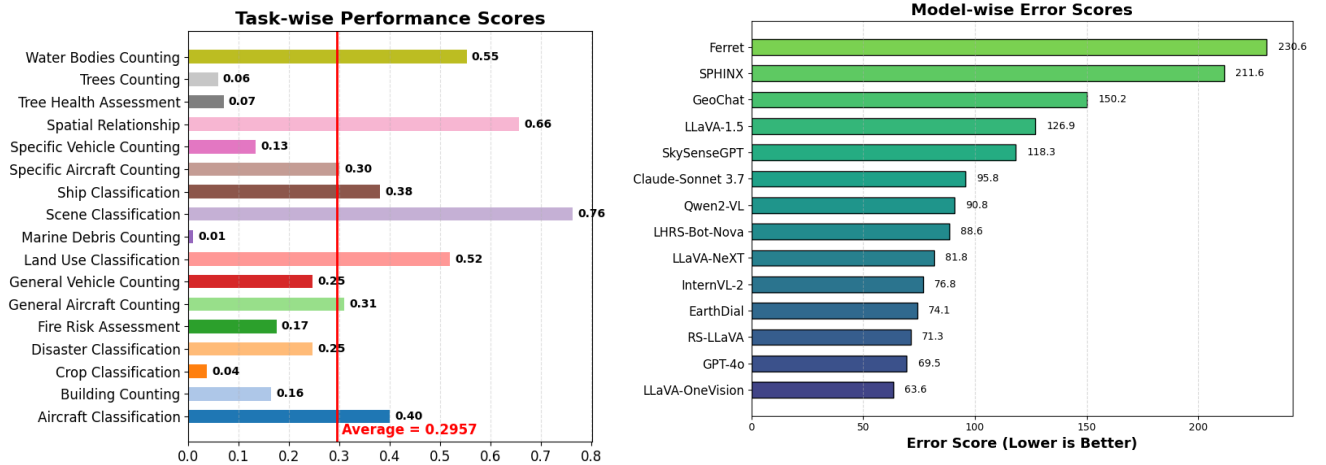


Figure A1. **Additional Comparisons.** (left) Claude-Sonnet 3.7 task-wise performance scores. (right) Model-wise error scores (lower is better).

ison highlights the importance of contextual information for accurate scene classification. In the last image of Fig. A2, ambiguous scenes such as the “ferry terminal” where all

models except EarthDial [20] failed, the misclassification is likely due to overlapping visual cues. The visual similarity between a harbor and a ferry terminal makes it challenging

for models to differentiate between these categories. For the counting tasks in Fig. A3, almost all models struggled, with wrongly estimating due to difficulty in differentiating objects in complex environments.

Fig. A4 shows that the models performed well in the first two images, probably due to familiar contextual clues. The “atago-class destroyer” and “small civil transport/utility” aircraft are common object types with distinct characteristics, making them easier for models to recognize. However, in the last two images, none of the models successfully identified the “murasame-class destroyer” or “garibaldi aircraft carrier” which are rarer categories. The failure is likely due to insufficient exposure to these specific classes in training datasets, coupled with the overlapping features of the objects that require advanced fine-grained recognition.

As shown in Fig. A5, models performed well on disaster assessment with relatively clear indicators, such as “fire” damage. For the second image, depicting “flooding”, Ferret and LHRS-Bot-Nova struggled. The third image depicts “tsunami” damage, characterized by disrupted layouts, scattered debris, and damaged buildings, which are often visually similar to flooding. Models may misclassify this due to overlapping features, and insufficient tsunami-specific training data. For the last image, only Qwen2VL identified the “seismic activity”, as others likely misclassified it due to overlapping features with “precipitation-related events”.

In Fig. A6, for the first image, a few models performed well because the objects are close to each other, easy to identify, and have minimal visual complexity. In the remaining images, models struggled because the objects were farther apart, making it harder to identify their spatial relationships. The cluttered environments and larger spatial gaps made it difficult for the models to accurately understand the relationships between the objects. Fig. A9 shows an aerial image alongside its ground truth caption and responses from various models. The ground truth provides a detailed and accurate description of the scene, while the model generated captions vary in capturing key elements such as urban & natural features, pathways, and architectural structures. This comparison highlights differences in model responses for image captioning tasks.

