J Intel Connect Veh

Research Article

VTF-Flow: Unified Terrain–Vehicle Kinematic Feasibility Guidance for Generative Trajectory Planning in Unstructured Off-Road Environments

Linheng Li[1], Weijie Wang[1], Tiantong Nan[1], Jing Gan[2], Xu Qu[1, ✉], Bin Ran[3]

Cite this article: https://doi.org/10.26599/JICV.2026.9210112

ABSTRACT: Trajectory planning in unstructured off-road environments must handle weak road boundaries, continuous terrain variation, and multimodal local motion. Although generative planners can represent multiple futures, terrain feasibility is often treated as a static mask or post-generation check, separating it from candidate-induced motion states and trajectory geometry. We propose VTF-Flow, a Vehicle-Conditioned Terrain Feasibility-Guided Flow Planning framework that embeds a unified terrain–vehicle kinematic potential into trajectory generation. It constructs a continuous static potential from LiDAR geometry and semantic labels, conditions it on candidate speed and reliable heading, and adds soft curvature and lateral-acceleration penalties. A bounded feasibility score regularises Flow Matching during training, while an endpoint-preserving trajectory-coordination operator injects the potential gradient into every inference-time integration step. Across 1,909 paired diagnostics on RELLIS-3D sequence 00004, all controlled perturbations reduced relative feasibility; obstacle-directed perturbations achieved a 77.58% win rate, and five-step gradient descent increased feasibility while reducing occupancy exposure. In a sequence-level outer-holdout benchmark of 7,063 planning scenes, VTF-Flow achieved the lowest minADE@8 and unified TVK potential, as well as the highest validation-derived q80 candidate compliance and goal-conditioned compliant coverage at K=8. Relative to matched unguided Flow Matching, it reduced ADE-0, minADE@8, the TVK potential, and curvature violation by 1.31%, 1.07%, 0.99%, and 17.70%, respectively; the principal feasibility metrics improved in all nine seed–held-out-sequence pairs. Continuous in-flow correction therefore improved the availability of low-potential, goal-consistent trajectories while retaining the learned multi-candidate prior.

KEYWORDS: Unstructured Off-road Environments; Off-road Trajectory Planning; Flow Matching; Vehicle-conditioned Feasibility; Generative Motion Planning; Terrain Traversability

1 Introduction

Autonomous vehicles deployed in agricultural and forestry operations, mining transport, disaster response, and field inspection must often operate beyond structured traffic environments with clearly marked lanes, regular road boundaries, and stable surface conditions. Unlike urban roads, the available motion space in unstructured off-road environments is jointly shaped by vegetation, mud, gravel, slopes, ditches, and local obstacles. Its boundaries are discontinuous and vary with the observable region. Natural scenes captured in RUGD, RELLIS-3D, and WildScenes further demonstrate that such environments combine diverse terrain classes, irregular semantic boundaries, pronounced geometric relief, and sparse sensor observations (Jiang et al., 2020; Vidanapathirana et al., 2025; Wigness et al., 2019). Off-road trajectory planning must therefore address more than how to reach a local goal. Within a limited sensing horizon, it must also coordinate terrain exposure, vehicle operating state, trajectory geometry, and the diversity of candidate routes.

Existing off-road navigation systems typically convert elevation, slope, roughness, semantic class, or obstacle information into grid-based traversability or risk costs, after which graph search, sampling, trajectory optimisation, or model predictive control is used to obtain a local path (Fan et al., 2021; Jiang et al., 2025; Khatib, 1986; Meng et al., 2023; Zhu et al., 2020). With the development of data-driven approaches, GA-Nav, RoadRunner, and EVORA have learned off-road operating costs from semantic navigability, multimodal bird's-eye-view traversability, and uncertainty-aware traction, respectively (Cai et al., 2024; Frey et al., 2024; Guan et al., 2022). These methods allow environmental risk to enter the planning process explicitly. Nevertheless, their terrain representations are often reduced to binary masks, discrete cost levels, or static position-dependent costs that are separated from the motion state of a particular candidate. In practice, the suitability of the same terrain region changes with speed, direction of travel, and trajectory curvature. A binary traversable/non-traversable description cannot fully express the continuous difference between crossing rough terrain at low and high speed, traversing a slope laterally, or turning through a narrow clearance. Recent work on motion-aware traversability likewise indicates that terrain cost should vary with the actual motion state rather than remain a fixed spatial scalar (Zhao et al., 2026). A trajectory planner therefore requires a feasibility representation that extends beyond position-dependent terrain cost to a continuous, candidate-dependent operating potential.

1 School of Transportation, Southeast University, Nanjing 211189, China. 2 School of Modern Posts & School of Intelligent Transportation, Nanjing University of Posts and Telecommunications, Nanjing 211189, China. 3 Department of Civil and Environmental Engineering, University of Wisconsin–Madison, Madison WI 53706, USA.

✉ Corresponding author. E-mail: quxu@seu.edu.cn

Received: September 4, 2026; Revised: September 21, 2026; Accepted: October 28, 2026

© The Author(s) 2026. This is an open access article under the terms of the Creative Commons Attribution 4.0 International License (CC BY 4.0, http://creativecommons.org/licenses/by/4.0/).

In Accepted

Unstructured environments also rarely provide a unique reference centreline. Given the same local goal, a vehicle may pass on either side of an obstacle or adopt different trade-offs between terrain cost and travel distance. Single-output regression tends to reproduce the conditional mean of the demonstrations, whereas graph search and local optimisation are largely determined by prescribed costs and initialisation. Neither formulation readily represents multiple reasonable driving choices within one forward pass. Research on multimodal trajectory prediction and generative planning has shown that directly modelling a trajectory distribution can produce candidate sets with distinct geometric modes (Chai et al., 2020; Chi et al., 2023; Janner et al., 2022; Jiang et al., 2023; Phan-Minh et al., 2020). Diffusion models and Flow Matching have subsequently been applied to autonomous driving trajectory generation to improve candidate coverage, goal conditioning, interaction modelling, and sampling efficiency (Liao et al., 2025; Tan et al., 2025; Xing et al., 2025; Zheng et al., 2025). Flow Matching learns a continuous transport from a simple base distribution to the data distribution by regressing a vector field along conditional probability paths, without numerically simulating those paths during training (Lipman et al., 2022). Rectified Flow further shows that nearly straight probability paths facilitate efficient transport under coarse temporal discretisation (Liu et al., 2023). These developments provide a direct basis for multi-candidate generation in trajectory space. However, a standard generative objective primarily fits the trajectory distribution and does not, by itself, ensure low operating cost with respect to local slope, roughness, obstacle proximity, and vehicle motion state.

