CAN VISION-LANGUAGE MODELS THINK FROM THE SKY? UNIFYING UAV REASONING AND GENERATION

Anonymous authors
Paper under double-blind review

ABSTRACT

Vision-Language Models have achieved strong progress in ground-view visual understanding, yet they remain brittle in high-altitude Unmanned Aerial Vehicle scenes, where objects are tiny and densely packed, textures are repetitive, and top-down orientations are ambiguous. We introduce UAVReason, a large-scale UAV-native dataset and evaluation suite for studying unified aerial reasoning and generation under this nadir-view domain shift. UAVReason aligns RGB imagery, depth maps, semantic segmentation masks, captions, and question-answer pairs within a consistent aerial domain. It contains 23.6K captioned frames, 273K VQA pairs including 68.2K two-frame temporal questions, and 188.8K cross-modal generation samples across RGB, depth, and segmentation modalities. We further adapt UAVReason-Bagel as a unified understanding-and-generation baseline that jointly optimizes language reasoning and dense visual generation objectives. Experiments show that general-purpose VLMs and off-the-shelf unified generators struggle with UAV-native grounding, while UAVReason-Bagel substantially improves over its pretrained counterpart, increasing VQA-1F F1 from 0.394 to 0.711, VQA-2F F1 from 0.427 to 0.822, and heading-aware VQA F1 from 0.798 to 0.973. For generation, it improves segmentation mIoU to 0.143 and reduces KID from 0.078 to 0.048 for depth-segmentation-text-conditioned RGB synthesis. More importantly, our ablations reveal a bidirectional synergy between synthesis and reasoning. Dense generation objectives improve temporal semantic consistency, while language-level reasoning regularizes sparse-condition image synthesis. These results suggest that unified reasoning and generation provide effective geometry-aware structural priors for physically grounded aerial intelligence. All data, code, and evaluation tools will be released.

1 INTRODUCTION

Vision-Language Models (VLMs) have reshaped visual understanding by aligning images and text at scale, enabling open-ended recognition, grounding, and reasoning across diverse tasks (Radford et al., 2021; OpenAI, 2023; Liu et al., 2023; Cheng et al., 2024). Such capabilities are increasingly central to embodied agents, including Unmanned Aerial Vehicles (UAVs), mobile robots, and autonomous systems, that perceive, reason, and act in complex, dynamic environments. However, today's VLMs remain largely ground-centric and often degrade in aerial scenarios, since most pretraining corpora are dominated by human-centric, ground-view imagery (Lin et al., 2014; Young et al., 2014). When deployed on high-altitude UAV platforms, models face a pronounced domain shift: objects appear tiny and densely packed; textures are repetitive (*e.g.*, rooftops and foliage); and the absence of a stable horizon yields ambiguous top-down orientation.

As summarized in Table 1, existing aerial datasets primarily focus on individual tasks such as detection (Zhu et al., 2018; Xia et al., 2018), tracking and 3D perception (Du et al., 2018; Ye et al., 2024), and segmentation (Lyu et al., 2020; Nigam et al., 2018; Waqas Zamir et al., 2019). While recent efforts explore monocular depth (Pan et al., 2024; Florea et al., 2021), embodied intelligence in urban spaces (Zhao et al., 2025), or dedicated UAV VQA (Li et al., 2024b; Wang et al., 2024), they typically lack integrated geometric supervision and cross-modal generation within a unified protocol. This fragmentation results in isolated training of understanding and generation tasks. Consequently, models fail to benefit from the complementary geometric cues shared across

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

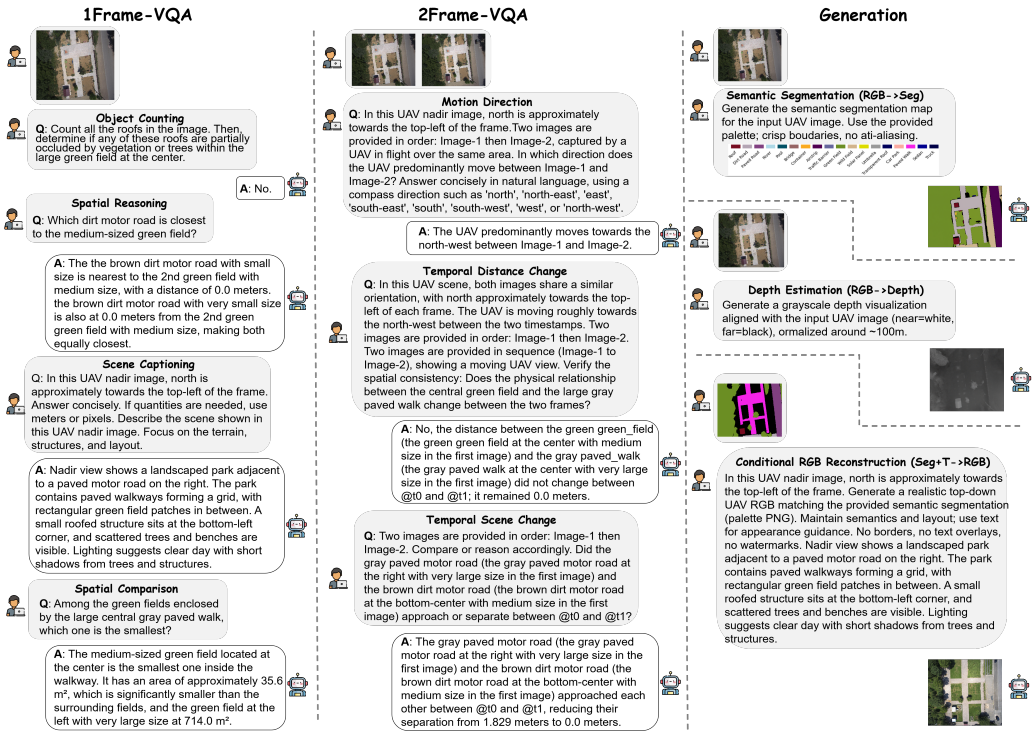

Figure 1: **UAVReason**, a unified benchmark for nadir-view spatio-temporal reasoning and cross-modal generation. **1Frame-VQA (Left):** Single-frame reasoning with orientation cues (*e.g.*, north), requiring object counting, scene captioning, spatial reasoning, and comparison ability. **2Frame-VQA (Middle):** Two-frame temporal reasoning over aligned viewpoints, probing motion direction and temporal changes in relative distance and scene changes. **Generation (Right):** Pixel-level conditional synthesis with diverse inputs, covering semantic segmentation (RGB→Seg), depth estimation (RGB→Depth), and conditional RGB reconstruction from segmentation and text prompts (Seg+T→RGB).

tasks, which are crucial for learning robust representations in nadir-view UAV scenarios. In this paper, we argue that UAV-native capabilities require jointly modeling "reasoning" and "generation". Perceiving small, ambiguous overhead objects requires precise pixel-level spatial reasoning beyond semantic alignment. For models that can answer VQA questions, those that can additionally generate segmentation maps from text exhibit a richer form of scene understanding: they capture not only what an object is, but also where it is and what its spatial shape looks like.

To fill this gap, we introduce **UAVReason**, a unified large-scale nadir-view benchmark for *multi-modal UAV scene reasoning and generation* (see Figure 1). Built from a single, high-fidelity UAV platform to ensure domain consistency, UAVReason consolidates dense RGB–Depth–Segmentation supervision with two complementary tracks: **Reasoning Track**: over **273,000** VQA pairs including **23,600** single frames each with detailed scene captions, and **68,223** temporal 2-Frame sequences. **Generation Track**: **188,800** cross-modal samples spanning tasks, *e.g.*, RGB→Segmentation, Depth/Segmentation→RGB, and text-conditioned generation. Our main contributions are as follows:

- **UAV-Native Evaluation Suite.** We introduce UAVReason, a large-scale unified evaluation suite that aligns RGB imagery, dense geometry including depth and segmentation, and language annotations to study multimodal models under nadir-view domain shift.
- **Unified Aerial Protocol.** UAVReason integrates 23.6K captioned frames, 273K VQA pairs including 68.2K two-frame temporal questions, and 188.8K cross-modal generation samples, enabling joint evaluation of 22 spatio-temporal reasoning types and geometrically grounded visual synthesis.
- **Reasoning and Generation Synergy.** We establish UAVReason-Bagel as a unified understanding-and-generation baseline and use it to examine the interaction between synthesis and reasoning, showing that dense generation objectives improve temporal semantic consistency while language-level reasoning regularizes sparse-condition image synthesis.

Table 1: **Comparison with existing UAV, aerial, and remote sensing benchmarks.** We summarize representative datasets by category and by supported supervision and tasks, including meta properties (Typ, Vw), dense geometry signals (Dep, Seg), reasoning capabilities (VQA, Tmp, Cap), and generation support (CMG), together with dataset scale and statistics. *Legend:* **Typ** denotes real (R) or synthetic (S); **Vw** denotes nadir (N), oblique (O), or mixed (M) views; **Dep** and **Seg** indicate dense depth and segmentation annotations; **VQA**, **Tmp**, and **Cap** indicate visual question answering, temporal reasoning, and captioning; **CMG** indicates cross-modal generation.

Category	Dataset	Meta		Dense Geom.		Reasoning			Gen.	Scale & Stats
		Typ	Vw	Dep	Seg	VQA	Tmp	Cap	CMG	
Detection / Segmentation	VisDrone (Zhu et al., 2018)	R	M	✗	✗	✗	✗	✗	✗	263 clips, 10.2k imgs
	UAVDT (Du et al., 2018)	R	M	✗	✗	✗	✗	✗	✗	∼80k frames
	UAVid (Lyu et al., 2020)	R	O	✗	✓	✗	✗	✗	✗	30 seqs / 300 labeled 4K imgs
	AeroScapes (Nigam et al., 2018)	R	M	✗	✓	✗	✗	✗	✗	3.3k images
	DOTA (Xia et al., 2018)	R	M	✗	✗	✗	✗	✗	✗	2,806 imgs / 188,282 inst.
	iSAID (Waqas Zamir et al., 2019)	R	M	✗	✓	✗	✗	✗	✗	2,806 imgs / 655,451 inst.
Geometry / Multi-sensor	UAVScenes (Wang et al., 2025a)	R	M	✓	✓	✗	✗	✗	✗	>120k frames (RGB+L)
	WildUAV (Florea et al., 2021)	R	M	✓	✗	✗	✗	✗	✗	>1.5k high-res imgs
	MoNA Bench (Pan et al., 2024)	R	M	✓	✗	✗	✗	✗	✗	MAV nav. benchmarks
Synthetic	Mid-Air (Fonder & Van Droogenbroeck, 2019)	S	M	✓	✓	✗	✗	✗	✗	Sim. trajectories
	TartanAir (Wang et al., 2020)	S	M	✓	✓	✗	✗	✗	✗	Sim. navigation
	SynDrone (Rizzoli et al., 2023)	S	M	✓	✓	✗	✗	✗	✗	Physics simulation
	University-1652 (Zheng et al., 2020)	R	M	✗	✗	✗	✗	✗	✗	72 univs / 1652 bldgs; Google Earth
VQA / RS	RSVQA-LR (Lobry et al., 2020)	R	N	✗	✗	✓	✗	✗	✗	772 imgs, 77k QA
	RSVQA-HR (Lobry et al., 2020)	R	N	✗	✗	✓	✗	✗	✗	10.7k imgs, 1.07M QA
	RSICD (Lu et al., 2018)	R	M	✗	✗	✗	✗	✓	✗	10.9k imgs, 5 Cap/img
	HRVQA(Li et al., 2024b)	R	N	✗	✗	✓	✗	✗	✗	53k imgs, 1M QA
	EarthVQA (Wang et al., 2024)	R	N	✗	✓	✓	✗	✗	✗	6k imgs, 208k QA
	RSIVQA (Zheng et al., 2022)	R	M	✗	✗	✓	✗	✗	✗	37k imgs, 111k QA
	VRSBench (Li et al., 2024c)	R	M	✗	✗	✓	✗	✓	✗	29k imgs, Cap+QA
Unified	**UAVReason**	**R**	**N**	✓	✓	✓	✓	✓	✓	**23.6k frames with Caps; 68.2k Seq 273k VQA; 188.8k Gen Pairs**

