BraintankAI

DEEP READ · DRIVING VLA MODELS · 8 MIN

Qwen-Drive adds 3D perception and a planner to a 4B VLM, keeping its architecture

Qwen researchers build on a 4-billion-parameter vision-language model to map scenes in 3D, answer driving questions and plan paths, with mixed results across benchmarks.

17 November 2026 · arXiv preprint, not yet peer reviewed



Photo: Mika Baumeister

A vision-language model, or VLM, is an AI system that takes in images and text and replies in text. Over the past few years, researchers have tried to turn such models into self-driving software. The usual goal is a vision-language-action model (VLA): the same system, extended so that it also outputs an action, such as the path a car should follow. In a 2026 preprint, which has not yet been peer reviewed, a team from Qwen and Huazhong University of Science and Technology describes Qwen-Drive-1.0. It starts from Qwen3.5-4B, an existing 4-billion-parameter VLM, and leaves its internal architecture untouched. The team adds two external modules: one that builds a 3D map of the car's surroundings from camera images, and one that draws the future path. The result is a single model that handles 3D perception, answers questions about driving scenes and plans motion.

The authors identify two weaknesses in the common way of building driving VLAs, which adapts a general model mainly by training it on driving questions and text answers. First, text answers say little about precise depth, distances or which parts of space are occupied. A model can therefore describe a scene fluently while placing objects inaccurately, and nobody can measure that error directly. Second, heavy specialisation risks catastrophic forgetting, where training on a new speciality wipes out skills the model had before. The authors argue that broad general knowledge matters for rare and unusual road situations. They also point to a practical motive. In what they call cockpit-driving integration, the in-car assistant and the driving system share one computer, so one model may need to hold a conversation and drive competently.

How it works

The base VLM sees camera images as ordinary input, labelled with plain-text tags. Tags such as one marking the front view tell the model which of eight directions an image faces, and frame tags mark which moment in time it comes from. The VLM's architecture is not altered. The 3D information comes from an attached perception head. The authors describe it as a perception probe: an add-on that reads the model's internal features and makes visible how much 3D structure they contain. It takes six or eight surround-view images. It lifts the vision encoder's features into a 3D grid of small cubes, called voxels, using a depth network that guesses how far away each part of the image is. That depth network learns without being given any true depth. The head also takes the VLM's own processed image features and expands them into several scales. A transformer fuses the two streams into a bird's-eye view, or BEV: a top-down map of the area around the car.

Three branches read that shared map. One detects and places 3D objects such as vehicles and pedestrians. One predicts semantic occupancy, which labels every voxel as empty, drivable surface, vehicle, pedestrian and so on. One segments the road map. Separately, a Planning Expert produces the driving path. It uses flow matching, a technique that starts from random noise and refines it, here in ten steps, into a future trajectory for the car. At each step it is guided by information the main model has already computed about the scene.

Training happens in four stages. First, only the perception head learns, on top of the frozen VLM. Second, the head, vision encoder and VLM train together on perception labels plus driving and general question-answering data. In this stage the head's learning rate is 20 times the VLM's, so the new module adapts more quickly. Third, the VLM is frozen and only the Planning Expert learns from recorded trajectories. Some training prompts also include a written reason for the manoeuvre; the resulting model is called Qwen-Drive-1.0-SFT. Fourth, reinforcement learning, which improves a system by rewarding good outcomes, fine-tunes only the Planning Expert, giving Qwen-Drive-1.0-RL.

The reinforcement step needs the planner to try varied paths. The authors therefore inject a small amount of randomness, but only in the last three of the ten refinement steps. They restrict that randomness to smooth, low-frequency shapes, so the alternatives differ as manoeuvres rather than as jittery wobbles. Rewards come from each dataset's own driving score, plus a shared term for staying close to the human driver's actual path. The team merged perception labels from the

nuScenes and OpenScene datasets into a common set of 7 object classes, 10 occupancy classes and 6 map classes. They also filtered 24 public driving question-answering datasets from 5.53 million samples down to 3.09 million, and drew on 2.83 million planning samples from public sources.

What the evidence shows

On 3D perception, the authors compare their model against several established camera-based methods on nuScenes. They retrained these baselines themselves under the unified labels, including versions using the same Qwen-initialised image encoder. Qwen-Drive scores 43.95 mAP, a standard measure of how accurately objects are found and placed. That is about 2 points above the best baseline. It also reaches 60.99 on map segmentation overlap, 3.37 points ahead. One instructive result is a version that trains only the perception head on the frozen VLM. It trails the strongest baseline by 6.34 mAP. The authors read this as evidence that an unmodified VLM's features do not directly expose enough 3D structure, and that joint training is needed. The same model also works on OpenScene's eight-camera rig, which the comparison methods could not be evaluated on.