Controllable sampling in generative planning offers a means of reconciling a data-driven prior with explicit operating constraints. Diffuser and MotionDiffuser demonstrate that gradients of rewards or differentiable costs can alter trajectory samples during iterative generation (Janner et al., 2022; Jiang et al., 2023). FISOR studies safety constraints in diffusion policies through feasible-set characterisation and energy-guided sampling (Zheng et al., 2024). Diffusion Planner, FeaXDrive, and GuideFlow further incorporate driving rules, drivable areas, curvature, or other physical constraints into generative vehicle planning (Liu et al., 2026; Wang et al., 2026; Zheng et al., 2025). Collectively, these studies indicate that feasibility need not be restricted to filtering or ranking candidates after generation. Existing vehicle-planning approaches, however, predominantly address structured urban roads, for which constraints are derived from road masks, traffic rules, or predefined physical terms. Off-road planning instead requires a continuous potential constructed from measured three-dimensional terrain attributes, conditioned on candidate vehicle states, and differentiable along the complete trajectory. The guidance mechanism must additionally suppress local zigzags, candidate collapse, and goal drift that can arise from independently applying pointwise gradients.

To address this gap, we propose VTF-Flow, a Vehicle-Conditioned Terrain Feasibility-Guided Flow Planning framework. Inspired by vector-field regression along conditional probability paths in Flow Matching (Lipman et al., 2022) and goal-directed multi-trajectory generation (Xing et al., 2025), VTF-Flow learns a goal-conditioned generative prior in planning-ego trajectory space and produces multiple future candidates from distinct Gaussian initial states. The central contribution is not a new Flow Matching formulation, but a hierarchical terrain–vehicle kinematic feasibility potential that acts directly on the generative dynamics. Occupancy, non-traversability, slope, roughness, and obstacle proximity are first combined into a continuous static terrain-feasibility potential field. The terrain cost is then adjusted according to the speed and reliable heading induced by each candidate, and augmented with soft curvature and lateral-acceleration penalties to obtain the unified trajectory potential J_{TVK} . During training, its bounded relative feasibility score provides a soft regulariser. During inference, the gradient of J_{TVK} is evaluated on the current clean trajectory estimate at every Flow integration step, processed by the trajectory-coordination operator, and injected into the state update after endpoint-preserving projection. Feasibility therefore continuously modifies the direction of trajectory evolution during generation rather than acting as an independent validity check after generation.

The main contributions of this study are as follows:

- We formulate a hierarchical continuous feasibility representation for trajectory planning in unstructured off-road environments. Measured terrain geometry, semantics, and obstacle information are fused into a spatial potential and combined with candidate speed, reliable heading, curvature, and lateral acceleration to form a unified terrain–vehicle kinematic trajectory potential. This formulation extends a fixed position-dependent cost to a candidate-dependent operating cost.
- We introduce VTF-Flow, which combines bounded feasibility regularisation during training with endpoint-preserving in-flow guidance at inference. The feasibility gradient participates in state evolution at each numerical integration step. The trajectory-coordination operator applies trajectory-wise smoothing, per-candidate scale normalisation, low-gradient suppression, and trajectory-level norm clipping, while endpoint-preserving projection maintains global trajectory structure and goal conditioning. The resulting procedure is distinct from post-generation path replacement or candidate filtering.
- We establish a claim-aligned hierarchical evaluation protocol on RELLIS-3D. Controlled trajectory perturbations, gradient-direction tests, and local descent experiments first assess whether the feasibility potential is discriminative and numerically suitable for guidance. A sequence-level outer-holdout protocol then compares VTF-Flow with motion extrapolation, terrain-grid search, deterministic learning, and unguided Flow Matching, while component ablations isolate the roles of vehicle-state conditioning, the kinematic potential, training regularisation, and in-flow guidance.

The remainder of this paper is organised as follows.

Just Accepted

Section 2 reviews off-road terrain-feasibility representations, trajectory planning, and constraints imposed during generative sampling. Section 3 presents the problem formulation, the goal-conditioned Flow backbone, the hierarchical terrain–vehicle kinematic feasibility potential, bounded training regularisation, and endpoint-preserving in-flow guidance. Section 4 describes the processing of RELLIS-3D, trajectory construction, the sequence-level evaluation protocol, comparison methods, and evaluation metrics. Section 5 reports the preliminary potential diagnostics, unified benchmark, component ablations, and qualitative analysis in real scenes. Finally, Section 6 summarises the main findings, scope of applicability, and directions for future work.

2 Related Work

2.1 Off-Road Terrain Representation and Traversability Assessment

Off-road traversability assessment seeks to infer the relative cost of crossing different regions from observations of the natural terrain surrounding a vehicle. Classical approaches typically extract slope, elevation change, surface normals, and height variance from two-dimensional or 2.5D elevation maps, and combine them with obstacle distance or manually defined rules to construct cost maps. Artificial potential fields express goal attraction and obstacle repulsion as continuous spatial functions, providing a classical basis for gradient-based local navigation (Khatib, 1986). STEP combines collision, slope, roughness, and uncertainty in a 2.5 D terrain-risk map and solves trajectories using risk-aware geometric planning and kinematic model predictive control (Fan et al., 2021). For wheeled unmanned ground vehicles, risk-aware global–local planning has incorporated potential-field risk assessment into Coarse2fine A· and deterministic local sampling (Jiang et al., 2025). Real-time planning on deformable terrain has further used terramechanics models and online parameter estimation to adapt model predictive control (Yu et al., 2025). These methods are interpretable, but their performance depends on the selected terrain features, thresholds, vehicle models, and numerical optimisation settings.