2 Related Work

UAV Vision Benchmarks. The UAV aerial imagery landscape has evolved from detection-oriented benchmarks, *e.g.*, VisDrone (Zhu et al., 2018) and DOTA (Xia et al., 2018), to tracking and 3D perception benchmarks (Du et al., 2018; Ye et al., 2024; Naik et al., 2024), dense prediction datasets, *e.g.*, UAVid (Lyu et al., 2020), and high-level reasoning, *e.g.*, HRVQA (Li et al., 2024b). However, these datasets remain largely task-specific, isolating perception from reasoning and lacking integrated evaluation of spatio-temporal VQA with pixel-level generation. In contrast, UAVReason overcomes these limitations by unifying perception, reasoning, and generative capabilities on a consistent nadir-view platform.

Vision-Language Models. Standard VLMs, *e.g.*, LLaVA-OneVision (Li et al., 2025) and Qwen2.5-VL (Bai et al., 2025b), achieve remarkable general-purpose understanding (Li et al., 2024a; Chen et al., 2024). However, they suffer from pronounced ground-centric bias when transferred to nadir views. More critically, their discriminative nature restricts outputs to textual responses, thereby hindering the pixel-level synthesis (*e.g.*, depth/segmentation) required for dense spatial grounding in aerial environments. To transcend text-only outputs, recent works, *i.e.*, Janus (Wu et al., 2025), OmniGen (Xiao et al., 2025), and Bagel (Deng et al., 2025) have integrated generative capabilities into transformer-based models. Despite their promise, these general-domain models lack supervision from paired dense geometry under aerial viewpoints. Consequently, they remain vulnerable to hallucinations and unstable grounding when deployed in aerial scenarios.

Aerial-adapted VLMs. Recently, some approaches, *e.g.*, GeoChat (Kuckreja et al., 2024) and EarthGPT (Zhang et al., 2024), mitigate the domain shift via instruction tuning but lack support for controllable cross-modal generation, limiting their ability to enforce structural consistency. In contrast, we fine-tune a unified architecture on UAVReason, establishing a baseline that couples reasoning with pixel-level synthesis. We show that incorporating generative objectives provides effective structural priors, leading to improved spatial perception.



Figure 2: **Overview of UAVReason benchmark tasks and I/O protocols. Top (Reasoning):** Language-centric reasoning on nadir-view UAV imagery, including single-frame spatial queries (*e.g.*, counting/comparison and referring questions with orientation cues), global scene captioning, and two-frame temporal reasoning (motion direction and cross-time relational verification). **Middle (Geometry prediction):** Dense pixel-level supervision aligned to the same UAV frames, covering RGB→Depth (grayscale depth visualization, near-to-far) and RGB→Seg (palette PNG with crisp boundaries). **Bottom (Conditional generation):** Controllable RGB synthesis conditioned on geometric modalities and text, including Depth+Text→RGB, Seg+Text→RGB, and Depth+Seg+Text→RGB. Together, these tasks unify semantic reasoning with geometric/structural constraints, supporting unified training and evaluation of grounding fidelity, geometric consistency, and controllable synthesis in the UAV domain.

3 UAVREASON BENCHMARK

3.1 DATASET CONSTRUCTION

Different from existing benchmarks in Table 1, UAVReason is constructed as a unified benchmark for spatio-temporal reasoning and cross-modal generation in nadir-view UAV scenarios (see Figure 2). To support such capabilities, the benchmark is designed to provide not only semantic annotations, but also precise, metric-level geometric grounding across space and time. As shown in Figure 3, we develop an automated pipeline that systematically transforms raw multi-sensor UAV data into geometry-aware representations suitable for reasoning and generation. In particular, our construction leverages UAVScenes (Wang et al., 2025a) as a high-fidelity data source, which offers synchronized camera and LiDAR streams with accurate sensor calibration across diverse real-world environments, including towns, valleys, and airports. These properties enable dense geometric supervision to be derived directly from the underlying 3D environment, without reliance on manual semantic annotation, while preserving physical consistency.

Depth Rendering. We first obtain dense, pixel-aligned depth as the geometric foundation for subsequent reasoning and question generation. However, raw LiDAR point clouds are inherently sparse in the image plane. To address this sparsity, we render dense depth maps from the high-fidelity 3D environment meshes registered to the sensor trajectories in UAVScenes. We configure an off-screen renderer with the exact camera intrinsics and extrinsic pose for each frame. The renderer's Z-buffer provides accurate ground-truth depth supervision.

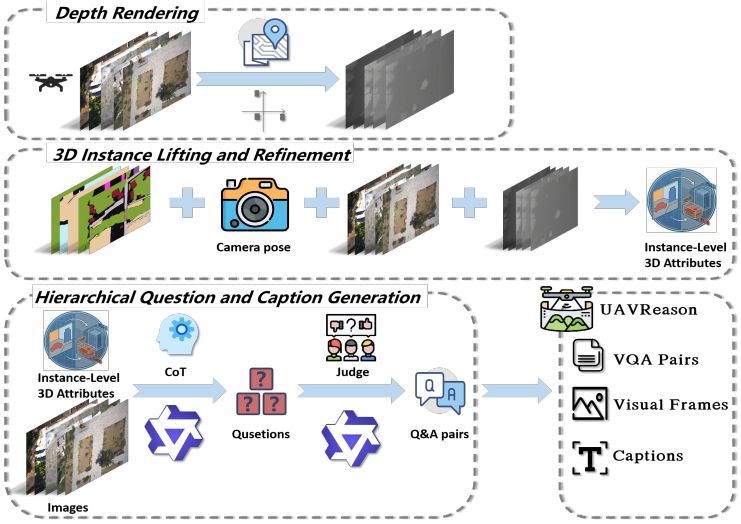

Figure 3: **UAVReason construction pipeline.** We apply the off-the-shelf tools to generate dense, physically grounded annotations for the nadir-view UAV imagery, including depth rendering, 3D instance lifting and refinement, and hierarchical question and caption generation.

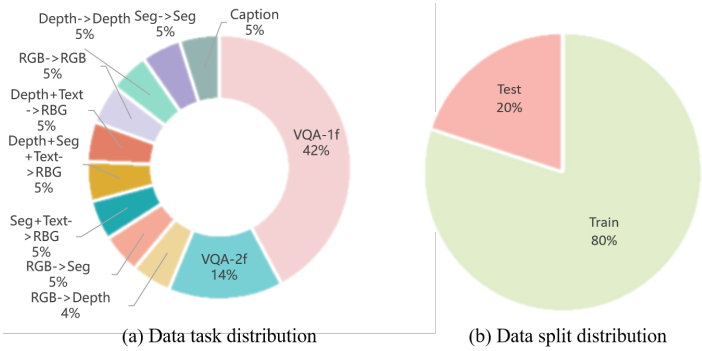

(a) Data task distribution (b) Data split distribution

Figure 4: Task and Data Distribution in UAVReason.

3D Instance Lifting and Refinement. With dense depth maps available, we derive explicit object-level ground truth to define instance identity and 3D localization, enabling geometry-grounded question–answer pairs on object identity, location, and quantitative spatial relations. We first extract candidate object regions via connected-component analysis on semantic segmentation masks. In nadir-view imagery, occlusions and limited resolution often merge distinct objects in the image plane. To recover accurate instance boundaries, we leverage depth discontinuities to split connected regions into spatially coherent components, followed by refinement to suppress noisy detections. Specifically, we employ a hybrid non-maximum suppression with a tightness criterion to eliminate erroneous enclosures of smaller instances within larger regions. Each refined 2D instance mask is then lifted to 3D by back-projecting its pixels into the world coordinate system using calibrated camera intrinsics and rendered depth maps. From the resulting 3D point set, we fit an oriented bounding box (OBB) and compute metric attributes, including the 3D centroid and orientation. These geometry-grounded object representations serve as the foundation for QA pairs involving object counting, distance estimation, comparative spatial relations, and other quantitative reasoning tasks.

Hierarchical Question and Caption Generation. Building on the geometry-grounded 3D object instances, we construct questions and captions via a three-step pipeline: **(1) Contextual Enrichment.** For each frame, we first generate a global scene description to establish overall context. At the instance level, we retrieve object attributes and spatial relations directly from the associated 3D representations, ensuring all references remain faithful to the underlying scene geometry. **(2) Taxonomy-Guided Question Generation.** To promote diverse reasoning demands, we create questions spanning multiple levels: basic object perception, metric spatial reasoning, and multi-step logical inference. By leveraging persistent 3D instance identities across frames, we further generate temporal questions that probe object motion, trajectory, and state evolution over time. **(3) Programmatic Answer**

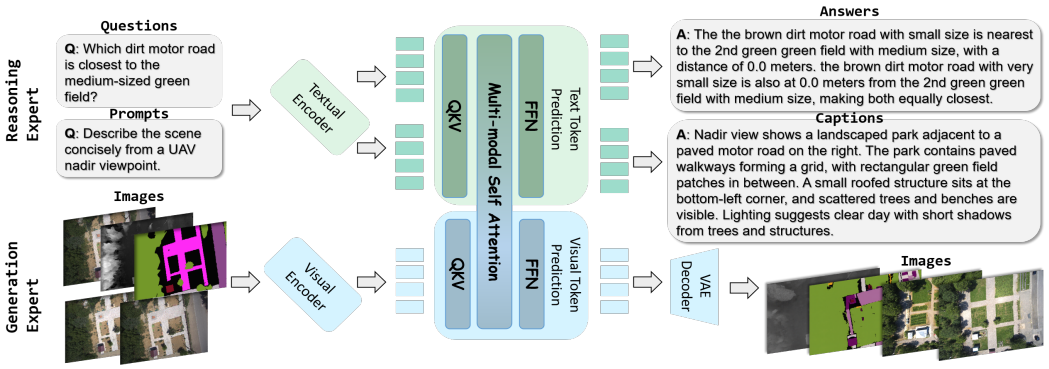

Figure 5: **Unified multi-task baseline for UAVReason.** We adopt a shared multi-modal transformer backbone with two task-specialized experts: an **Understanding Expert** (top) optimized for next-token prediction over language outputs, and a **Generation Expert** (bottom) optimized for Rectified-Flow velocity prediction over continuous visual latents. Text is tokenized and fused with visual features via multi-modal self-attention, while expert-specific QKV projections and FFNs specialize computation for understanding versus generation. The understanding branch consumes UAV RGB frames with questions (and captioning prompts) to produce textual outputs, whereas the generation branch conditions on RGB/depth/segmentation (and optional text) to synthesize target modalities. This design enables a single backbone to jointly learn language reasoning and pixel-level generation under a consistent UAV domain.