## S5. Multi-spectral

In this section, we compare how GPT-4o and Qwen2-VL [23] perform on Crop Type Classification and Land Use Classification tasks using RGB and multispectral (MS) data (Fig. A7). The models perform much better with RGB inputs because they are designed and trained specifically for RGB images. The accuracy drops significantly for multi-spectral data, especially in crop-type classification. To use MS data with these models, Sentinel-2 bands were combined into three channels sequentially to mimic RGB inputs. For land use classification, which depends more on

spatial patterns than detailed spectral information, the drop in performance is smaller. These results show the need for improved methods to adapt MS data for such tasks.

## S6. Bi-temporal vs. Multi-temporal

We compare bi-temporal and multi-temporal image classification performance for Crop Type and Land Use classification tasks (Fig. A8). Multi-temporal data outperforms bi-temporal data for land use classification, suggesting that more timestamps are sufficient to capture key temporal changes. For crop classification, multi-temporal inputs, reflect the improvement in GPT-4o and LLaVA-OneVision.

## S7. Geographical Analysis

In this section, we detail the geographic distribution of benchmarking datasets used in the studied geospatial tasks. It categorizes datasets into global/diverse datasets and regional/localized datasets. Global datasets provide extensive coverage with samples from over 100 countries or diverse regions worldwide. On the other hand, regional and localized datasets, are tailored to specific tasks. The map in Fig. A10 highlights that our benchmark is well represented across the globe.

## S8. Word Cloud

The breakdown in Fig. A11a leverages the word cloud as part of evaluating VLMs in geospatial tasks, with image captioning being one of the key areas of interest. The word cloud highlights terms commonly used in captions describing aerial or geospatial imagery. Words like “aerial”, “surrounding” and “residential” reflect spatial and contextual elements frequently addressed in such descriptions, while terms such as “harbor”, “ship”, “tennis court” and “greenery” represent specific features often observed in geospatial data. This provides a basis for understanding the capabilities and limitations of these models in capturing spatial relationships and identifying key features within geospatial tasks.

The word cloud in Fig. A11b shows the terms used in MCQs. Keywords such as “large vehicle”, “transport utility”, “harbor”, “bridge” and “small vehicle” emphasize important categories and features frequently mentioned in the questions. Additional terms like “aircraft carrier”, “runway”, “basketball court” and “helicopter” represent a blend of transportation, infrastructure, and activity-based elements often linked to geospatial data. The use of varied and domain-specific vocabulary ensures the MCQs encompass a broad spectrum of scenarios for testing model capabilities.