The emergence of off-road datasets and multimodal perception has motivated learning-based traversability estimation from semantic and geometric cues. RUGD provides pixel-wise semantic annotations of unstructured natural scenes; RELLIS-3D adds synchronised camera images, LiDAR scans, poses, and off-road semantic labels; and WildScenes extends two-dimensional–three-dimensional semantic correspondences to large-scale natural environments (Jiang et al., 2020; Vidanapathirana et al., 2025; Wigness et al., 2019). At the image level, GA-Nav uses group-wise attention to organise off-road semantic classes into navigability levels (Guan et al., 2022), whereas self-supervised traversability prediction based on reconstruction error reduces dependence on manual labels by exploiting previously traversed regions (Schmid et al., 2022). In bird's-eye view and three dimensions, TerrainNet converts semantic and geometric predictions into planning cost maps (Meng et al., 2023); RoadRunner directly predicts elevation and traversability from cameras and LiDAR (Frey et al., 2024); and ORDformer uses semantic scene completion to estimate dense three-dimensional traversable occupancy, with RELLIS-OCC labels incorporating slope, unevenness, and vehicle geometry constraints (Chen et al., 2024). In parallel, maximum-entropy inverse reinforcement learning has been used to infer off-road traversability costs from expert trajectories (Zhu et al., 2020), while EVORA models terrain traction together with aleatoric and epistemic uncertainty for risk-aware planning (Cai et al., 2024). These studies have advanced the field from binary road segmentation towards continuous costs, three-dimensional occupancy, and risk-aware representations. Nevertheless, perception outputs, trajectory-distribution learning, and candidate-level feasibility correction usually remain separate computational stages.

Vehicle-state-dependent terrain cost is the line of research most closely related to the present study. Motion-aware Traversability models the cost of each terrain region as a function of speed, representing situations in which the outcome of crossing a ditch changes with the commanded motion state (Zhao et al., 2026). This finding independently supports the view that traversability should not depend on position alone. VTF-Flow addresses a different problem: it combines speed- and reliable-heading-dependent terrain cost with soft curvature and lateral-acceleration penalties to define a candidate trajectory potential, and studies how this potential can provide structure-preserving correction during the evolution of a learned multi-candidate Flow planner.

2.2 From Search and Optimisation to Learning-Based Multi-Candidate Trajectory Planning

Classical motion planning for vehicles and mobile robots encompasses graph search, sampling-based planning, and trajectory optimisation. A· searches for a minimum-cost path in a discrete graph using a heuristic cost (Hart et al., 1968). Artificial potential fields realise goal attraction and obstacle repulsion through continuous gradients (Khatib, 1986). CHOMP, STOMP, and TrajOpt address smoothness, collision, and other trajectory costs through covariant gradients, stochastic trajectory perturbations, and sequential convex optimisation, respectively (Kalakrishnan et al., 2011; Ratliff et al., 2009; Schulman et al., 2013). For vehicles with nonlinear dynamics and complex costs, information-theoretic model predictive control uses path-integral sampling for receding-horizon optimisation and has been demonstrated in aggressive driving (Williams et al., 2017). These methods permit explicit terrain-risk and vehicle constraints, but their outputs can be sensitive to map discretisation, candidate libraries, dynamics models, initialisation, and manually selected weights. When several acceptable routes exist, a single optimisation run typically returns only one local solution, whereas repeated search increases computational and parameter-design burdens.

Learning-based trajectory models absorb driving priors from demonstrations (Farahnakian et al., 2025; Ndunguru et

Just Accepted

al., 2025). Deterministic regression is simple and can reproduce recorded vehicle futures, but may collapse distinct modes towards a conditional mean when the future is multimodal. MultiPath and CoverNet represent multiple futures using distributions over trajectory anchors and classification over trajectory sets, respectively (Chai et al., 2020; Phan-Minh et al., 2020). These methods were developed primarily for traffic-participant motion prediction, and their output forms and evaluation objectives are not identical to off-road ego-vehicle planning. They nevertheless demonstrate why a single average trajectory is insufficient to describe multiple plausible futures.

Generative trajectory models instead formulate planning as sampling from a trajectory distribution. Diffuser unifies sampling and planning through iterative denoising, Diffusion Policy applies conditional diffusion to multimodal robot action sequences, and MotionDiffuser supports differentiable cost constraints over multi-agent trajectory distributions (Chi et al., 2023; Janner et al., 2022; Jiang et al., 2023). For autonomous driving, Diffusion Planner uses classifier guidance to realise flexible closed-loop behaviour, while DiffusionDrive combines truncated diffusion with multimodal anchors to improve sampling efficiency (Liao et al., 2025; Zheng et al., 2025; Cheng et al., 2025). GoalFlow employs goal-directed Flow Matching and candidate scoring for multimodal trajectory generation, whereas Flow Planner combines segment-level trajectory representations, scene-interaction fusion, and classifier-free guidance to improve closed-loop interactive behaviour (Tan et al., 2025; Xing et al., 2025). These Flow-based planners build on general developments in conditional continuous transport and nearly straight probability-path learning (Lipman et al., 2022; Liu et al., 2023). VTF-Flow adopts the ability of generative models to represent multi-candidate trajectory distributions, but restricts the planning object to off-road future trajectories in the planning-ego coordinate frame. Its focus is not a new general-purpose generative backbone, but the direct influence of measured terrain and candidate-induced motion states on the direction of trajectory transport.

2.3 Feasibility Constraints and Guidance in Generative Planning

A data-driven generative prior can be coupled with explicit planning objectives through conditioning, training regularisation, guided sampling, or post-generation optimisation. Diffuser uses classifier guidance and conditional inpainting to satisfy task objectives; MotionDiffuser introduces gradients of differentiable losses during sampling to incorporate rules and physical priors; and Motion Planning Diffusion combines a learned trajectory prior with gradients of collision and other costs for robot motion planning (Carvalho et al., 2024; Janner et al., 2022; Jiang et al., 2023). FISOR uses reachability analysis to characterise the feasible region supported by offline data and employs an energy-guided diffusion policy to coordinate reward, behaviour regularisation, and safety requirements (Zheng et al., 2024). These methods demonstrate that external differentiable objectives can act on generation in trajectory or action space. Whether the guidance signal is continuous, conditioned on vehicle state, and able to preserve trajectory structure under pointwise gradients nevertheless depends on the application design.

Recent work on structured-road driving has begun to treat trajectory feasibility explicitly within generative planning. FeaXDrive uses predicted clean trajectories as a common representation and combines curvature-constrained training, drivable-area guidance during reverse diffusion, and post-training optimisation (Wang et al., 2026). GuideFlow jointly trains Flow Matching and an energy model, and imposes safety and physical constraints during generation (Liu et al., 2026). These methods share an important premise with the present work: constraints should act on trajectory generation itself rather than only selecting among completed candidates. Their primary setting is end-to-end planning on structured roads with road boundaries, traffic rules, and urban-driving metrics. Off-road environments lack stable drivable-area boundaries, and their operating costs additionally depend on slope, roughness, semantic terrain, obstacle clearance, and candidate speed and heading.