Construction. For each question, the answer is obtained through explicit geometric computation on the 3D ground truth. A *Planner* parses the question into a symbolic execution program (*e.g.*, `calc_distance`, `compare_count`), which is evaluated by a deterministic *Executor* operating directly on the 3D ground truth. Finally, a *Writer* renders the computed result in natural language. This programmatic approach provides metric-grounded, reproducible, and consistent supervision for spatial and temporal reasoning evaluation.

3.2 Task Description

Task Formulation. Our benchmark includes two complementary tracks that share aligned RGB, depth, and segmentation signals. As shown in Figure 4, UAVReason contains 23,600 unique frames, with 19,901 for training and 3,699 for testing, sampled from 20 annotated UAV flight sequences. To reduce leakage between highly correlated neighboring frames, we adopt a sequence-level split rather than a frame-level split, reserving four complete flight sequences for testing.

Reasoning Track. To probe semantic grounding under the challenging nadir viewpoint, we provide rich language supervision organized into three core evaluation settings. UAVReason contains **273,000** VQA pairs, comprising: (i) **204,777** single-frame questions targeting static spatial and compositional reasoning (**VQA-1F**); and (ii) **68,223** two-frame questions designed for temporal dynamics. The latter is further stratified into standard temporal reasoning (**VQA-2F**) and harder navigational reasoning requiring implicit heading cues (**VQA-2F$_{\text{ihead}}$**). In addition, each of the **23,600** frames is annotated with a detailed scene-level caption that summarizes global context and salient object relations. The VQA taxonomy spans **22** sub-categories grouped into 4 major axes, such as *Object Counting*, *Spatial Reasoning*, *Spatial Comparison*, and *Temporal Reasoning*, enabling fine-grained diagnosis of diverse reasoning abilities.

Generation Track. To enforce pixel-level grounding, we pair each RGB frame with precisely aligned dense depth maps and semantic segmentation masks, supporting eight multimodal generation tasks with a total of **188,800** samples. These tasks are grouped into three families. First, *cross-modal translation* (RGB→Depth and RGB→Segmentation) requires recovering geometric and semantic structures from visual appearance alone. Second, *conditional synthesis* (Text+Depth→RGB, Text+Segmentation→RGB, and Text+Depth+Segmentation→RGB) assesses the generation of photorealistic UAV imagery consistent with both geometric constraints and linguistic intent. Third, *reconstruction* performs self-supervised reconstruction of RGB, depth, and segmentation, serving as an auxiliary signal to enforce consistency between high-level representations and pixel-level geometric and appearance details. Together, these families enable unified evaluation of semantic

Table 2: **Results on UAVReason for spatio-temporal VQA and scene captioning.** We evaluate 3 reasoning settings: single-frame VQA (VQA-1F), two-frame VQA (VQA-2F), and two-frame VQA with implicit heading cues (VQA-2F$_{\text{ihead}}$), alongside scene captioning. For VQA, we report Exact Match (EM), token-level F1, and an LLM-based judge score (LLM-J) to measure semantic correctness beyond lexical overlap. Methods are grouped into general-purpose VLMs, UAV/remote-sensing VLMs, unified generation models, and our unified framework. **UAVReason-Bagel** achieves strong overall performance while jointly supporting reasoning and generation. Compared with the *Reasoning only* ablation, joint training with generation improves semantic consistency on temporal reasoning under LLM-J.

Method	VQA-1F			VQA-2F			VQA-2F$_{\text{ihead}}$			Caption		
	EM ↑	F1 ↑	**LLM-J** ↑	EM ↑	F1 ↑	**LLM-J** ↑	EM ↑	F1 ↑	**LLM-J** ↑	B4 ↑	C ↑	**LLM-J** ↑
Group A: Off-the-shelf General VLMs / MLLMs												
Qwen2.5-VL-7B	0.041	0.217	0.355	0.000	0.138	0.454	0.000	0.088	0.000	0.015	0.321	7.02
Qwen2.5-VL-3B	0.052	0.163	0.325	0.000	0.236	0.323	0.000	0.086	0.050	0.019	0.319	6.33
Qwen3-VL-8B	0.031	0.113	0.325	0.000	0.207	0.475	0.000	0.088	0.080	0.435	0.041	7.01
Qwen3-VL-4B	0.031	0.132	0.340	0.000	0.279	0.414	0.000	0.041	0.070	0.037	0.513	6.99
LLaVA-OneVis-7B	0.031	0.109	0.247	0.000	0.099	0.126	0.000	0.030	0.060	0.007	0.221	5.89
InternVL3.5-1B	0.000	0.205	0.299	0.000	0.244	0.333	0.000	0.027	0.060	0.027	0.395	6.57
InternVL3.5-2B	0.062	0.163	0.294	0.000	0.140	0.414	0.140	0.865	0.180	0.024	0.416	7.03
InternVL3.5-4B	0.021	0.329	0.340	0.000	0.151	0.439	0.000	0.038	0.070	0.017	0.393	6.67
InternVL3.5-8B	0.010	0.215	0.314	0.000	0.069	0.409	0.000	0.042	0.130	0.022	0.384	6.82
Group B: UAV / Remote-Sensing VLMs												
GeoChat-7B	0.031	0.371	0.175	0.000	0.281	0.505	0.000	0.050	0.060	0.004	0.162	4.34
RS-LLaVA-1.5-7B	0.021	0.232	0.263	0.000	0.309	0.374	0.000	0.334	0.040	0.000	0.192	6.37
Group C: Unified Generation Models (Off-the-shelf)												
Janus-Pro-7B	0.021	0.141	0.294	0.000	0.423	0.253	0.000	0.215	0.090	0.021	0.327	6.77
Janus-1.3B	0.010	0.354	0.160	0.000	0.532	0.444	0.000	0.448	0.090	0.012	0.330	6.01
OmniGen2	0.000	0.208	0.299	0.000	0.425	0.298	0.000	0.080	0.050	0.014	0.294	6.01
Show-o2-7B	0.052	0.106	0.286	0.000	0.081	0.455	0.000	0.078	0.340	0.011	0.179	4.95
Group D: Our Unified Framework												
Bagel (Pretrain)	0.000	0.394	0.28	0.000	0.427	0.250	0.000	0.798	0.090	0.022	0.387	6.68
UAVReason-Bagel	**0.196**	0.711	0.397	0.186	0.822	**0.600**	0.657	0.973	0.657	0.216	1.530	**7.69**
Reasoning only	0.165	**0.734**	**0.454**	**0.234**	**0.855**	0.571	**0.770**	**0.982**	**0.770**	**0.242**	**1.741**	7.68
Generation only	0.052	0.422	0.330	0.000	0.401	0.240	0.000	0.801	0.145	0.017	0.371	6.63

reasoning and pixel-level generation under a consistent UAV data distribution and shared geometric constraints.

4 METHOD

In this work, rather than pursuing the strongest possible backbone, we focus on quantifying relative gains from jointly training complementary reasoning and generation tasks. We introduce **UAVReason-Bagel** (Figure 5), a specialized adaptation of the unified multimodal foundation model BAGEL (Deng et al., 2025). Building upon BAGEL's unified understanding-and-generation paradigm, UAVReason-Bagel employs a shared multimodal transformer backbone while decoupling task-specific computation through a Mixture-of-Transformer-Experts (MoT) design. In particular, a dedicated Generation Expert is optimized for velocity prediction over continuous visual latents within a Rectified Flow framework, while a dedicated Understanding Expert is optimized for discrete next-token prediction over language outputs. This architectural separation alleviates gradient conflicts arising from jointly optimizing discrete linguistic objectives and continuous visual generation objectives, thereby enabling a single backbone to effectively acquire both high-level spatio-temporal reasoning and pixel-level grounding in the UAV domain.

Table 3: **Dense perception and conditional cross-modal generation on UAVReason.** We report 18-class mIoU for semantic segmentation (Perc.). For conditional image generation, we evaluate three settings: depth+segmentation+text to RGB (dSeg+T→RGB), segmentation+text to RGB (Seg+T→RGB), and depth+text to RGB (d+T→RGB). Quality is measured by KID (↓), CLIP score (↑), and DINO similarity (↑). UAVReason-Bagel substantially outperforms off-the-shelf unified generators on dense perception and conditional generation. The *Generation only* ablation further indicates that reasoning supervision improves generation quality under sparse conditioning (d+T→RGB), yielding lower KID and higher DINO similarity.

Method	**Perc.**	**Conditional Generation Quality**								
	Seg	**dSeg+T→RGB**			**Seg+T→RGB**			**d+T→RGB**		
	mIoU↑	KID↓	CLIP↑	DINO↑	KID↓	CLIP↑	DINO↑	KID↓	CLIP↑	DINO↑
Off-the-shelf Unified Generators										
OmniGen2	0.037	0.094	25.46	0.317	0.093	24.29	0.507	0.094	24.40	0.378
Show-o2-7B	0.023	–	–	–	0.109	25.67	**0.537**	0.118	23.81	0.507
Our Unified Framework										
Bagel (Pretrain)	0.033	0.078	25.69	0.513	0.083	26.97	0.509	0.082	26.38	0.565
UAVReason-Bagel	**0.143**	**0.048**	**29.52**	0.648	**0.050**	**29.41**	0.524	**0.036**	**29.23**	**0.609**
Reasoning only	0.044	0.082	25.41	0.513	0.080	26.74	0.522	0.084	25.76	0.565
Generation only	0.143	0.051	28.59	**0.657**	0.050	29.01	0.536	0.052	28.86	0.594

Unified Architecture. We follow the MoT architecture and unify all UAVReason tasks under a single sequence-to-sequence formulation. For textual inputs (*e.g.*, questions, instructions, or prompts), the text encoder produces corresponding token sequences. For visual inputs (*e.g.*, RGB images, depth maps, segmentation masks, or temporal sequences), the visual encoder transforms them into visual representations, which are subsequently fused with text tokens. For reasoning-oriented tasks, textual outputs are generated autoregressively through next-token prediction. For generation tasks, the Generation Expert predicts continuous visual latents under the Rectified Flow formulation, which are subsequently decoded into images or dense geometric maps via the VAE decoder. This unified backbone enables language-level reasoning and pixel-level generation to share representations and supervisory signals within a single model, fostering mutual reinforcement across modalities and task families.