## References

- [1] Swedish Forest Agency. Forest damages – larch casebearer 1.0, 2021. National Forest Data Lab. Dataset. 2
- [2] Gerald Baier, Antonin Deschemps, Michael Schmitt, and Naoto Yokoya. Synthesizing optical and sar imagery from land cover maps and auxiliary raster data. *IEEE Transactions on Geoscience and Remote Sensing*, 60:1–12, 2022. 2
- [3] Yakoub Bazi, Laila Bashmal, Mohamad Mahmoud Al Rahhal, Riccardo Ricci, and Farid Melgani. Rs-llava: A large vision-language model for joint captioning and question answering in remote sensing imagery. *Remote Sensing*, 16(9): 1477, 2024. 1
- [4] Gong Cheng, Junwei Han, and Xiaoqiang Lu. Remote sensing image scene classification: Benchmark and state of the art. *Proceedings of the IEEE*, 105(10):1865–1883, 2017. 1, 2
- [5] Gong Cheng, Jiabao Wang, Ke Li, Xingxing Xie, Chunbo Lang, Yanqing Yao, and Junwei Han. Anchor-free oriented proposal generator for object detection. *CoRR*, abs/2110.01931, 2021. 1, 2
- [6] CSE499DeforestationSatellite. Deforestation-satellite-imagery dataset. <https://universe.roboflow.com/cse499deforestationssatellite/deforestation-satellite-imagery-335n4>. 2
- [7] Ilke Demir, Krzysztof Koperski, David Lindenbaum, Guan Pang, Jing Huang, Saikat Basu, Forest Hughes, Devis Tuia, and Ramesh Raskar. Deepglobe 2018: A challenge to parse the earth through satellite images. *CoRR*, abs/1805.06561, 2018. 2
- [8] Yanghua Di, Zhiguo Jiang, and Haopeng Zhang. A public dataset for fine-grained ship classification in optical remote sensing images. *Remote Sensing*, 13(4), 2021. 2
- [9] Ritwik Gupta, Richard Hosfelt, Sandra Sajeev, Nirav Patel, Bryce Goodman, Jigar Doshi, Eric T. Heim, Howie Choset, and Matthew E. Gaston. xbd: A dataset for assessing building damage from satellite imagery. *CoRR*, abs/1911.09296, 2019. 1, 2
- [10] Hannah Kerner, Snehal Chaudhari, Aninda Ghosh, Caleb Robinson, Adeel Ahmad, Eddie Choi, Nathan Jacobs, Chris Holmes, Matthias Mohr, Rahul Dodhia, Juan M. Lavista Ferres, and Jennifer Marcus. Fields of the world: A machine learning benchmark dataset for global agricultural field boundary segmentation, 2024. 1, 2
- [11] Ziyi Lin, Chris Liu, Renrui Zhang, Peng Gao, Longtian Qiu, Han Xiao, Han Qiu, Chen Lin, Wenqi Shao, Keqin Chen, et al. Sphinx: The joint mixing of weights, tasks, and visual embeddings for multi-modal large language models. *arXiv preprint arXiv:2311.07575*, 2023. 1
- [12] Jinglei Liu, Weixun Zhou, Haiyan Guan, and Wenzhi Zhao. Similarity learning for land use scene-level change detection. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 17:6501–6513, 2024. 2
- [13] Gabriel L. S. Machado, Edemir Ferreira, Keiller Nogueira, Hugo N. Oliveira, Pedro H. T. Gama, and Jefersson A. dos Santos. Airound and cv-brct: Novel multi-view datasets for scene classification. *CoRR*, abs/2008.01133, 2020. 1, 2
- [14] T. Nathan Mundhenk, Goran Konjevod, Wesam A. Sakla, and Kofi Boakye. A large contextual dataset for classification, detection and counting of cars with deep learning. In *Computer Vision – ECCV 2016*, pages 785–800, Cham, 2016. Springer International Publishing. 2
- [15] Daniele Rege Cambrin and Paolo Garza. Quakeset: A dataset and low-resource models to monitor earthquakes through sentinel-1. *Proceedings of the International IS-CRAM Conference*, 2024. 1, 2
- [16] Vivien Sainte Fare Garnot and Loic Landrieu. Panoptic segmentation of satellite image time series with convolutional temporal attention networks. *ICCV*, 2021. 1, 2
- [17] A. Shah, L. Thomas, and M. Maskey. Marine debris dataset for object detection in planetscope imagery, 2021. 2
- [18] Shuchang Shen, Sachith Seneviratne, Xinye Wanyan, and Michael Kirley. Firerisk: A remote sensing dataset for fire risk assessment with benchmarks using supervised and self-supervised learning, 2023. 2
- [19] Jacob Shermeyer, Thomas Hossler, Adam Van Etten, Daniel Hogan, Ryan Lewis, and Daeil Kim. Rareplanes dataset, 2020. 2
- [20] Sagar Soni, Akshay Dudhane, Hiyam Debary, Mustansar Fiaz, Muhammad Akhtar Munir, Muhammad Sohail Danish, Paolo Fraccaro, Campbell D Watson, Levente J Klein, Fahad Shahbaz Khan, et al. Earthdial: Turning multi-sensory earth observations to interactive dialogues. *arXiv preprint arXiv:2412.15190*, 2024. 2
- [21] Xian Sun, Peijin Wang, Zhiyuan Yan, F. Xu, Ruiping Wang, W. Diao, Jin Chen, Jihao Li, Yingchao Feng, Tao Xu, M. Weinmann, S. Hinz, Cheng Wang, and K. Fu. Fair1m: A benchmark dataset for fine-grained object recognition in high-resolution remote sensing imagery. *Isprs Journal of Photogrammetry and Remote Sensing*, 2021. 1, 2
- [22] Chintan Tundia, Rajiv Kumar, Om Damani, and G. Sivakumar. Fpcd: An open aerial vhr dataset for farm pond change detection. In *Proceedings of the 18th International Joint Conference on Computer Vision, Imaging and Computer Graphics Theory and Applications*, page 862–869. SCITEPRESS - Science and Technology Publications, 2023. 1, 2
- [23] Peng Wang, Shuai Bai, Sinan Tan, Shijie Wang, Zhihao Fan, Jinze Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, et al. Qwen2-vl: Enhancing vision-language model’s perception of the world at any resolution. *arXiv preprint arXiv:2409.12191*, 2024. 3
- [24] Gui-Song Xia, Xiang Bai, Jian Ding, Zhen Zhu, Serge Belongie, Jiebo Luo, Mihai Datcu, Marcello Pelillo, and Liangpei Zhang. Dota: A large-scale dataset for object detection in aerial images. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3974–3983, 2018. 1, 2
- [25] Haoxuan You, Haotian Zhang, Zhe Gan, Xianzhi Du, Bowen Zhang, Zirui Wang, Liangliang Cao, Shih-Fu Chang, and Yinfei Yang. Ferret: Refer and ground anything anywhere at any granularity. *arXiv preprint arXiv:2310.07704*, 2023. 1
- [26] Xiaokang Zhang, Weikang Yu, Man-On Pun, and Wenzhong Shi. Cross-domain landslide mapping from large-scale re-