The proposed VTF-Flow further focuses this line of research on continuous terrain–vehicle kinematic coordination in unstructured off-road environments. Its static terrain-feasibility potential is constructed from real-world LiDAR geometry and semantic information provided by RELLIS-3D, rather than directly treating a road or traversability mask as a safety label. The vehicle-state-dependent component, together with the curvature and lateral-acceleration terms, is consistently evaluated on the denoised trajectory estimate. The same potential is incorporated into training regularization through a bounded relative score and, during inference, is continuously injected into the Flow integration through an endpoint-preserving trajectory-coordination operator.

3 Method

3.1 Problem formulation and framework overview

Unstructured off-road environments generally lack explicit lane centrelines, regular road boundaries, and predefined drivable corridors. The space available to a vehicle is therefore continuous, open, and strongly dependent on local terrain conditions. For the same local goal, a vehicle may follow trajectories with different directions and geometric forms. Their suitability for vehicle motion depends not only on spatial location, but also on terrain slope, surface roughness, obstacle distribution, and the motion state induced by each candidate. Accordingly, this study addresses the generation of multiple goal-directed vehicle trajectories and the coordination of terrain conditions and vehicle kinematics during generation, rather than the geometric fitting of a single reference path.

Within a complete intelligent-vehicle planning stack, VTF-Flow serves as a local reference-trajectory generator. Its candidate set must be passed to an online candidate selector and a downstream trajectory-tracking controller. The present model does not directly issue steering, braking, or propulsion commands. This role supports short-horizon navigation in agricultural transport, mining-

Just Accepted

road travel, field inspection, and emergency mobility, provided that platform-specific dynamics and controller constraints are enforced downstream.

At the current planning time t, the vehicle receives a local goal $\mathbf{g}_t$ and an ego-centred local terrain representation $\mathbf{M}_t$. The planner produces K candidate future trajectories,

$$\mathcal{S}_t = \{\mathbf{g}_t, \mathbf{M}_t\}, \qquad \boldsymbol{\tau}_t^{(k)} = \left\{\mathbf{p}_{t+i}^{(k)}\right\}_{i=1}^{H}, \qquad \mathbf{p}_{t+i}^{(k)} = \left(x_{t+i}^{(k)}, y_{t+i}^{(k)}, z_{t+i}^{(k)}\right), \tag{1}$$

where H is the number of discrete future points within the planning horizon. All trajectories are expressed in the current planning-ego coordinate frame. The current ego position is mapped to the origin, with the x-axis pointing forward, the y-axis pointing to the left of the vehicle, and the z-axis pointing upward. Although the software interface retains an ego-history input, the frozen implementation evaluated in this study does not use measured ego history as an effective condition. The planning input is therefore restricted to the local goal and terrain representation.

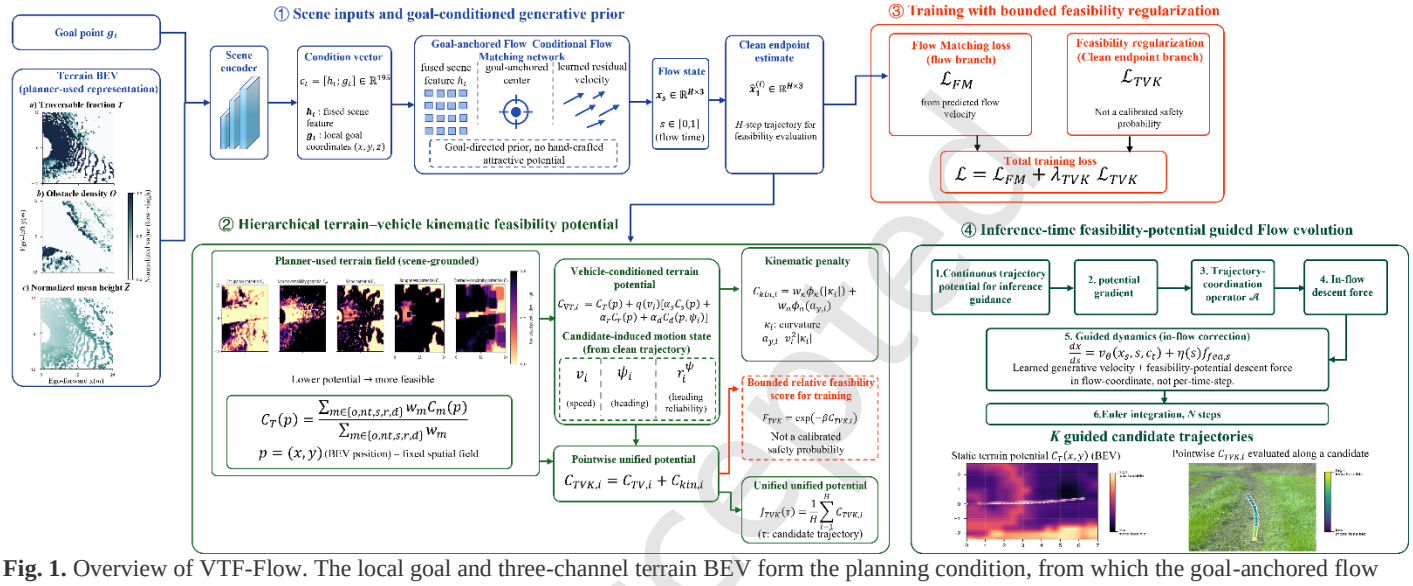


Fig. 1. Overview of VTF-Flow. The local goal and three-channel terrain BEV form the planning condition, from which the goal-anchored flow network generates multiple future vehicle trajectories. The terrain BEV is converted into a continuous spatial potential from occupancy, non-traversability, slope, roughness, and obstacle proximity. Each candidate trajectory further induces speed, reliable heading, curvature, and lateral acceleration, resulting in a unified trajectory-dependent potential. During training, a bounded feasibility score regularises the clean trajectory estimate. During inference, the gradient of the unified trajectory potential is evaluated at every Euler step, coordinated along the trajectory, projected to preserve the goal-anchored terminal waypoint, and applied directly to the current flow state. Feasibility therefore modifies the trajectory-generation dynamics rather than acting as a post-generation filter. The local terrain and trajectories in the diagram illustrate information flow and do not constitute a quantitative comparison.