Multi-Task Learning Strategy. We train UAVReason-Bagel by jointly optimizing reasoning and generation objectives. The total training loss is formulated as:

$$\mathcal{L}_{\text{total}} = \lambda_{\text{CE}}\,\mathcal{L}_{\text{CE}} + \lambda_{\text{MSE}}\,\mathcal{L}_{\text{MSE}}, \tag{1}$$

where $\mathcal{L}_{\text{CE}}$ denotes the next-token cross-entropy loss for language outputs, and $\mathcal{L}_{\text{MSE}}$ denotes the mean-squared velocity-prediction loss used by the continuous visual generation branch under the Rectified Flow formulation. To counteract the *semantic floating* phenomenon prevalent in nadir-view UAV imagery, where objects often lack strong contextual cues due to limited viewpoint variation, we deliberately emphasize geometry-aware supervision. This emphasis is achieved by setting $\lambda_{\text{MSE}} : \lambda_{\text{CE}} = 2 : 1$ and adopting a 2:1 data sampling ratio favoring generation tasks over reasoning tasks. Intuitively, stronger pixel-level geometric supervision encourages the model to internalize robust structural priors (*e.g.*, depth discontinuities, relative scaling, and layout consistency) directly from dense ground-truth signals. These priors, in turn, anchor high-level semantic reasoning to physically plausible configurations, reducing hallucinations and promoting spatially grounded interpretations under the nadir viewpoint.

5 Experiment

5.1 Spatio-temporal VQA and Captioning

We compare UAVReason-Bagel against a broad spectrum of baselines, including (i) off-the-shelf general VLMs trained primarily on ground-centric data; (ii) specialized UAV/remote-sensing models, including GeoChat (Kuckreja et al., 2024) and RS-LLaVA (Bazi et al., 2024); and (iii) unified understanding-generation models, including Janus-Pro (Chen et al., 2025), Janus (Wu et al., 2025), OmniGen2 (Wu et al., 2026), and Show-o2 (Xie et al., 2025). As shown in Group A, most general

VLMs obtain near-zero EM across VQA-1F and VQA-2F, suggesting that producing *exact*, physically grounded answers remains difficult under tiny object scales and ambiguous orientations. While token-level F1 can be non-trivial for some models (*e.g.*, InternVL3.5-4B (Wang et al., 2025b) reaches 0.329 on VQA-1F), the corresponding LLM-J scores are typically modest (about 0.25–0.35), indicating that partial lexical overlap does not reliably translate to semantically correct spatial reasoning. This gap becomes more evident in the orientation-sensitive VQA-2F$_{\text{ihead}}$ setting, where LLM-J remains low for most general VLMs, reflecting limited ability to leverage heading cues for distinguishing relative spatial relations. In contrast, unified fine-tuning substantially improves spatio-temporal reasoning. Specifically, UAVReason-Bagel (Group D) achieves consistent gains over the pre-trained baseline. Fine-tuning improves VQA-1F from 0.394 to 0.711 in F1 (EM from 0.000 to 0.196), and yields stronger temporal reasoning on VQA-2F (0.822 F1 and 0.600 LLM-J versus 0.427 F1 and 0.250 LLM-J for Bagel (Pretrain)). Notably, with implicit heading cues (VQA-2F$_{\text{ihead}}$), UAVReason-Bagel reaches 0.973 F1 and 0.657 LLM-J, suggesting that the model can effectively leverage orientation signals to reason about relative spatial relations in aerial scenes.

For captioning, content relevance matters beyond n-gram fluency. UAVReason-Bagel achieves the highest LLM-J score of 7.69, while the *Reasoning only* variant obtains higher BLEU-4 and CIDEr scores. This difference further reflects the trade-off between task-specialized optimization and unified multi-task learning. CIDEr and LLM-J better reflect content relevance and semantic correctness. Beyond captioning quality, we further observe that generation objectives improve semantic robustness in temporal reasoning. The *reasoning only* ablation reveals a non-trivial interaction between understanding and generation supervision. Removing generation objectives slightly increases certain lexical-overlap metrics (*e.g.*, VQA-1F F1 0.734 vs. 0.711), but reduces semantic correctness on temporal reasoning under LLM-J (VQA-2F LLM-J drops from 0.600 to 0.571). These results suggest that dense generation losses encourage learning geometry-aware structural priors, which in turn improve semantically consistent reasoning in challenging spatio-temporal settings.

5.2 Dense Perception and Generation

As shown in Table 3, we summarize performance on dense perception (RGB→Seg) and three conditional generation settings. We observe that off-the-shelf unified generators struggle with aerial dense grounding. Models trained on broad internet distributions exhibit limited capability in UAV dense prediction. OmniGen2 (Wu et al., 2026) and Show-o2-7B (Xie et al., 2025) obtain very low segmentation performance (mIoU 0.037 and 0.023), and even the pretrained Bagel baseline remains low (0.033), indicating that nadir-view dense grounding requires domain-specific supervision. In contrast, UAVReason-Bagel improves dense perception and geometry-conditioned synthesis. With unified multi-task fine-tuning, UAVReason-Bagel achieves 0.143 mIoU on segmentation, substantially outperforming all baselines. This improved geometric grounding ability further translates to better conditional synthesis quality.

Please refer to the Appendix for more generalization results and ablation studies.

6 Conclusion

Standard VLMs struggle to generalize to the aerial domain, primarily due to *insufficient spatial grounding* under the ambiguous nadir viewpoint. To bridge this gap, we introduce **UAVReason**, a unified benchmark that jointly evaluates spatio-temporal reasoning and pixel-level generation in UAV scenes. Our baseline model, **UAVReason-Bagel**, shows that dense geometric supervision acts as a robust structural prior. Empirical results support our hypothesis that *"Synthesis promotes Analysis"*: generative objectives improve semantic robustness in temporal reasoning while maintaining geometry-consistent synthesis. We hope this work paves the way for future research on physically grounded embodied aerial agents.

7 Limitations

UAVReason is designed to study unified reasoning and generation in nadir-view UAV imagery, and the conclusions of this work should therefore be interpreted within this problem setting. UAVReason-Bagel further builds upon a pretrained unified multimodal foundation model and consequently inherits limitations of the underlying backbone, including its representation biases and generation errors.

AI USE STATEMENT

In this work, generative AI tools were used to assist with dataset construction, including hierarchical image enrichment, question generation, and answer synthesis and validation (Appendix E), and to provide feedback on research methodology and the interpretation of experimental results. Qwen3-VL-235B was used as an LLM-based semantic judge in the evaluation protocol (Appendix B). Generative AI tools were also used for manuscript editing. All AI-assisted work was reviewed by the authors against the underlying data, implementation, and experimental results. The authors take responsibility for the final content of this work.

ETHICS STATEMENT

UAVReason is constructed from the previously released UAVScenes data and is intended for research on aerial perception, reasoning, and generation. As aerial vision technologies may have privacy-sensitive or surveillance-related applications, users should follow the applicable data-use terms, privacy requirements, and relevant laws when deploying such systems.

REPRODUCIBILITY STATEMENT

UAVReason's dataset construction and processing pipeline are described in Section 3.1, with task definitions and data splits provided in Section 3.2. The unified model and training objectives are described in Section 4. Optimization and input-handling details are provided in Appendix A; the evaluation protocols and metrics are documented in Appendix B; ablation studies are reported in Appendix C; and the prompts used for data construction are included in Appendix E. We will release the data, code, and evaluation tools.

REFERENCES

Shuai Bai, Yuxuan Cai, Ruizhe Chen, Keqin Chen, Xionghui Chen, Zesen Cheng, Lianghao Deng, Wei Ding, Chang Gao, Chunjiang Ge, Wenbin Ge, Zhifang Guo, Qidong Huang, Jie Huang, Fei Huang, Binyuan Hui, Shutong Jiang, Zhaohai Li, Mingsheng Li, Mei Li, Kaixin Li, Zicheng Lin, Junyang Lin, Xuejing Liu, Jiawei Liu, Chenglong Liu, Yang Liu, Dayiheng Liu, Shixuan Liu, Dunjie Lu, Ruilin Luo, Chenxu Lv, Rui Men, Lingchen Meng, Xuancheng Ren, Xingzhang Ren, Sibo Song, Yuchong Sun, Jun Tang, Jianhong Tu, Jianqiang Wan, Peng Wang, Pengfei Wang, Qiuyue Wang, Yuxuan Wang, Tianbao Xie, Yiheng Xu, Haiyang Xu, Jin Xu, Zhibo Yang, Mingkun Yang, Jianxin Yang, An Yang, Bowen Yu, Fei Zhang, Hang Zhang, Xi Zhang, Bo Zheng, Humen Zhong, Jingren Zhou, Fan Zhou, Jing Zhou, Yuanzhi Zhu, and Ke Zhu. Qwen3-VL technical report. arXiv preprint arXiv:2511.21631, 2025a. URL https://arxiv.org/abs/2511.21631.

Shuai Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Sibo Song, Kai Dang, Peng Wang, Shijie Wang, Jun Tang, Humen Zhong, Yuanzhi Zhu, Mingkun Yang, Zhaohai Li, Jianqiang Wan, Pengfei Wang, Wei Ding, Zheren Fu, Yiheng Xu, Jiabo Ye, Xi Zhang, Tianbao Xie, Zesen Cheng, Hang Zhang, Zhibo Yang, Haiyang Xu, and Junyang Lin. Qwen2.5-VL technical report. arXiv preprint arXiv:2502.13923, 2025b. URL https://arxiv.org/abs/2502.13923.

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. doi: 10.3390/rs16091477.

Lin Chen, Jinsong Li, Xiaoyi Dong, Pan Zhang, Yuhang Zang, Zehui Chen, Haodong Duan, Jiaqi Wang, Yu Qiao, Dahua Lin, and Feng Zhao. Are we on the right way for evaluating large vision-language models? In *Advances in Neural Information Processing Systems*, volume 37, pp. 27056–27087. Curran Associates, Inc., 2024. doi: 10.52202/079017-0850.

Xiaokang Chen, Zhiyu Wu, Xingchao Liu, Zizheng Pan, Wen Liu, Zhenda Xie, Xingkai Yu, and Chong Ruan. Janus-Pro: Unified multimodal understanding and generation with data and model scaling. arXiv preprint arXiv:2501.17811, 2025. URL https://arxiv.org/abs/2501.17811.