mote sensing images using prototype-guided domain-aware progressive representation learning. *ISPRS Journal of Photogrammetry and Remote Sensing*, 197:1–17, 2023. [1](#), [2](#)

- [27] Weixun Zhou, Shawn D. Newsam, Congmin Li, and Zhenfeng Shao. Patternnet: A benchmark dataset for performance evaluation of remote sensing image retrieval. *CoRR*, abs/1706.03424, 2017. [2](#)
- [28] Xiao Xiang Zhu, Jingliang Hu, Chunping Qiu, Yilei Shi, Jian Kang, Lichao Mou, Hossein Bagheri, Matthias Haberle, Yuansheng Hua, Rong Huang, Lloyd Hughes, Hao Li, Yao Sun, Guichen Zhang, Shiyao Han, Michael Schmitt, and Yuanyuan Wang. So2sat lc42: A benchmark data set for the classification of global local climate zones [software and data sets]. *IEEE Geoscience and Remote Sensing Magazine*, 8(3):76–89, 2020. [1](#), [2](#)

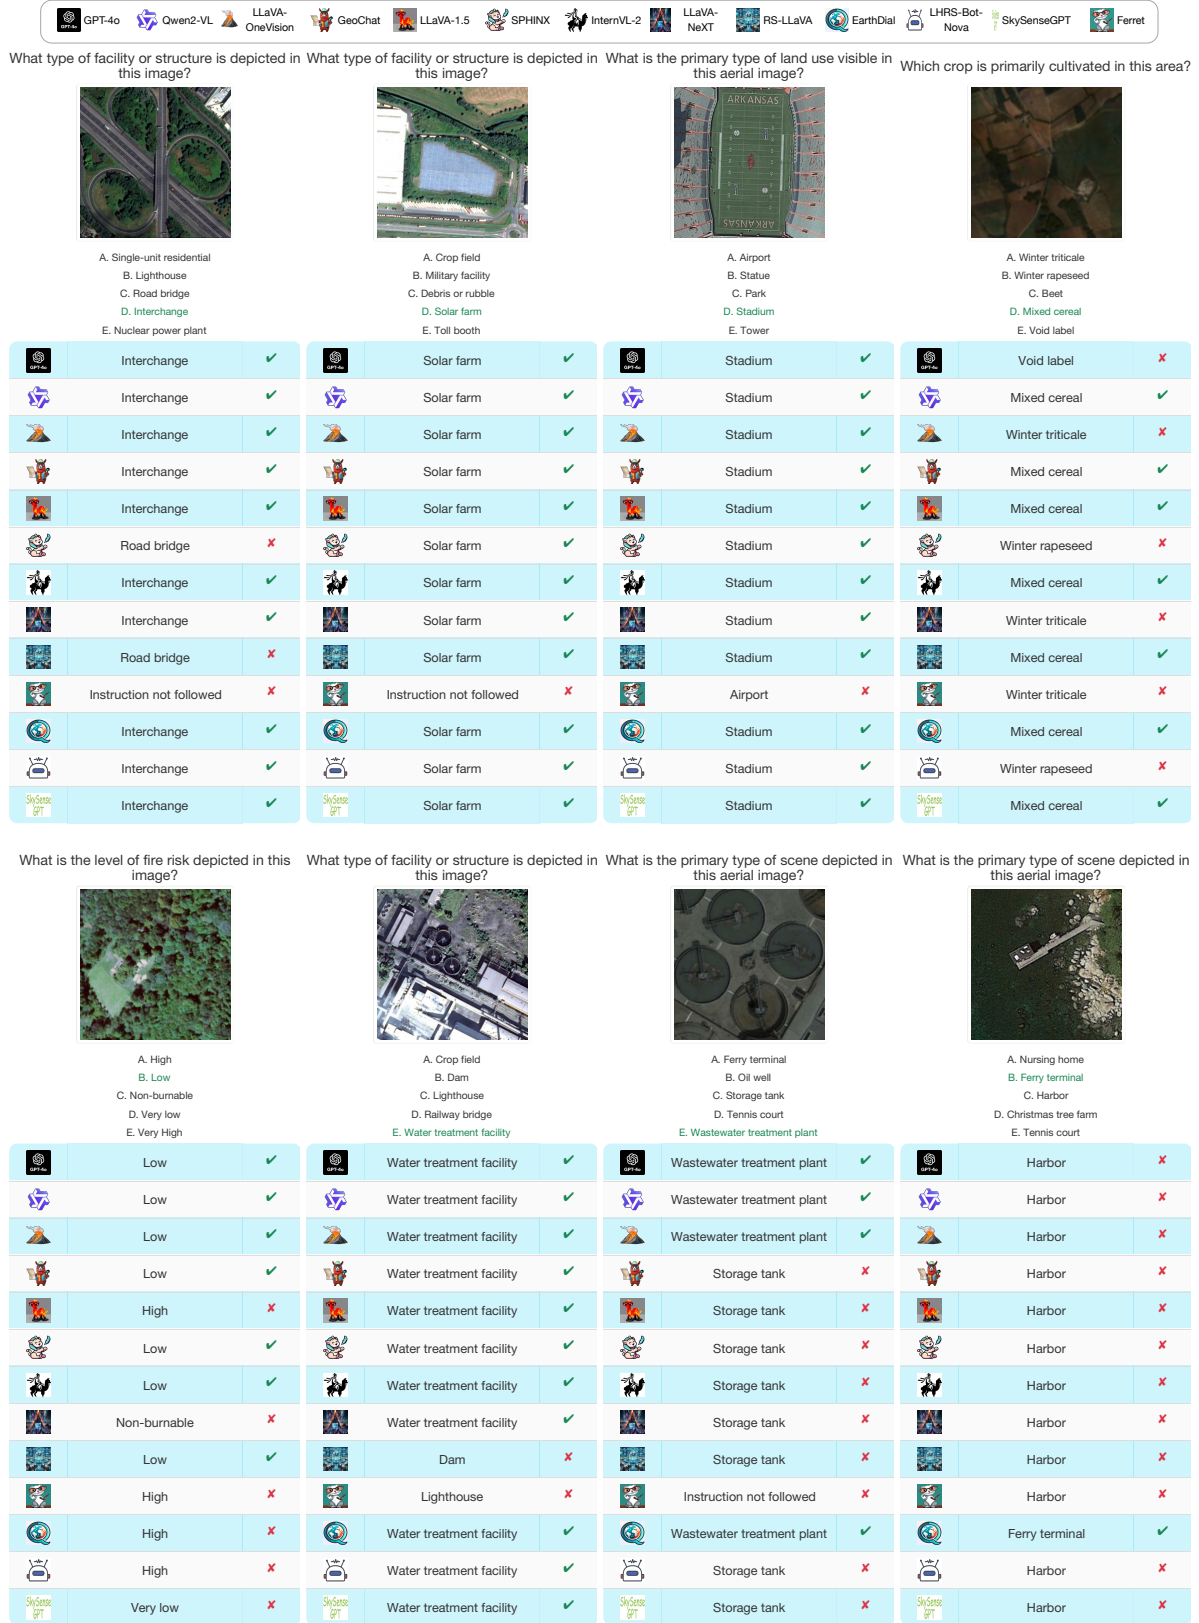


Figure A2. Scene Understanding: This illustrates model performance on geospatial scene understanding tasks, highlighting successes in clear contexts and challenges in ambiguous scenes. The results emphasize the importance of contextual reasoning and addressing overlapping visual cues for accurate classification.



Figure A3. Counting: The figure showcases model performance on counting tasks.