We propose VTF-Flow, short for Vehicle-Conditioned Terrain Feasibility-Guided Flow Planning. Instead of treating terrain feasibility as a validity check after trajectory generation, VTF-Flow formulates a differentiable terrain–vehicle kinematic feasibility potential and embeds its gradient directly into the generative dynamics of Flow Matching. The conditional Flow Matching backbone first learns a data-driven trajectory prior from demonstrated vehicle motion. The local terrain is then converted into a continuous spatial potential field and extended to a trajectory-dependent potential using the speed, reliable heading, curvature, and lateral acceleration induced by each candidate. At every numerical integration step, the planner recovers a clean trajectory estimate from the current flow state, evaluates the unified trajectory potential, and uses its gradient to steer the same state update towards a lower-cost region. Thus, the generative backbone identifies trajectories that are consistent with the demonstrated driving distribution, whereas the feasibility branch identifies how the potential decreases and how the current candidate should be corrected.

Inspired by the simulation-free vector-field regression over conditional probability paths introduced by Flow Matching (Lipman et al., 2022), we formulate the transport process in goal-conditioned trajectory space and use it as the multi-candidate generation backbone. Building on this generative prior, the methodological advance of VTF-Flow consists of three connected levels. First, a hierarchical terrain–vehicle kinematic potential-field representation encodes occupancy, non-traversability, slope, roughness, and obstacle proximity as a continuous spatial field and lifts it to a trajectory-dependent potential using candidate speed, reliable heading, curvature, and lateral acceleration. Second, conditional generation is performed directly in trajectory space, allowing scene and kinematic constraints to act on the entire future trajectory instead of on isolated locations or a post-generation ranking stage. Third, the gradient of the unified trajectory potential with respect to the current flow state is processed by a trajectory-coordination operator and injected throughout numerical integration. VTF-Flow therefore combines a learned trajectory prior with in-flow potential-field correction, in contrast to a pipeline based on a static traversability map followed by post-generation checking.

The complete framework is shown in Fig. 1. The blue

Just Accepted

branch denotes stochastic trajectory generation conditioned on the local goal and terrain bird's-eye-view (BEV) representation. The green branch constructs the unified terrain–vehicle kinematic feasibility. The red branch introduces feasibility regularisation during training, and the teal branch performs feasibility-guided flow evolution during inference. Sections 3.2 and 3.3 describe the generative process and feasibility representation, respectively. Sections 3.4 and 3.5 then specify how these components are coupled during training and inference.

The Conditional Flow Matching backbone supplies the stochastic transport model and the goal-conditioned multi-candidate prior. It maps Gaussian initial trajectory states to the distribution of recorded future trajectories through the learned conditional velocity field. The contributions of this study begin after this backbone is defined: the hierarchical $\mathbf{c}_T$ - $\mathbf{c}_{VT}$ - $\mathbf{c}_{TVK}$ feasibility representation, the bounded F_{TVK} training regulariser, and the endpoint-preserving gradient update based on J_{TVK} during inference. These additions do not alter the Conditional Flow Matching probability path or target velocity.

3.2 Goal-conditioned trajectory-space generation

In off-road environments without lane topology, a local goal specifies the general direction of travel but does not uniquely determine a traversable path. To retain both goal directedness and candidate diversity, VTF-Flow models the future trajectory distribution using goal-anchored conditional Flow Matching. A scene encoder maps the terrain BEV and local goal to a condition vector $\mathbf{c}_t$. The terrain features describe the locally available traversal space, whereas the metric goal coordinates specify the desired travel direction and distance of the current planning task.

We use $s \in [0,1]$ to denote flow time and distinguish it from the physical time t of the traffic scene. Let the clean supervised trajectory be $\mathbf{x}_1 = \boldsymbol{\tau}_t^{\mathrm{gt}}$, and sample the initial random trajectory state from $\mathbf{x}_0 \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$. Training constructs a linear probability path and regresses the corresponding transport velocity (Lipman et al., 2022):

$$\begin{aligned} x_s &= (1-s)x_0 + sx_1, \\ \mathcal{L}_{\mathrm{FM}} &= \mathbb{E}\left[\left\| v_\theta(x_s, s, c_t) - (x_1 - x_0) \right\|_2^2\right]. \end{aligned} \tag{2}$$

A velocity field conditioned only through generic scene features may not retain a sufficiently strong directional constraint over a long planning horizon. VTF-Flow therefore introduces a goal-derived reference transport and learns spatial and velocity residuals relative to that reference:

$$\begin{aligned} \mu_{\theta,i} &= \frac{i}{H} g_t + r_{\theta,i}(c_t), \\ v_\theta &= \frac{\mu_\theta - x_s}{\max(1-s, \epsilon)} + \Delta v_\theta(x_s, s, c_t), \qquad i = 1, \ldots, H. \end{aligned} \tag{3}$$

This parameterisation encourages overall progress towards the local goal without hard-constraining the final waypoint to coincide with it, thereby preserving room for obstacle avoidance and terrain adaptation. At an arbitrary flow time, the clean trajectory estimate is recovered as

$$\hat{\mathbf{x}}_1^{(s)} = \mathbf{x}_s + (1-s)\mathbf{v}_\theta(\mathbf{x}_s, s, \mathbf{c}_t). \tag{4}$$

This estimate represents the complete future vehicle trajectory implied by the current flow state and provides the common trajectory-space quantity from which speed, heading, and feasibility costs are evaluated. During inference, parallel integration from different Gaussian initial states produces K candidates. Their differences represent variability learned from the trajectory distribution and are not assigned predefined semantic driving-mode labels.

Equation (4) first recovers the clean trajectory estimate from the current flow state and the predicted conditional velocity. During training, the terrain and kinematic quantities are evaluated on this clean estimate. The feasibility gradient therefore propagates through the bilinear terrain queries and the finite-difference speed, reliable-heading, curvature, and lateral-acceleration calculations to the clean estimate, and then through Equation (4) and the conditional velocity field to the network parameters. During inference, automatic differentiation evaluates the gradient of J_{TVK} with respect to the current flow state through the same chain. The trajectory-coordination operator and endpoint projection are then applied before the correction enters the Euler update. No GT waypoint is used in this inference-time gradient pathway.

3.3 Hierarchical terrain–vehicle kinematic feasibility potentials

To clarify the dependency hierarchy and the distinct training and inference roles of the five feasibility quantities, Table 1 summarises their definitions, primary dependencies, and operational roles.

where H is the number of discrete future points within the planning horizon. All trajectories are expressed in the current planning-ego coordinate frame. The current ego position is mapped to the origin, with the x-axis pointing forward, the y-axis pointing to the left of the vehicle, and the z-axis pointing upward. Although the software interface retains an ego-history input, the frozen implementation evaluated in this study does not use measured ego history as an effective condition. The planning input is therefore restricted to the local goal and terrain representation.

Just Accepted

Table 1. Definitions, dependencies, and operational roles of the hierarchical feasibility quantities.