An-Chieh Cheng, Hongxu Yin, Yang Fu, Qiushan Guo, Ruihan Yang, Jan Kautz, Xiaolong Wang, and Sifei Liu. SpatialRGPT: Grounded spatial reasoning in vision-language models. In *Advances in Neural Information Processing Systems*, volume 37, pp. 135062–135093. Curran Associates, Inc., 2024. doi: 10.52202/079017-4293.

Chaorui Deng, Deyao Zhu, Kunchang Li, Chenhui Gou, Feng Li, Zeyu Wang, Shu Zhong, Weihao Yu, Xiaonan Nie, Ziang Song, Guang Shi, and Haoqi Fan. Emerging properties in unified multimodal pretraining. arXiv preprint arXiv:2505.14683, 2025. URL https://arxiv.org/abs/2505.14683.

Dawei Du, Yuankai Qi, Hongyang Yu, Yifan Yang, Kaiwen Duan, Guorong Li, Weigang Zhang, Qingming Huang, and Qi Tian. The unmanned aerial vehicle benchmark: Object detection and tracking. In *Proceedings of the European Conference on Computer Vision*, pp. 370–386, 2018. doi: 10.1007/978-3-030-01249-6_23.

Horaţiu Florea, Vlad-Cristian Miclea, and Sergiu Nedevschi. WildUAV: Monocular UAV dataset for depth estimation tasks. In *2021 IEEE 17th International Conference on Intelligent Computer Communication and Processing*, pp. 291–298, 2021. doi: 10.1109/ICCP53602.2021.9733671.

Michaël Fonder and Marc Van Droogenbroeck. Mid-air: A multi-modal dataset for extremely low altitude drone flights. In *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pp. 553–562, 2019. doi: 10.1109/CVPRW.2019.00081.

Kartik Kuckreja, Muhammad Sohail Danish, Muzammal Naseer, Abhijit Das, Salman Khan, and Fahad Shahbaz Khan. GeoChat: Grounded large vision-language model for remote sensing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 27831–27840, 2024. doi: 10.1109/CVPR52733.2024.02629.

Baiqi Li, Zhiqiu Lin, Wenxuan Peng, Jean de Dieu Nyandwi, Daniel Jiang, Zixian Ma, Simran Khanuja, Ranjay Krishna, Graham Neubig, and Deva Ramanan. NaturalBench: Evaluating vision-language models on natural adversarial samples. In *Advances in Neural Information Processing Systems*, volume 37, pp. 17044–17068. Curran Associates, Inc., 2024a. doi: 10.52202/079017-0542.

Bo Li, Yuanhan Zhang, Dong Guo, Renrui Zhang, Feng Li, Hao Zhang, Kaichen Zhang, Peiyuan Zhang, Yanwei Li, Ziwei Liu, and Chunyuan Li. LLaVA-OneVision: Easy visual task transfer. *Transactions on Machine Learning Research*, 2025. URL https://openreview.net/forum?id=zKv8qULV6n.

Kun Li, George Vosselman, and Michael Ying Yang. HRVQA: A visual question answering benchmark for high-resolution aerial images. *ISPRS Journal of Photogrammetry and Remote Sensing*, 214:65–81, 2024b. doi: 10.1016/j.isprsjprs.2024.06.002.

Xiang Li, Jian Ding, and Mohamed Elhoseiny. VRSBench: A versatile vision-language benchmark dataset for remote sensing image understanding. In *Advances in Neural Information Processing Systems*, volume 37, pp. 3229–3242. Curran Associates, Inc., 2024c. doi: 10.52202/079017-0106.

Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C. Lawrence Zitnick. Microsoft COCO: Common objects in context. In *Computer Vision – ECCV 2014*, pp. 740–755, 2014. doi: 10.1007/978-3-319-10602-1_48.

Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. Visual instruction tuning. In *Advances in Neural Information Processing Systems*, volume 36, pp. 34892–34916. Curran Associates, Inc., 2023. doi: 10.52202/075280-1516.

Sylvain Lobry, Diego Marcos, Jesse Murray, and Devis Tuia. RSVQA: Visual question answering for remote sensing data. *IEEE Transactions on Geoscience and Remote Sensing*, 58(12):8555–8566, 2020. doi: 10.1109/TGRS.2020.2988782.

Xiaoqiang Lu, Binqiang Wang, Xiangtao Zheng, and Xuelong Li. Exploring models and data for remote sensing image caption generation. *IEEE Transactions on Geoscience and Remote Sensing*, 56(4):2183–2195, 2018. doi: 10.1109/TGRS.2017.2776321.

Ye Lyu, George Vosselman, Gui-Song Xia, Alper Yilmaz, and Michael Ying Yang. UAVid: A semantic segmentation dataset for UAV imagery. *ISPRS Journal of Photogrammetry and Remote Sensing*, 165:108–119, 2020. doi: 10.1016/j.isprsjprs.2020.05.009.

Hemal Naik, Junran Yang, Dipin Das, Margaret C. Crofoot, Akanksha Rathore, and Vivek Hari Sridhar. BuckTales: A multi-UAV dataset for multi-object tracking and re-identification of wild antelopes. In *Advances in Neural Information Processing Systems*, volume 37, pp. 81992–82009. Curran Associates, Inc., 2024. doi: 10.52202/079017-2606.

Ishan Nigam, Chen Huang, and Deva Ramanan. Ensemble knowledge transfer for semantic segmentation. In *2018 IEEE Winter Conference on Applications of Computer Vision*, pp. 1499–1508, 2018. doi: 10.1109/WACV.2018.00168.

OpenAI. GPT-4V(ision) system card. Technical report, 2023. URL https://cdn.openai.com/papers/GPTV_System_Card.pdf.

Yongzhou Pan, Binhong Liu, Zhen Liu, Hao Shen, Jianyu Xu, Wenxing Fu, and Tao Yang. MoNA Bench: A benchmark for monocular depth estimation in navigation of autonomous unmanned aircraft system. *Drones*, 8(2):66, 2024. doi: 10.3390/drones8020066.

Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. Learning transferable visual models from natural language supervision. In *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pp. 8748–8763. PMLR, 2021. URL https://proceedings.mlr.press/v139/radford21a.html.

Giulia Rizzoli, Francesco Barbato, Matteo Caligiuri, and Pietro Zanuttigh. SynDrone – multi-modal UAV dataset for urban scenarios. In *2023 IEEE/CVF International Conference on Computer Vision Workshops*, pp. 2202–2212, 2023. doi: 10.1109/ICCVW60793.2023.00235.

Junjue Wang, Zhuo Zheng, Zihang Chen, Ailong Ma, and Yanfei Zhong. EarthVQA: Towards queryable earth via relational reasoning-based remote sensing visual question answering. *Proceedings of the AAAI Conference on Artificial Intelligence*, 38(6):5481–5489, 2024. doi: 10.1609/aaai.v38i6.28357. URL https://ojs.aaai.org/index.php/AAAI/article/view/28357.

Sijie Wang, Siqi Li, Yawei Zhang, Shangshu Yu, Shenghai Yuan, Rui She, Quanjiang Guo, JinXuan Zheng, Ong Kang Howe, Leonrich Chandra, Shrivarshann Srijeyan, Aditya Sivadas, Toshan Aggarwal, Heyuan Liu, Hongming Zhang, Chujie Chen, Junyu Jiang, Lihua Xie, and Wee Peng Tay. UAVScenes: A multi-modal dataset for UAVs. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 28946–28958, 2025a. doi: 10.1109/ICCV51701.2025.02688.

Weiyun Wang, Zhangwei Gao, Lixin Gu, Hengjun Pu, Long Cui, Xingguang Wei, Zhaoyang Liu, Linglin Jing, Shenglong Ye, Jie Shao, Zhaokai Wang, Zhe Chen, Hongjie Zhang, Ganlin Yang, Haomin Wang, Qi Wei, Jinhui Yin, Wenhao Li, Erfei Cui, Guanzhou Chen, Zichen Ding, Changyao Tian, Zhenyu Wu, Jingjing Xie, Zehao Li, Bowen Yang, Yuchen Duan, Xuehui Wang, Songze Li, Xiangyu Zhao, Haodong Duan, Nianchen Deng, Bin Fu, Yinan He, Yi Wang, Conghui He, Botian Shi, Junjun He, Yingtong Xiong, Han Lv, Lijun Wu, Wenqi Shao, Kaipeng Zhang, Huipeng Deng, Biqing Qi, Jiaye Ge, Qipeng Guo, Wenwei Zhang, Wanli Ouyang, Limin Wang, Min Dou, Xizhou Zhu, Tong Lu, Dahua Lin, Jifeng Dai, Bowen Zhou, Weijie Su, Kai Chen, Yu Qiao, Wenhai Wang, and Gen Luo. InternVL3.5: Advancing open-source multimodal models in versatility, reasoning, and efficiency. arXiv preprint arXiv:2508.18265, 2025b. URL https://arxiv.org/abs/2508.18265.

Wenshan Wang, Delong Zhu, Xiangwei Wang, Yaoyu Hu, Yuheng Qiu, Chen Wang, Yafei Hu, Ashish Kapoor, and Sebastian Scherer. TartanAir: A dataset to push the limits of visual SLAM. In *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 4909–4916, 2020. doi: 10.1109/IROS45743.2020.9341801.

Syed Waqas Zamir, Aditya Arora, Akshita Gupta, Salman Khan, Guolei Sun, Fahad Shahbaz Khan, Fan Zhu, Ling Shao, Gui-Song Xia, and Xiang Bai. iSAID: A large-scale dataset for instance segmentation in aerial images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pp. 28–37, 2019.

Chengyue Wu, Xiaokang Chen, Zhiyu Wu, Yiyang Ma, Xingchao Liu, Zizheng Pan, Wen Liu, Zhenda Xie, Xingkai Yu, Chong Ruan, and Ping Luo. Janus: Decoupling visual encoding for unified multimodal understanding and generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 12966–12977, 2025. doi: 10.1109/CVPR52734.2025.01210.

Chenyuan Wu, Jiahao Wang, Pengfei Zheng, Ruiran Yan, Shitao Xiao, Xin Luo, Yueze Wang, Wanli Li, Xiyan Jiang, Yexin Liu, Junjie Zhou, Ziyi Xia, Ze Liu, Chaofan Li, Haoge Deng, Kun Luo, Bo Zhang, Jiajun Zhang, Dong Liu, Defu Lian, Xinlong Wang, Zhongyuan Wang, Tiejun Huang, and Zheng Liu. OmniGen2: Towards instruction-aligned multimodal generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 21964–21975, 2026.

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*, pp. 3974–3983, 2018. doi: 10.1109/CVPR.2018.00418.

Shitao Xiao, Yueze Wang, Junjie Zhou, Huaying Yuan, Xingrun Xing, Ruiran Yan, Chaofan Li, Shuting Wang, Tiejun Huang, and Zheng Liu. OmniGen: Unified image generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 13294–13304, 2025. doi: 10.1109/CVPR52734.2025.01241.