<div>  GPT-4o          Qwen2-VL          LLaVA-OneVision          GeoChat          LLaVA-1.5          SPHINX          InternVL-2          LLaVA-NeXT          RS-LLaVA          EarthDial          LHRS-Bot-Nova          SkySenseGPT          Ferret       </div>															
What type of ship is visible in this image?				What type of aircraft is visible in this image?				What type of ship is visible in this image?				What type of ship is visible in this image?			
A. Atago-class destroyer B. Type 45 destroyer C. Mega yacht D. Civil yacht E. Mistral-class amphibious assault ship				A. Military Bomber B. Medium Civil Transport/Utility C. Military Trainer D. Small Civil Transport/Utility E. Large Civil Transport/Utility				A. Murasame-class destroyer B. Kongo-class destroyer C. Civil yacht D. Arleigh Burke-class destroyer E. Kitty Hawk-class aircraft carrier				A. Civil yacht B. INS Vikramaditya carrier C. Atago-class destroyer D. Garibaldi aircraft carrier E. Kitty Hawk-class aircraft carrier			
	Atago-class destroyer	✓			Small Civil Transport/Utility	✓			Kitty Hawk-class aircraft carrier	✗			INS Vikramaditya carrier	✗	
	Atago-class destroyer	✓			Small Civil Transport/Utility	✓			Kitty Hawk-class aircraft carrier	✗			Kitty Hawk-class aircraft carrier	✗	
	Atago-class destroyer	✓			Small Civil Transport/Utility	✓			Kitty Hawk-class aircraft carrier	✗			INS Vikramaditya carrier	✗	
	Atago-class destroyer	✓			Small Civil Transport/Utility	✓			Kitty Hawk-class aircraft carrier	✗			Kitty Hawk-class aircraft carrier	✗	
	Atago-class destroyer	✓			Small Civil Transport/Utility	✓			Kitty Hawk-class aircraft carrier	✗			Kitty Hawk-class aircraft carrier	✗	
	Type 45 destroyer	✗			Small Civil Transport/Utility	✓			Arleigh Burke-class destroyer	✗			INS Vikramaditya carrier	✗	
	Atago-class destroyer	✓			Small Civil Transport/Utility	✓			Kitty Hawk-class aircraft carrier	✗			Kitty Hawk-class aircraft carrier	✗	
	Atago-class destroyer	✓			Small Civil Transport/Utility	✓			Arleigh Burke-class destroyer	✗			Kitty Hawk-class aircraft carrier	✗	
	Instruction not followed	✗			Small Civil Transport/Utility	✓			Kitty Hawk-class aircraft carrier	✗			Instruction not followed	✗	
	Atago-class destroyer	✓			Instruction not followed	✗			Kitty Hawk-class aircraft carrier	✗			INS Vikramaditya carrier	✗	
	INS Vikramaditya aircraft carrier	✗			Military Trainer	✗			Arleigh Burke-class destroyer	✗			INS Vikramaditya aircraft carrier	✗	
	INS Vikramaditya aircraft carrier	✗			Medium Civil Transport/Utility	✗			Arleigh Burke-class destroyer	✗			INS Vikramaditya aircraft carrier	✗	
	Atago-class destroyer	✓			Small Civil Transport/Utility	✓			Arleigh Burke-class destroyer	✗			Atago-class destroyer	✗	

Figure A4. Object Classification: The figure highlights model performance on object classification, showing success with familiar objects like the “atago-class destroyer” and “small civil transport/utility” aircraft. However, models struggled with rarer objects like the “murasame-class destroyer” and “garibaldi aircraft carrier” indicating a need for improvement on less common classes and fine-grained recognition.



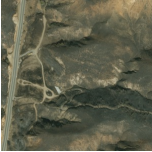



GPT-4o          Qwen2-VL          LLaVA-OneVision          GeoChat          LLaVA-1.5          SPHINX          InternVL-2          LLaVA-NeXT          RS-LLaVA          EarthDial          LHRS-Bot-Nova          SkySenseGPT          Ferret											
What type of disaster is responsible for the visible damage in this image? 			What type of disaster is responsible for the visible damage in this image? 			What type of disaster is responsible for the visible damage in this image? 			What was the primary trigger for this landslide? 		
A. volcano B. flooding C. tsunami D. fire E. earthquake			A. earthquake B. tsunami C. volcano D. wind E. flooding			A. flooding B. volcano C. earthquake D. wind E. tsunami			A. Soil Erosion B. Human Activities C. Snow and Glacier Melting D. Seismic Activity E. Precipitation-Related Events		
	fire	✓		flooding	✓		earthquake	✗		Precipitation-Related Events	✗
	fire	✓		flooding	✓		earthquake	✗		Seismic Activity	✓
	fire	✓		flooding	✓		flooding	✗		Precipitation-Related Events	✗
	fire	✓		flooding	✓		earthquake	✗		Precipitation-Related Events	✗
	fire	✓		flooding	✓		flooding	✗		Soil Erosion	✗
	earthquake	✗		flooding	✓		flooding	✗		Human Activities	✗
	fire	✓		flooding	✓		flooding	✗		Precipitation-Related Events	✗
	fire	✓		flooding	✓		earthquake	✗		Precipitation-Related Events	✗
	fire	✓		flooding	✓		earthquake	✗		Precipitation-Related Events	✗
	Instruction not followed	✗		Instruction not followed	✗		Instruction not followed	✗		Instruction not followed	✗
	volcano	✗		flooding	✓		tsunami	✓		Soil Erosion	✗
	volcano	✗		tsunami	✗		wind	✗		Snow and Glacier Melting	✗
	fire	✓		flooding	✓		earthquake	✗		Soil Erosion	✗