Symbol	Definition	Primary dependencies	Operational role
C_T	Static terrain cost	Occupancy, traversability, slope, roughness, and proximity	Continuous BEV potential
C_{VT}	Vehicle-conditioned terrain cost	C_T, speed, reliable heading, and look-ahead	Candidate-dependent terrain cost
C_{TVK}	Unified pointwise cost	C_{VT}, curvature, and lateral acceleration	Terrain and kinematic integration
J_{TVK}	Trajectory-level TVK potential	Mean C_{TVK} along one candidate	Training diagnostic and inference objective
F_{TVK}	Bounded relative feasibility score	Monotone mapping of C_{TVK}	Training regulariser; not a safety probability

3.3.1 Continuous terrain-feasibility potential field

To make generated trajectories responsive to off-road terrain, the planner uses a three-channel terrain BEV represented in the same local coordinate frame as the trajectories:

$$\mathbf{M}_t = \{T, O, \bar{Z}\}, \tag{5}$$

where T is the fraction of traversable points within each grid cell, O is the obstacle-point density, and $\bar{Z}$ is the normalised mean elevation. Point-wise semantic labels are used only during offline construction of T and O; the planner does not receive a separate semantic channel. Unobserved cells are represented conservatively through a low traversable fraction rather than through an additional learned safety label. These base channels are used to derive the occupancy cost C_o, the non-traversability cost C_{nt}, the slope cost C_s, the roughness cost C_r, and the obstacle-proximity cost C_d. The components describe whether the local space is occupied, whether the surface supports vehicle traversal, the local elevation gradient, the degree of surface variation, and the available spatial margin to nearby obstacles, respectively.

After normalisation to comparable scales, the terrain-feasibility potential field and its corresponding continuous feasibility score are

$$C_T(p) = \frac{\sum_{m\in\{o,nt,s,r,d\}} w_m\, C_m(p)}{\sum_m w_m}, \qquad F_T(p) = \exp[-\beta C_T(p)]. \tag{6}$$

where $w_m \geq 0$ weights each terrain factor and $\beta > 0$ controls the mapping from cost to feasibility. Terrain attributes at trajectory points are queried by bilinear interpolation. Consequently, C_T and F_T remain continuous with respect to the planar trajectory coordinates and can provide spatial gradients. Under a potential-field interpretation, the scalar function $C_T(\mathbf{p})$ is the local terrain-feasibility potential: its value describes the relative operating cost of a location, and $-\nabla_{\mathbf{p}} C_T$ gives a local direction of decreasing terrain cost. The bounded value $F_T \in (0,1]$ is only a monotone mapping of this potential and represents a relative traversal preference, not a calibrated probability of safety. Unlike a classical artificial potential field (Khatib, 1986), C_T contains no hand-crafted goal-attraction term. Goal directedness is supplied by the conditional generative backbone in Section 3.2, whereas the terrain potential supplies only the feasibility correction. This separation preserves the learned driving prior rather than replacing it with manually specified attraction–repulsion rules.

3.3.2 Vehicle-state-conditioned terrain-feasibility potential

The same terrain location need not have the same operational implication under different vehicle motion states. Traversing a slope or rough surface at higher speed generally warrants a larger penalty than traversing it slowly, and the desired forward clearance may also increase with speed. VTF-Flow therefore does not treat terrain cost as a static position-only map. Instead, it conditions the cost on the motion state induced by each candidate trajectory.

Finite differences between consecutive points of the clean trajectory estimate yield the candidate speed v_i and tangential heading ψ_i at each planning step. When the displacement between adjacent points is small, a continuous heading-reliability term attenuates directional information so that unstable heading estimates at near-zero speed do not dominate the terrain query. After querying look-ahead obstacle proximity along the reliable heading direction, the vehicle-conditioned terrain cost is written as

$$\begin{aligned} C_{\mathrm{VT},i} &= C_T(p_i) + q(v_i)\big[\alpha_s C_s(p_i) + \alpha_r C_r(p_i) \\ &\quad + \alpha_d C_{d,i}^{\mathrm{off}}(p_i, \psi_i)\big], \\ F_{\mathrm{VT},i} &= \exp\big[-\beta C_{\mathrm{VT},i}\big]. \end{aligned} \tag{7}$$

where $q(v_i)$ is a non-negative differentiable function that increases with speed. The coefficients α_s, α_r, and α_d control the speed-dependent amplification of slope, roughness, and obstacle-proximity costs. Consequently, $C_{\mathrm{VT}}(\mathbf{p}; v, \psi)$ defines a vehicle-state-conditioned terrain potential: the same spatial location may receive different potential values for candidates with different speeds or headings. The complete vehicle-conditioned potential therefore cannot be interpreted as a fixed two-dimensional heatmap independent of the candidate; only its static spatial component $C_T(x, y)$ can be displayed directly in BEV.

3.3.3 Trajectory kinematic cost and unified trajectory potential

Pointwise terrain costs alone do not describe the geometric continuity of an entire trajectory or the coupling between path geometry and speed. We construct a planar path through the current ego origin and future waypoints, estimate the curvature κ_i from consecutive point triplets, and compute the lateral acceleration proxy $a_{y,i} = v_i^2 |\kappa_i|$. Curvature and lateral acceleration are converted into dimensionless penalties through differentiable soft-excess functions $\phi_\kappa(\cdot)$ and $\phi_a(\cdot)$ relative to nominal limits. When either adjacent segment is too short to support a stable turning estimate, a continuous displacement-reliability gate attenuates κ_i and the corresponding lateral-acceleration proxy. This prevents near-stationary pose jitter from dominating the unified objective while retaining differentiability. The penalty remains small below the corresponding nominal limit and increases continuously beyond it. The unified cost and bounded feasibility score are defined as

$$\begin{aligned} C_{\mathrm{TVK},i} &= C_{\mathrm{VT},i} + w_\kappa \phi_\kappa(|\kappa_i|) + w_a \phi_a\big(a_{y,i}\big), \\ F_{\mathrm{TVK},i} &= F_{\mathrm{VT},i} \exp\big[-w_\kappa \phi_\kappa(|\kappa_i|) - w_a \phi_a\big(a_{y,i}\big)\big], \end{aligned} \tag{8}$$

Just Accepted

where w_κ and w_a control the contributions of curvature and lateral acceleration. Averaging $C_{\mathrm{TVK},i}$ along a candidate gives the unified trajectory potential J_{TVK}. Accordingly, C_T represents the static spatial potential field, C_{VT} represents the state-conditioned terrain potential, and J_{TVK} represents the generalised potential of an entire candidate under the current scene and motion state. This hierarchy jointly captures terrain location, candidate speed and heading, and curvature–speed coupling while remaining differentiable with respect to trajectory coordinates.