Jinheng Xie, Zhenheng Yang, and Mike Zheng Shou. Show-o2: Improved native unified multimodal models. In *Advances in Neural Information Processing Systems*, volume 38, pp. 53002–53030, 2025. doi: 10.52202/085713-1584.

Hui Ye, Rajshekhar Sunderraman, and Shihao Ji. UAV3D: A large-scale 3d perception benchmark for unmanned aerial vehicles. In *Advances in Neural Information Processing Systems*, volume 37, pp. 55425–55442. Curran Associates, Inc., 2024. doi: 10.52202/079017-1760.

Peter Young, Alice Lai, Micah Hodosh, and Julia Hockenmaier. From image descriptions to visual denotations: New similarity metrics for semantic inference over event descriptions. *Transactions of the Association for Computational Linguistics*, 2:67–78, 2014. doi: 10.1162/tacl_a_00166.

Wei Zhang, Miaoxin Cai, Tong Zhang, Yin Zhuang, and Xuerui Mao. EarthGPT: A universal multimodal large language model for multisensor image comprehension in remote sensing domain. *IEEE Transactions on Geoscience and Remote Sensing*, 62:1–20, 2024. doi: 10.1109/TGRS.2024.3409624.

Baining Zhao, Jianjie Fang, Zichao Dai, Ziyou Wang, Jirong Zha, Weichen Zhang, Chen Gao, Yue Wang, Jinqiang Cui, Xinlei Chen, and Yong Li. UrbanVideo-Bench: Benchmarking vision-language models on embodied intelligence with video data in urban spaces. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 32400–32423. Association for Computational Linguistics, 2025. doi: 10.18653/v1/2025.acl-long.1558. URL https://aclanthology.org/2025.acl-long.1558/.

Xiangtao Zheng, Binqiang Wang, Xingqian Du, and Xiaoqiang Lu. Mutual attention inception network for remote sensing visual question answering. *IEEE Transactions on Geoscience and Remote Sensing*, 60:1–14, 2022. doi: 10.1109/TGRS.2021.3079918.

Zhedong Zheng, Yunchao Wei, and Yi Yang. University-1652: A multi-view multi-source benchmark for drone-based geo-localization. In *Proceedings of the 28th ACM International Conference on Multimedia*, pp. 1395–1403, 2020. doi: 10.1145/3394171.3413896.

Pengfei Zhu, Longyin Wen, Xiao Bian, Haibin Ling, and Qinghua Hu. Vision meets drones: A challenge. arXiv preprint arXiv:1804.07437, 2018. URL https://arxiv.org/abs/1804.07437.

APPENDIX

Contents

A Implementation Details **14**

B Evaluation Protocol **14**

C Ablation Studies **15**

D More Qualitative Analysis **16**

D.1 Spatio-temporal VQA and Captioning Qualitative Analysis 16

D.2 Dense Perception and Generation Qualitative Analysis 19

E Prompt Engineering Details **19**

A IMPLEMENTATION DETAILS

Optimization Setup. We perform full-parameter fine-tuning of the Bagel-7B-MoT backbone on 8 GPUs for 30,000 steps. We use AdamW with a peak learning rate of 8×10^{-6} and a 1,000-step linear warmup, followed by a cosine decay schedule. To test our hypothesis that geometry supervision provides a structural constraint for aerial reasoning, we enforce the loss weighting in Eq. 1 as $\lambda_{\text{MSE}} : \lambda_{\text{CE}} = \mathbf{2} : \mathbf{1}$. This design targets the *semantic floating* issue by prioritizing pixel-level fidelity over language-only objectives, which helps reduce the tendency to overfit to dataset-specific language priors and encourages grounded adaptation to the nadir UAV domain.

Efficient Input Handling. UAVReason involves heterogeneous input patterns, ranging from single-frame spatial queries to two-frame temporal sequences. To process these mixed-length multimodal samples efficiently, we adopt dynamic sequence packing with a maximum budget of 11,520 tokens per batch. In particular, since spatio-temporal VQA consumes a pair of frames (X_{v1}, X_{v2}) together with question tokens, this strategy ensures high training throughput without truncating the long-context inputs required for temporal reasoning.

B EVALUATION PROTOCOL

We adopt a strict, task-aligned evaluation protocol to handle the open-ended nature of aerial reasoning and to quantify generative fidelity. Beyond naive string matching, we report conventional automatic metrics together with Large Language Model (LLM) based semantic evaluation. For LLM based evaluation, we use **Qwen3-VL-235B** (Bai et al., 2025a) as the judge model and apply fixed prompt templates and rubrics to ensure consistent and reproducible scoring. We denote these semantic scores as **LLM-Judge** (LLM-J), with task-specific scales used across VQA and captioning.

VQA (Single-frame and Two-frame). We report **Exact Match (EM)** and **Token-level F1** after standard answer normalization, including lowercasing, punctuation and article removal, and whitespace canonicalization. For numerically grounded questions (*e.g.*, distance, area, speed), we additionally report a tolerance-aware correctness criterion, where a prediction is considered correct if it falls within a relative tolerance $\tau = 5\%$ or an absolute tolerance ϵ of the ground truth, with ϵ depending on the physical unit. To assess semantic correctness beyond lexical overlap, we further use an LLM judge, *i.e.*, Qwen3-VL-235B (Bai et al., 2025a), which compares the model prediction with the ground-truth answer under two complementary notions of correctness. The judge outputs two boolean fields in JSON format: `semantic_correct` and `strict_correct`. `semantic_correct` allows paraphrases and minor numerical deviations as long as key facts and relations are preserved, while `strict_correct` requires exact matches for categorical and counting answers. Based

Table 4: **Effect of input image resolution on UAVReason understanding tasks.** Resolution 384 achieves the strongest captioning scores and competitive VQA performance. Higher resolution (512) specifically benefits heading-aware VQA, likely due to better visibility of subtle orientation cues.

Resolution	VQA-1F		VQA-2F		VQA-2F$_{\text{ihead}}$		Caption	
	EM ↑	F1 ↑	EM ↑	F1 ↑	EM ↑	F1 ↑	B4 ↑	CIDEr ↑
256	**0.196**	0.711	0.186	0.822	0.657	0.973	0.216	1.530
384	0.186	**0.727**	0.197	**0.850**	0.793	0.980	**0.254**	**1.800**
512	0.189	0.725	**0.200**	0.847	**0.868**	**0.990**	0.236	1.663

Table 5: **Effect of input resolution on dense perception and conditional generation.** Higher resolutions generally improve generation fidelity, with 512 achieving the lowest KID and highest DINO across the three synthesis settings. However, segmentation performance peaks at 384.

Resolution	Perc.	Conditional Generation Quality								
	Seg	dSeg+T→RGB			Seg+T→RGB			d+T→RGB		
	mIoU↑	KID↓	CLIP↑	DINO↑	KID↓	CLIP↑	DINO↑	KID↓	CLIP↑	DINO↑
256	0.143	0.048	**29.52**	0.648	0.050	29.41	0.524	0.036	29.23	0.609
384	**0.194**	0.041	29.14	0.554	0.032	**30.23**	0.696	0.028	29.20	0.582
512	0.144	**0.024**	28.02	**0.697**	**0.024**	29.52	**0.701**	**0.018**	**29.38**	**0.724**

on these two fields, we compute the VQA LLM-J score by mapping {(`semantic_correct`, `strict_correct`)} to a three-level score: 1.0 if both are true, 0.5 if only `semantic_correct` is true, and 0.0 otherwise, then averaging over the test set.

Scene Captioning. We evaluate caption quality using **BLEU-4** for n-gram fluency and **CIDEr** for content relevance. CIDEr is particularly important for UAV imagery because it upweights informative geometry-related terms and relations (*e.g.*, *paved road*, *roof*, *intersection*) over generic descriptions. Similarly, we report an LLM-based semantic score. The judge is provided with the candidate caption, a set of reference captions, and the image, and assigns an integer score on a 1–10 scale. The rubric evaluates semantic similarity to the references, coverage of salient objects and relations, and penalizes hallucinations of objects not supported by the image or references. The final caption LLM-J score is the average judge score across the test set.

Geometry Perception (RGB→Segmentation). For semantic segmentation, we report **mean Intersection over Union (mIoU)** across 18 classes. Predictions are converted from the segmentation color palette to discrete class indices via nearest-neighbor matching prior to evaluation, providing a robust measure of boundary adherence and region accuracy against ground-truth masks.

Conditional Generation. Since pixel-wise metrics (*e.g.*, PSNR) often penalize plausible multimodal outputs, we prioritize feature-level and distribution-level metrics. We report **KID** (↓) to measure distribution distance between generated and real UAV images, **CLIP Score** (↑) to assess text-image semantic alignment, and **DINOv2 Similarity** (↑) to quantify structural fidelity to the paired ground truth. Together, these metrics capture complementary aspects of generation quality, including realism, semantic controllability, and geometric consistency.

C Ablation Studies

Detailed conditional-generation analysis. For dSeg+T→RGB, UAVReason-Bagel reduces KID from 0.078 to 0.048 while increasing CLIP score from 25.69 to 29.52 and DINO score from 0.513 to 0.648, indicating better realism, text alignment, and structural fidelity with the corresponding ground

Table 6: **Impact of lr on understanding tasks.** A lower lr (5×10^{-6}) performs best on most multi-frame VQA and captioning metrics, whereas our selected lr (8×10^{-6}) achieves the highest exact-match score on single-frame VQA. Together with its stronger overall dense perception and generation performance in Table 7, we use 8×10^{-6} as the default setting.

lr	VQA-1F		VQA-2F		VQA-2F$_{\text{ihead}}$		Caption	
	EM ↑	F1 ↑	EM ↑	F1 ↑	EM ↑	F1 ↑	B4 ↑	CIDEr ↑
5×10^{-6}	0.189	0.725	**0.200**	**0.847**	**0.868**	**0.990**	**0.236**	**1.663**
$\mathbf{8 \times 10^{-6}}$	**0.196**	0.711	0.186	0.822	0.657	0.973	0.216	1.530
1×10^{-5}	0.186	**0.726**	0.159	0.839	0.725	0.725	0.225	1.540

truth. Similar improvements are observed for Seg+T→RGB and d+T→RGB, especially in CLIP scores, suggesting stronger controllability under UAV prompts.

Understanding supervision as a regularizer for generation. We further observe that supervision from understanding tasks provides an implicit regularization effect on generation. To examine this effect, we perform a *generation-only* ablation by removing all language-based understanding objectives (*i.e.*, VQA and captioning) during training, while keeping the generation losses unchanged. Although segmentation performance remains stable (mIoU = 0.143), generation quality degrades under depth-only conditioning (d+T→RGB): KID increases from 0.036 to 0.052, and DINO similarity drops from 0.609 to 0.594. These results indicate that language-level understanding objectives regularize the shared representation space, enabling the model to infer correct semantics and realistic textures even when visual inputs are sparse or ambiguous.