Figure A5. Event Detection: Model performance on disaster assessment tasks, with success in scenarios like 'fire' and 'flooding' but challenges in ambiguous cases like 'tsunami' and 'seismic activity'. Misclassifications highlight limitations in contextual reasoning and insufficient exposure on overlapping disaster features.

<div>  GPT-4o            Qwen2-VL            LLaVA-OneVision            GeoChat            LLaVA-1.5            SPHINX            InternVL-2            LLaVA-NeXT            RS-LLaVA            EarthDial            LHRS-Bot-Nova            SkySenseGPT            Ferret         </div>			
What is the relationship between object in green box and object in red box in this image?			
A. A large vehicle is aligned with the bridge. B. A helicopter is positioned next to the helipad. C. A overpass leads to the golf field. D. A tennis court is beside the basketball court. E. A runway connects to the airport.			
	C.	✗	
	C.	✗	
	C.	✗	
	C.	✗	
	C.	✗	
	C.	✗	
	C.	✗	
	C.	✗	
	C.	✗	
	C.	✓	
	C.	✓	
	C.	✓	
What is the relationship between object in green box and object in red box in this image?			
A. A helicopter is below the airport. B. A small vehicle is to the left of the large vehicle. C. A ship is to the right of the a small-vehicle. D. A large-vehicle is positioned in front of the bridge. E. A helicopter is above the helipad.			
	B.	✗	
	E.	✗	
	B.	✗	
	D.	✗	
	D.	✗	
	C.	✗	
	A.	✗	
	B.	✗	
	E.	✗	
	C.	✓	
	C.	✓	
	D.	✗	
What is the relationship between object in green box and object in red box in this image?			
A. An a350 is aligned with the runway. B. A plane is parked near the runway. C. A large-vehicle is driving by the storage-tank. D. A large-vehicle is moving towards the storage-tank. E. A large-vehicle is moving away from the a roundabout.			
	C.	✗	
	E.	✗	
	B.	✗	
	D.	✗	
	B.	✗	
	B.	✗	
	E.	✗	
	B.	✗	
	D.	✗	
	E.	✓	
	B.	✗	
	B.	✗	
What is the relationship between object in green box and object in red box in this image?			
A. A small-vehicle is driving by the bridge. B. A plane is parked near the runway. C. A helicopter is moving away from the helipad. D. A helicopter is positioned beside the helipad. E. A large-vehicle is moving towards the a roundabout.			
	C.	✗	
	D.	✗	
	D.	✗	
	D.	✗	
	D.	✗	
	D.	✗	
	A.	✗	
	D.	✗	
	A.	✗	
	C.	✗	
	A.	✗	
	D.	✗	

Figure A6. Spatial Relations: The figure demonstrates model performance on spatial relationship tasks, with success in close-object scenarios and struggles in cluttered environments with distant objects.

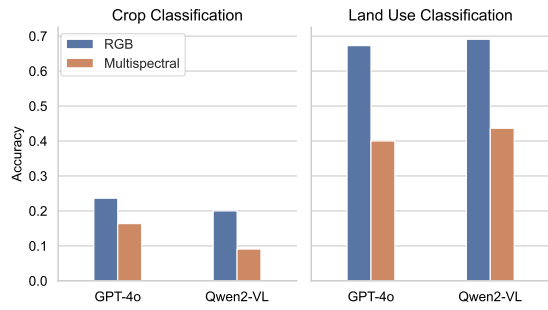


Figure A7. It compares RGB and multispectral performance for Crop Type and Land Use Classification, showing a performance drop in multispectral accuracy.