Importantly, F_{TVK} is a relative continuous feasibility score and not a calibrated safety probability. Curvature and lateral acceleration are kinematic proxies that can be estimated consistently from a candidate trajectory.

More specifically, F_{TVK} is a bounded monotone transformation that expresses relative preference among candidates under the selected potential weights. It is not calibrated against collision or failure events and must not be interpreted as a safety probability. The curvature and lateral-acceleration terms are trajectory-derived kinematic proxies. Without vehicle mass, steering-system, suspension, tyre-ground friction, and roll-stability parameters, they do not establish full dynamic feasibility or stability.

The unified feasibility representation assumes complementary forms during training and inference. Training uses the bounded score F_{TVK} as a stable soft regulariser that changes the model's relative preference among trajectory regions. Inference uses the continuous cost J_{TVK} to retain informative spatial gradients for in-flow guidance. Thus, the representation defined in this section connects directly to the joint training objective in Section 3.4 and the guided flow dynamics in Section 3.5.

3.4 Flow Matching training with unified feasibility regularisation

VTF-Flow retains the original Flow Matching objective and introduces unified terrain–vehicle kinematic feasibility as an additional soft regulariser rather than redefining the generative task. For each training iteration, a Gaussian initial state and a flow time are sampled to form the linear probability path, after which the conditional velocity is predicted. Equation (4) then recovers the clean trajectory estimate, from which speed, reliable heading, curvature, and lateral acceleration are computed. The value of F_{TVK} is queried at every trajectory point, giving the total training objective

$$\mathcal{L} = \mathcal{L}_{\mathrm{FM}} + \lambda_{\mathrm{TVK}} \frac{1}{BH} \sum_{b=1}^{B} \sum_{i=1}^{H} \left[1 - F_{\mathrm{TVK},b,i}\right]. \tag{9}$$

All operations remain within the same automatic-differentiation graph, allowing the feasibility regulariser to propagate through $\hat{\mathbf{x}}_1^{(s)} \rightarrow \mathbf{v}_\theta \rightarrow \theta$. The linear probability path, target velocity, and mean-squared error of the base Flow Matching objective remain unchanged. The additional term only adjusts the relative preference assigned to different trajectory regions. The exponential mapping bounds the pointwise regulariser and reduces the risk that a small number of extreme costs overwhelms trajectory-distribution learning early in training. From a transport-planning perspective, the term promotes trajectories with lower terrain and kinematic operating costs; it does not define demonstrated trajectories or any score threshold as an absolute safety label. λ_{TVK} controls the balance between distribution learning and feasibility preference. Any additional smoothness regulariser is retained as a separate loss term rather than being incorporated into F_{TVK}, so that the contribution of each constraint remains interpretable.

Algorithm 1. Training VTF-Flow with unified terrain–vehicle kinematic feasibility regularisation

Require: Training scene $\mathcal{S}_t = \{\mathbf{g}_t, \mathbf{M}_t\}$, supervised trajectory $\mathbf{x}_1$, conditional velocity field $\mathbf{v}_\theta$, and feasibility weight λ_{TVK}
Ensure: Updated model parameters θ
For each training batch **do**
Encode the local goal and terrain BEV to obtain $\mathbf{c}_t$
Sample $\mathbf{x}_0 \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ and $s \sim \mathcal{U}(0,1)$
Construct $\mathbf{x}_s = (1-s)\mathbf{x}_0 + s\mathbf{x}_1$ and $\mathbf{u}_s = \mathbf{x}_1 - \mathbf{x}_0$
Compute $\mathbf{v}_s = \mathbf{v}_\theta(\mathbf{x}_s, s, \mathbf{c}_t)$ and $\mathcal{L}_{\mathrm{FM}} = \| \mathbf{v}_s - \mathbf{u}_s \|_2^2$
Recover $\hat{\mathbf{x}}_1^{(s)} = \mathbf{x}_s + (1-s)\mathbf{v}_s$
Estimate speed, reliable heading, curvature, and lateral acceleration
Query F_{TVK} and compute $\mathcal{L}_{\mathrm{TVK}} = \mathrm{mean}(1 - F_{\mathrm{TVK}})$
Compute $\mathcal{L} = \mathcal{L}_{\mathrm{FM}} + \lambda_{\mathrm{TVK}} \mathcal{L}_{\mathrm{TVK}}$
Update θ by gradient descent
End for

Algorithm 1 places transport learning and unified terrain–vehicle kinematic preference within the same parameter update. The feasibility branch modifies only the training objective, not the probability path or target velocity. This separation permits the contributions of the generative backbone, terrain conditioning, and kinematic regularisation to be identified through controlled ablation.

3.5 Feasibility-potential guidance during flow evolution

Training-time soft regularisation cannot ensure that every stochastic sample has a low operating cost. During inference, VTF-Flow therefore introduces the unified trajectory potential J_{TVK} as an online correction to the generative dynamics. For each scene, the planner first samples K Gaussian initial trajectory states in parallel. At every numerical integration step, the network predicts the generative velocity from the current state, flow time, and scene condition, and then recovers the corresponding clean trajectory estimate. Candidate speed, reliable heading, curvature, and lateral acceleration are subsequently computed, C_{TVK} is queried along the trajectory, and the gradient of the trajectory potential is evaluated. This gradient is propagated through the complete computational chain $\mathbf{x}_s \rightarrow \mathbf{v}_\theta \rightarrow \hat{\mathbf{x}}_1^{(s)} \rightarrow J_{\mathrm{TVK}}$ to the current flow state. It therefore modifies the direction of the next generative update rather than post-processing the completed trajectory.