Ablation on Image Resolution. Resolution trades off efficiency, dense grounding, and heading-aware reasoning. Tables 4 and 5 summarize the effect of input resolution. For understanding, 384 improves captioning (CIDEr 1.800) while remaining competitive on VQA (Table 4). Meanwhile, 512 yields the strongest heading-aware setting (VQA-2F$_{\text{ihead}}$ with EM 0.868 and F1 0.990), suggesting that higher resolution benefits fine-grained orientation cues. For generation, higher resolution generally improves distribution and structural fidelity (Table 5), with 512 achieving the lowest KID and highest DINO across conditional synthesis settings. These results show that increasing the input resolution can provide additional gains on selected tasks, while resolution 256 is used as the common setting for the main experiments.

Ablation on Learning Rate. We investigate the impact of lr to identify an appropriate optimization setting for jointly learning multimodal understanding and pixel-level generation. As shown in Table 6 and Table 7, the choice of lr reveals a trade-off between the two capabilities. A lower learning rate (lr=5×10^{-6}) tends to favor text-heavy tasks such as captioning and multi-frame VQA, likely because it better preserves the pretrained language representations. However, for dense perception and generation tasks, our default setting of lr=8×10^{-6} provides stronger overall performance. Specifically, lr=8×10^{-6} yields the best segmentation performance (0.143 mIoU), the lowest KID for dSeg+T→RGB (0.048) and Seg+T→RGB (0.050), and the highest DINO scores across all three generation conditions. Although lr=1×10^{-5} achieves the lowest KID for d+T→RGB (0.028), its segmentation performance and DINO similarity are lower than those obtained with lr=8×10^{-6}. Moreover, lr=8×10^{-6} maintains the highest exact-match score on single-frame VQA (0.196). We therefore adopt lr=8×10^{-6} as the default learning rate, as it provides the most balanced performance across understanding and generation tasks.

D More Qualitative Analysis

D.1 Spatio-temporal VQA and Captioning Qualitative Analysis

We further provide concrete examples in Figure 6 and Figure 7 to visualize the limitations of general-domain VLMs in aerial scenarios. Despite their strong general capabilities, models like Qwen2.5-VL (Bai et al., 2025b) and InternVL3.5 (Wang et al., 2025b) exhibit unreliable grounding

Table 7: **Effect of lr on dense perception and generation.** Our selected lr (8×10^{-6}) achieves the highest segmentation mIoU, the lowest KID for dSeg+T→RGB and Seg+T→RGB, and the highest DINO similarity across all three generation settings. Although 1×10^{-5} obtains the lowest KID for d+T→RGB, 8×10^{-6} provides the strongest overall performance across dense perception and conditional generation.

lr	Perc.	Conditional Generation Quality								
	Seg	dSeg+T→RGB			Seg+T→RGB			d+T→RGB		
	mIoU ↑	KID ↓	CLIP ↑	DINO ↑	KID ↓	CLIP ↑	DINO ↑	KID ↓	CLIP ↑	DINO ↑
5×10^{-6}	0.135	0.067	28.85	0.582	0.071	**29.49**	0.514	0.062	28.84	0.507
8×10^{-6}	**0.143**	**0.048**	29.52	**0.648**	**0.050**	29.41	**0.524**	0.036	**29.23**	**0.609**
1×10^{-5}	0.093	0.070	**29.73**	0.563	0.073	28.91	0.443	**0.028**	28.23	0.486

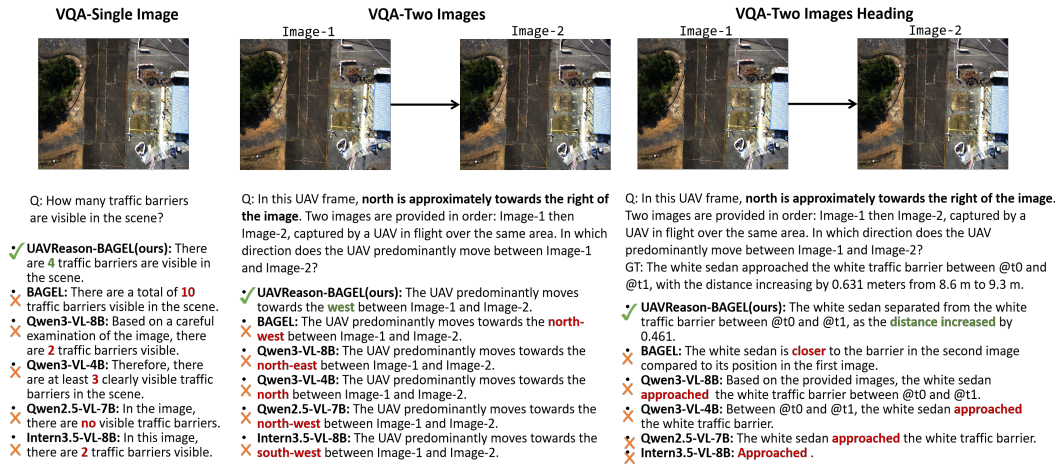

Figure 6: **Qualitative comparison on fine-grained spatio-temporal reasoning.** We evaluate **UAVReason-Bagel** (Ours) against leading general-domain VLMs (*e.g.*, Qwen2.5-VL (Bai et al., 2025b), InternVL3.5 (Wang et al., 2025b)) across three challenging aerial tasks. **Left (Counting):** Our model accurately counts tiny, densely packed instances (traffic barriers), whereas baselines suffer from severe miss-detection or count hallucinations. **Middle (Navigation):** Guided by explicit orientation cues (*e.g.*, "North is right"), our model correctly aligns the heading to infer the "West" trajectory, while general VLMs fail to adapt compass directions to the nadir view. **Right (Temporal Dynamics):** Our model successfully detects fine-grained status changes (separation) and quantifies distance shifts, whereas baselines often hallucinate the opposite motion trend (*e.g.*, "approaching").

when facing the unique challenges of UAV imagery. First, in VQA tasks, they struggle with *scale and density*, frequently hallucinating counts for tiny objects like traffic barriers (Figure 6, Left). Second, they lack *viewpoint adaptability*, failing to map explicit orientation prompts (*e.g.*, rotated compasses) to the nadir coordinate system, leading to incorrect navigation inferences (Figure 6, Middle). Third, regarding *temporal dynamics*, baselines often misinterpret fine-grained geometric changes, such as hallucinating an "approaching" motion when objects are clearly separating (Figure 6, Right). Finally, this grounding gap extends to *scene description* (Figure 7). While baselines often resort to coarse abstractions (*e.g.*, labeling specific sedans merely as "vehicles") or suffer from semantic drift (*e.g.*, Qwen2.5-VL misidentifying an airstrip as a "construction site"), UAVReason-Bagel preserves *instance-level granularity*. It correctly resolves specific categories like "sedans" and captures the precise land use, confirming that our unified training strategy effectively mitigates domain-specific hallucinations across both reasoning and generation tasks.

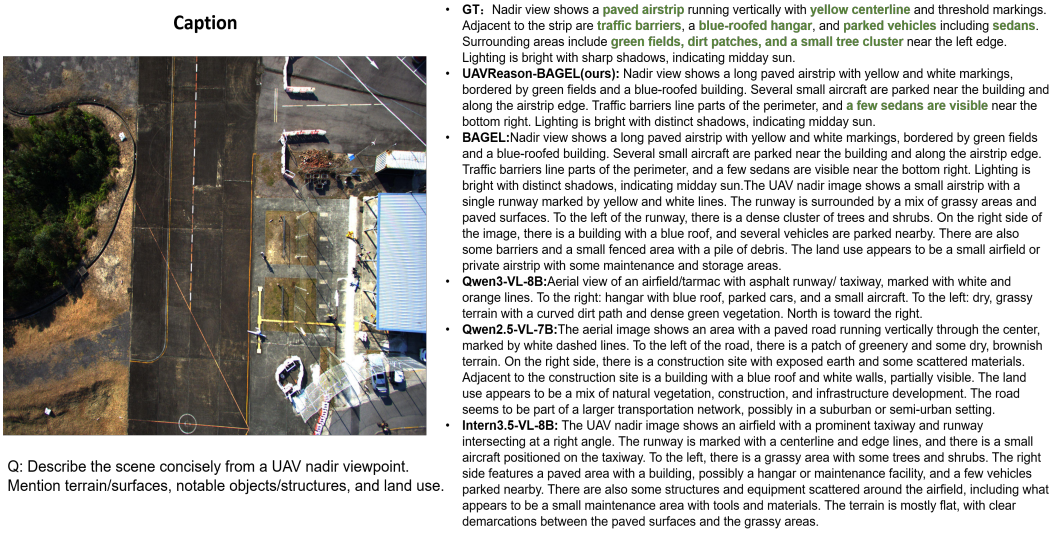

Figure 7: **Qualitative comparison on fine-grained scene description.** On a nadir-view airstrip, **UAVReason-Bagel (Ours)** achieves superior *semantic granularity*, correctly identifying the context ("airstrip") and resolving specific instances (*e.g.*, "sedans"). In contrast, general-domain VLMs exhibit severe semantic drift, often hallucinating non-existent scenes (*e.g.*, "construction site") or resorting to coarse abstractions (generic "vehicles").

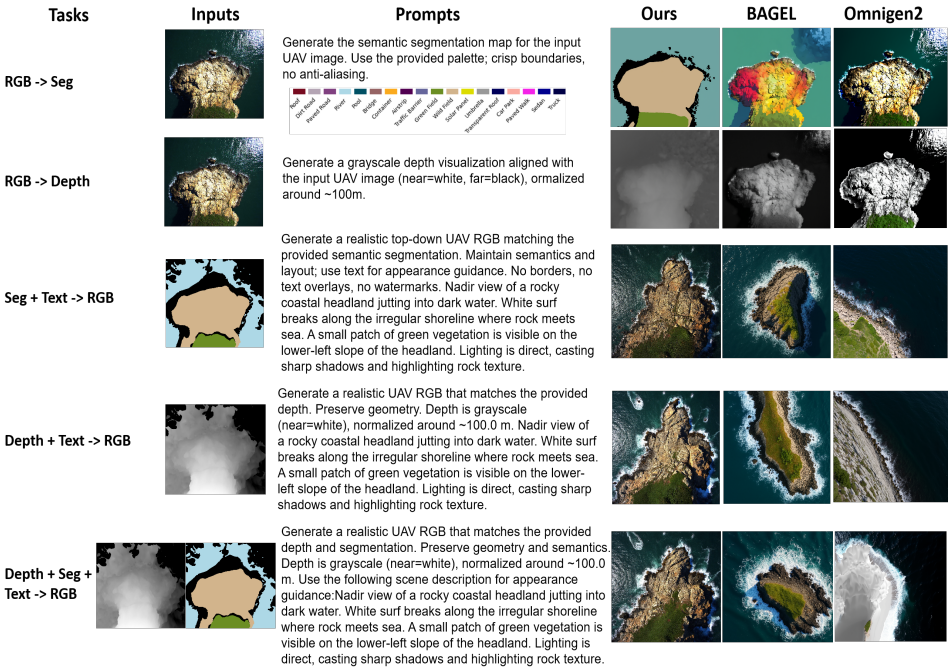

Figure 8: **Qualitative comparison of dense perception and conditional synthesis.** We visualize predictions from **UAVReason-Bagel** versus general-domain baselines (BAGEL, OmniGen2). In **perception tasks** (Top rows), our model recovers fine-grained object boundaries and smooth depth gradients, whereas baselines exhibit significant noise and semantic errors. In **conditional synthesis** (Bottom rows), our model generates photorealistic textures that more closely adhere to the input geometric layouts (depth/segmentation). In contrast, baselines suffer from severe domain shift, resulting in structural hallucinations and failure to follow the nadir-view control signals.