On driving question-answering and spatial reasoning, the model averages 69.43 across the benchmarks used, against 63.52 for the untouched Qwen3.5-4B. The next-best model, MiMo-Embodied-7B, scores 60.32. On Ego3D-Bench, its error in estimating distances is about 41 per cent lower than the base model's, although its accuracy on that benchmark is slightly lower. The largest gap appears on chain-of-causation questions. In these, the model must name which things in a scene caused a driving decision. Qwen-Drive's causal average is 58.30, against 22.05 for the runner-up, Gemma4-12B. The authors report that the model's general vision-language skills stay within about one point of the base model on knowledge, reasoning and recognition, and improve on spatial understanding and grounding.

Planning was tested in four settings of increasing realism. Open-loop tests predict a path and compare it with what the human did. On the Waymo end-to-end benchmark, WOD-E2E, the reinforcement-learning version reaches a Rater Feedback Score of 7.91, against 7.87 for the strongest rival, MindVLA-U1. That score compares a plan with several human-rated acceptable paths. NAVSIM runs a pseudo-closed-loop test, simulating the planned path once while other traffic replays its recording. There the model scores 90.7 on PDMS, a combined score covering collisions, staying on the road, progress, time-to-collision and comfort. ExploreVLA and EponaV2 both score 90.4. Adding written reasoning to the imitation-trained model lifted NAVSIM by 0.4.

What it does not show

The 3D perception claim holds mainly for object detection and maps. On nuScenes semantic occupancy, the authors' own multi-task baseline is clearly better, at 25.72 mIoU against 19.82. The perception baselines were reproduced by the team, trained only on nuScenes and limited to single frames. Detection is scored on remapped labels with a modified score, so the figures cannot be compared with public leaderboards. On a second planning dataset, PAI-AV, the model's average path error is 0.42 metres at three seconds, against 0.36 for NVIDIA's Alpamayo-1.5, which was trained on far more driving data. Reinforcement learning made PAI-AV errors slightly worse. The model's six candidate paths were also clustered closely around similar motions.

The toughest test is closed-loop simulation, where the model drives continuously and its mistakes compound. Across 916 scenarios in the AlpaSim simulator, reinforcement learning halved the off-road rate from 24 to 12 per cent. However, it also cut progress from 54 to 48 per cent. The model trails Alpamayo-R1, a model about twice its size, on close encounters, progress and overall score. Several headline margins are tiny: 0.04 on WOD-E2E and 0.3 on NAVSIM. One reported NAVSIM figure of 91.4 picks, for each scene, whichever of six plans scores highest on PDMS itself. The authors note that the WOD-E2E validation result, which beats the human-driver reference, is in-sample, because that split supplied the training reward.

The large causal-reasoning lead comes from a benchmark the authors built from the same source and format as their own training data. It was graded by a Qwen-family model, as was LingoQA. The authors also state that the model's planned paths do not always follow its written rationale. When causes at different distances coexist, such as a red light ahead and a child much closer, its reasoning may not capture which one matters now. Part of the planning gain from reasoning may come from extra internal information the generated text adds, rather than the reasoning itself. Their own data ablation found that the 3D supervision made only a small difference to planning. They say it does not establish 3D training as the source of the improvement.

Where it fits

The authors situate the work within the shift from modular driving pipelines, which hand results between separate perception, prediction and planning systems, towards unified learned systems such as UniAD. Among VLA models, they set it against recipes that adapt a VLM mostly with driving question-answering text, such as DriveLM, DriveGPT4 and OmniDrive. They also contrast it with designs such as UniDriveVLA, which split perception, understanding and planning into separate experts trained on driving data alone. Their alternative keeps the language model intact. It reuses established components, including depth-based lifting, transformer-based BEV fusion and detection decoders, as external heads, and mixes general data back in to protect existing skills.

Their written planning rationales borrow the chain-of-causation idea from Alpamayo-R1. Alpamayo-R1 and its successor Alpamayo-1.5 are also models Qwen-Drive trails in several planning tests. The authors describe their model as, to their knowledge, the first driving foundation model to unify 3D perception, question-answering and planning in one pretrained VLM. Taken on its own evidence, the paper shows that such heads can be attached to a 4-billion-parameter model without large losses in general ability. The resulting system is competitive on several planning benchmarks, though its claims to lead are best read benchmark by benchmark.

PAPER	*Qwen-Drive-1.0: An Initial Step towards a Vision-Language Foundation Model for Autonomous Driving*
AUTHORS	Xin Zhou, Zongchuang Zhao, Zhibo Yang, Mingsheng Li, Humen Zhong, Shuai Bai, Du Chu, Ruizhe Chen, Zhaohai Li, Jun Tang, Qiuyue Wang, Mingkun Yang et al.
VENUE	arXiv preprint — not yet peer reviewed · 31 August 2026
LINK	arxiv.org/abs/2609.00111
INDEX	stacks.braintank.ai
CODE	github.com/QwenLM/Qwen-Drive-1.0