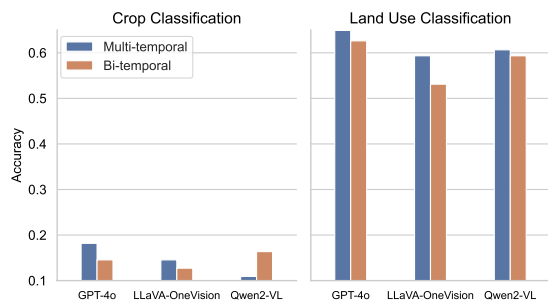
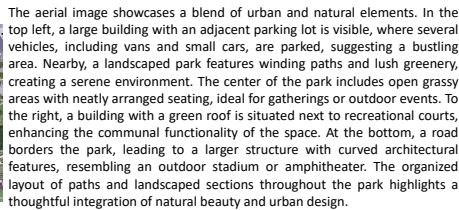


Figure A8. Bi-temporal and Multi-temporal performance



The image is an aerial view of a landscaped area and surrounding structures, likely a botanical garden or park. Key features in this image include: **1. Pathways:** There are several winding pathways throughout the landscape, designed for walking and exploration. **2. Gardens and Spaces:** The area is divided into various plant beds, trees, and open lawn areas. **3. Landmarks:** There are likely fountains, bridges, and other ornamental features. **4. Accessibility:** The pathways appear to be well-maintained and accessible, suggesting the presence of an outdoor amphitheater, often used for seminars or performances. **5. Buildings and Structures:** There is a building with a turquoise roof, possibly a visitor center or administrative building. **6. Parking:** A parking lot is visible, bordered by a white road. **7. Surrounding Area:** The landscape area is bordered by a road and other structures, indicating accessibility for visitors arriving by car. **8. Nearby Roads:** A road runs along the edge of the property, providing access to the park and parking areas. **9. Sports Field:** There is a small field or court area, likely used for sports or recreational activities. Overall, this image depicts a well-maintained recreational or botanical space, designed for public use and enjoyment.

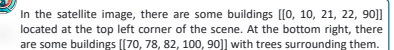
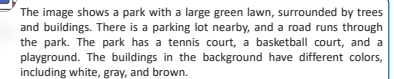
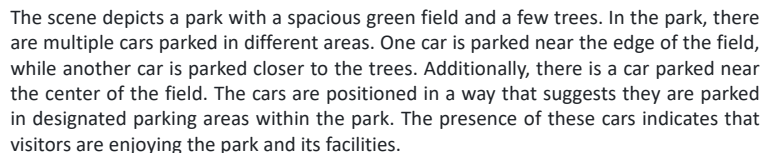
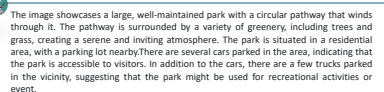
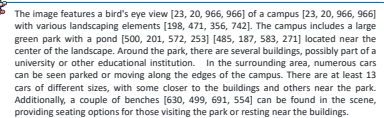
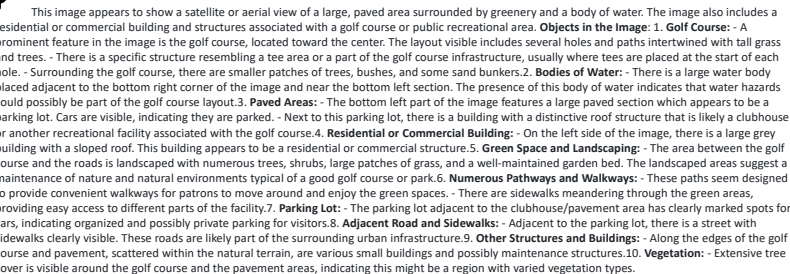
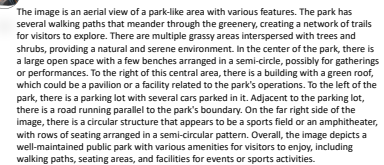
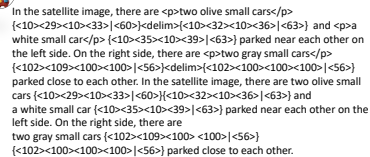
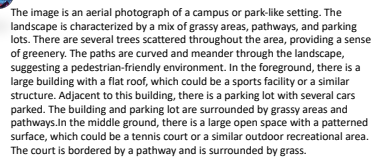
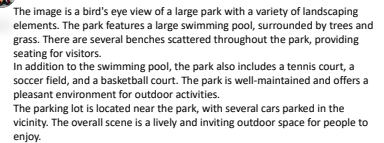
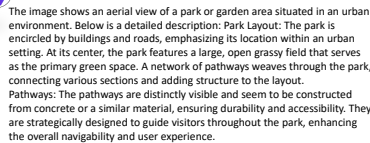


Figure A9. Image Captioning: Example response of different models.