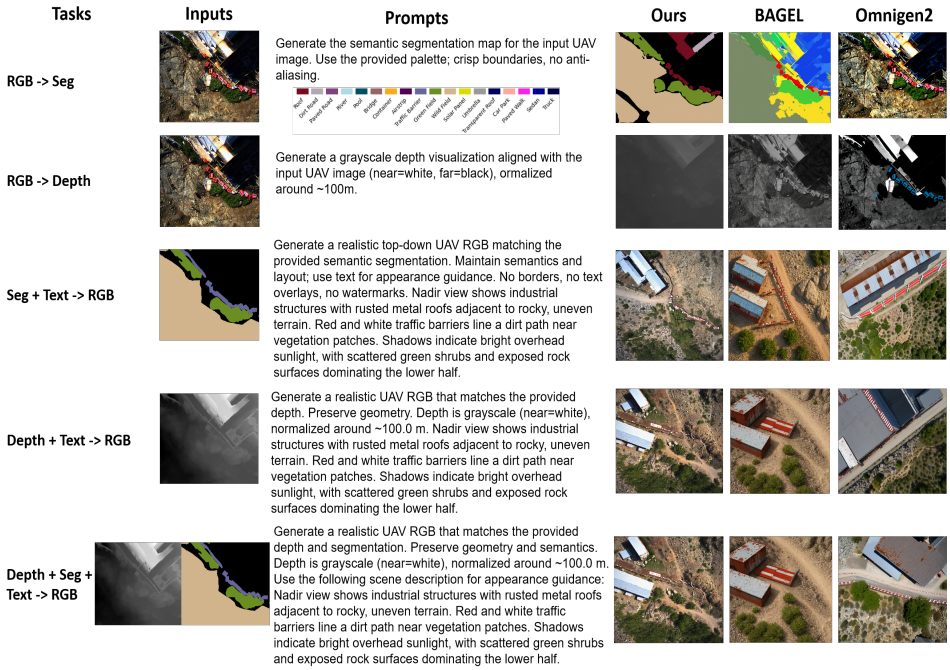

Figure 9: **Qualitative comparison of dense perception and conditional synthesis.** We visualize predictions from **UAVReason-Bagel** versus general-domain baselines (BAGEL, OmniGen2). In **perception tasks** (Top rows), our model recovers fine-grained object boundaries and smooth depth gradients, whereas baselines exhibit significant noise and semantic errors. In **conditional synthesis** (Bottom rows), our model generates photorealistic textures that strictly adhere to the input geometric layouts (depth/segmentation). In contrast, baselines suffer from severe domain shift, resulting in structural hallucinations and failure to follow the nadir-view control signals.

D.2 Dense Perception and Generation Qualitative Analysis

We provide a visual comparison in Figure 8 and Figure 9 to intuitively demonstrate the superiority of our unified approach. **(1) Dense Perception:** As shown in the top rows, general-domain models (OmniGen2, BAGEL) struggle with the nadir viewpoint, often producing fragmented segmentation masks and noisy depth estimations that fail to capture the flat, layout-centric nature of UAV scenes. In contrast, UAVReason-Bagel makes clear, semantically coherent predictions with high structural fidelity. **(2) Conditional Synthesis:** The bottom rows highlight a critical advantage in controllability. When generating RGB images from geometric conditions (Depth/Seg), baselines frequently hallucinate non-existent objects or distort the landscape, leading to weak alignment between the control signal and the generated output. Our model, however, maintains strict *geometric adherence*, synthesizing realistic textures that perfectly align with the provided layout. This qualitative evidence validates our hypothesis that unified training on UAVReason effectively bridges the domain gap, mitigating the "semantic floating" issue prevalent in existing VLMs.

E Prompt Engineering Details

To ensure that the generated Question-Answer (QA) pairs remain grounded in the observed UAV imagery and associated geometric information, we employ a coarse-to-fine prompting pipeline. For answer synthesis, we further adopt a program-aided strategy in which numerical and geometric facts are computed deterministically rather than inferred directly by the language model. The pipeline consists of three stages: hierarchical image enrichment, constraint-based QA generation, and program-aided answer synthesis.

E.1 Hierarchical Image Enrichment

We first transform the raw visual data into a structured textual representation through global scene parsing and fine-grained instance reasoning.

PROMPT: STEP-1 GLOBAL SCENE PARSING

```
INSTRUCTIONS (STEP-1 ONLY)
You are a professional UAV nadir/near-nadir image-to-text
annotator.
 • Based on the aerial image and the CANDIDATE instances list,
   output STRICT JSON with ONLY: "version", "image_size",
   "scene".
 • DO NOT output "instances" or "relations" in this step.
HARD RULES
 • Image coords: origin top-left; x → right, y → down.
 • Ontology: Do NOT invent new labels; ontology is fixed.
 • Language: Concise, factual, pixel-grounded (nadir cues).
   Avoid hedging.
 • Directions: Only left/right/top/bottom, near/far; never
   N/E/S/W.
OUTPUT JSON SCHEMA
{
   "version": "<model_version>",
   "image_size": { "width": [W], "height": [H] },
   "scene": {
     "scene_type": "residential | industrial | ...",
     "global_caption": "3-4 sentences with nadir cues...",
     "main_surfaces": ["paved_motor_road", ...],
     "time_of_day": "day | dusk | night | unknown"
   }
}
```

PROMPT: STEP-2 INSTANCE ATTRIBUTE & RELATION REASONING

```
INSTRUCTIONS (STEP-2 ONLY)
You are a professional UAV nadir/near-nadir image-to-text
annotator.
 • Using ONLY the provided STEP-1 scene JSON and
   ACCEPTED_PROPOSALS, return STRICT JSON with: "by_obj" and
   "relations".
 • DO NOT return bbox/label/src_id in "by_obj".
by_obj VALUE SCHEMA
{
   "attributes": {
     "color": "...", "condition": "intact | damaged",
     "state": "stationary | moving", "occluded": true | false
   },
   "attribute_sentence": "1-2 factual sentences...",
   "spatial_context_sentence": "1-2 sentences using relative
directions...",
   "evidence": "≤ 10 words"
}
HARD RULES
 • Produce an entry in "by_obj" for EVERY obj_id given.
 • Use only relative directions (left/right/top/bottom).
 • Relations must reference existing obj_ids and be
   pixel-grounded.
```

E.2 CONSTRAINT-BASED QA GENERATION

We generate verifiable single-frame and temporal questions from the structured scene representation while enforcing explicit object grounding and task-specific constraints.

PROMPT: SINGLE-FRAME VQA GENERATION

ROLE: You are a professional UAV nadir/near-nadir VQA author.
GOAL
- Design **EXACTLY 9 verifiable questions** using ONLY the provided scene JSON.
- Output JSON ONLY with: {"generated_questions":[...]}.

CATEGORY & COUNT (STRICT)
- Common Scenes: "Existence", "Count", "Proportion", "Attribute"
- Spatial Reasoning: "Adjacency", "Distance", "Direction", "Occlusion"
- CoT Multi-Hop: "Composition", "Constraint", "Filter-Then-Measure"

STRONG CONSTRAINTS
- Each question MUST reference an existing obj_id or label.
- Keep each question $\leq$ 22 words; concise and unambiguous.
- Favor **pixel-space geometry**. Do not use meters unless explicitly provided.

PROMPT: TEMPORAL (MULTI-FRAME) VQA GENERATION

ROLE: You are a professional UAV nadir/near-nadir VQA author.
GOAL
- Design **EXACTLY 3 verifiable questions** using ONLY the provided TWO-FRAME JSON.
- Categories: "Identity", "DistanceChange", "Speed", "RelationChange", "Approach/Separate".

TEMPORAL RULES
- EVERY question MUST cite **BOTH frames** (@t0 and @t1).
- Prefer using "tracks" (trk#k) to refer to the same instance across frames.
- Temporal reasoning requires strict object grounding across timestamps.

E.3 PROGRAM-AIDED ANSWER SYNTHESIS

To reduce numerical and geometric hallucinations, we use a *Plan–Execute–Verify* pipeline. The Planner converts each question into a structured operation, a deterministic Python executor computes the required quantities from ground-truth data, and the Writer generates the final response, which is subsequently checked by the Validator.

PROMPT: ANSWER PLANNER

INSTRUCTIONS
You are a planner that converts a UAV VQA question into a computable plan.
- Use ONLY obj_ids and labels that exist in the provided scene JSON.
- Output a SINGLE JSON object. Keep rationale short and PRIVATE.
- Prefer 'world_distance' if 3D centers exist; otherwise 'pixel_distance'.

ALLOWED OPERATIONS
["existence", "count", "proportion", "left_right_top_bottom", "nearest", "abs_distance", "ordering", "size_compare", "adjacency_touching", "overlaps", "on_inside", "occlusion_flag", "bottom"]
JSON SCHEMA

```
{
   "op": "<operation>",
   "targets": {
     "subject_ids": [...], "object_ids": [...],
     "subject_filter": { "label": "...", "attributes": {...}
},
     "object_filter": {...}
   },
   "metric": "world_distance | pixel_distance | auto",
   "constraints": { "k": 1, "relation": "..." }
}
```

PROMPT: ANSWER WRITER (FACT-CONDITIONED)

RULES

- Use **ONLY** the provided computed facts; **NEVER** change numbers.
- Write 1–3 fluent sentences. No bullets, no JSON.
- **Units**: Always express distances in meters (m) if available.
- If metrics.proportion exists, report numerator/denominator/ratio.
- Keep reasoning PRIVATE; do not reveal step-by-step thoughts.

INPUT DATA

Question: [Input Question]
Facts: [Computed JSON from Python Executor, e.g., distance_m, counts, booleans]

PROMPT: ANSWER VALIDATOR

TASK

- Compare the DRAFT ANSWER against the FACTS.
- **Do NOT invent numbers**. Only accept numbers that appear in explicit metric fields.
- If consistent: reply { "verdict": "ok" }.
- If correction needed: reply { "verdict": "fix", "rewrite": "..." }.

REWRITE RULES

- Never add new numbers not present in FACTS.
- If distances are pixel-only, do NOT output numeric distance.
- Keep it short ($\leq$ 2 sentences).