Raw pointwise gradients may contain high-frequency variation caused by local changes in the grid field. To coordinate corrections across neighbouring waypoints, we define a trajectory-coordination operator $\mathcal{A}(\cdot)$ that sequentially applies smoothing along the trajectory dimension, per-candidate scale normalisation, low-gradient suppression, and trajectory-level norm clipping. The final implementation additionally uses the linear projection $\mathcal{P}_g$

Just Accepted

that sets the terminal-waypoint correction to zero. This projection preserves the endpoint supplied by the goal-conditioned generative prior without using any intermediate GT waypoint. The guided continuous dynamics are

$$\frac{dx}{ds} = v_\theta(x_s, s, c_t) - \eta(s)\mathcal{P}_g\left[\mathcal{A}\left(\nabla_{x_s} J_{\mathrm{TVK}}\left(\hat{x}_1^{(s)}\right)\right)\right],$$
$$J_{\mathrm{TVK}} = \frac{1}{H}\sum_{i=1}^{H} C_{\mathrm{TVK},i}. \tag{10}$$

The first term in Eq. (10) supplies the data-driven generation direction learned from demonstrations. The quantity $-\mathcal{P}_g[\mathcal{A}(\nabla J_{\mathrm{TVK}})]$ can be interpreted as a generalised descent force induced by the feasibility potential in trajectory space, while $\eta(s)$ controls its strength over flow time. Both terms act on the same flow state within each Euler update. The potential-field correction therefore operates throughout the evolution from a random trajectory state to a complete candidate, rather than deleting, re-ranking, or independently optimising high-cost samples after generation. The time-dependent guidance coefficient $\eta(s)$ controls the strength of the correction over flow time. Earlier stages primarily preserve the global trajectory structure learned by the model, whereas later stages apply stronger local feasibility adjustment as the trajectory becomes better defined.

Algorithm 2. VTF-Flow evolution guided by unified terrain–vehicle kinematic feasibility

1. **Require:** Scene condition $\mathbf{c}_t$, terrain BEV $\mathbf{M}_t$, trained velocity field $\mathbf{v}_\theta$, number of candidates K, Euler steps N, and guidance schedule $\eta(s)$
2. **Ensure:** Candidate set $\{\boldsymbol{\tau}_t^{(k)}\}_{k=1}^{K}$
3. Sample $\mathbf{x}_0^{(k)} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ in parallel for $k = 1, \ldots, K$
4. **For** $n = 0, \ldots, N-1$ **do**
5. Set $s_n = n/N$ and compute $\mathbf{v}_n = \mathbf{v}_\theta(\mathbf{x}_n, s_n, \mathbf{c}_t)$
6. Recover $\hat{\mathbf{x}}_1^{(s_n)} = \mathbf{x}_n + (1 - s_n)\mathbf{v}_n$
7. Estimate speed, reliable heading, curvature, and lateral acceleration
8. Compute $J_{\mathrm{TVK}} = H^{-1}\sum_i C_{\mathrm{TVK},i}$
9. Evaluate $\mathbf{G}_n = \nabla_{\mathbf{x}_n} J_{\mathrm{TVK}}$
10. Coordinate and project the gradient as $\bar{\mathbf{G}}_n = \mathcal{P}_g[\mathcal{A}(\mathbf{G}_n)]$
11. Update $\mathbf{x}_{n+1} = \mathbf{x}_n + N^{-1}[\mathbf{v}_n - \eta(s_n)\bar{\mathbf{G}}_n]$
12. **End for**
13. **return** $\{\mathbf{x}_N^{(k)}\}_{k=1}^{K}$

Algorithm 2 shows that feasibility guidance is a component of every numerical integration step rather than an optimisation procedure independent of the generator. The process retains candidate differences arising from distinct random initial states while using controlled gradient updates to reduce the unified terrain–vehicle kinematic cost. For a fixed trained network, setting $\eta(s) = 0$ removes the inference-time correction exactly. Recovering the complete unguided Flow Matching baseline additionally requires $\lambda_{\mathrm{TVK}} = 0$ during training; the two switches therefore separate training regularisation from in-flow guidance in component-wise ablations.

VTF-Flow outputs a candidate trajectory tensor of shape $[B, K, H, 3]$. In this work, the phrase "evolution towards the feasible region" means that the trajectory is updated during generation towards lower J_{TVK} and higher relative continuous feasibility. It does not imply entry into a predefined hard-safe set. Likewise, "structure preserving" denotes trajectory-wise gradient smoothing, scale and norm control, low-gradient suppression, and endpoint-preserving projection to reduce local oscillation, abnormal displacement, and goal drift, whereas "kinematic feasibility" denotes continuous soft constraints on curvature and lateral-acceleration proxies. Neither constitutes a proof of full vehicle-dynamic executability.

4 Experimental setup

4.1 Evaluation objectives

The experiments were designed around the three connected claims of VTF-Flow. First, the continuous terrain-feasibility potential should distinguish recorded vehicle motion from controlled terrain-degrading perturbations and provide numerically stable gradients. Second, the vehicle-state and trajectory- kinematic terms should extend a static terrain representation without allowing one noisy component to dominate the objective. Third, embedding the resulting potential gradient in the Flow dynamics should improve the joint quality of the generated candidate set without replacing the data-driven trajectory prior. Accordingly, the evaluation addressed the following questions:

1. Is the terrain-feasibility potential discriminative, differentiable, and directionally meaningful under controlled trajectory perturbations?
2. Under a matched sequence-level protocol, how does VTF-Flow compare with motion extrapolation, terrain-grid search, deterministic learning, and unguided Flow Matching in terms of behavioural fidelity, candidate coverage, terrain exposure, and kinematic quality?
3. Does the final VTF-Flow configuration provide a directionally consistent improvement over its directly controlled Flow Matching baseline across held-out driving sequences?

The field diagnostics and the final planner benchmark are reported separately. Longer cached trajectories used in early field diagnostics were not included in the final method ranking, which used a common $H = 10$, 5 s protocol.

4.2 Dataset, coordinate alignment, and GT trajectory construction

Experiments used the RELLIS-3D dataset (Jiang et al., 2020). The locally available release contains five off-road driving sequences (00000–00004), with camera images, Ouster LiDAR point clouds, point-wise semantic labels, timestamped vehicle poses, and sensor calibration. LiDAR geometry, semantic labels, poses, and calibration were used to construct the planner input and recorded future trajectories. Camera images were used only for qualitative cross-view visualisation and were not provided to the planner.

The ground-truth future trajectory (GT) denotes the motion executed by the data-collection vehicle. It is neither a road centreline nor a manually labelled optimal path. Let

Just Accepted

$\mathbf{T}_{w\leftarrow e}(t)$ be the current ego pose in the world frame. A future pose was mapped into the current planning- ego frame as

$$\mathbf{T}_{\text{local}}(t+i) = \mathbf{T}_{w\leftarrow e}(t)^{-1}\mathbf{T}_{w\leftarrow e}(t+i). \quad (11)$$

The translation component provided (x_i, y_i, z_i), and the current pose was explicitly mapped to $(0,0,0)$. Linear interpolation between adjacent poses was used when a target sampling time did not coincide with a recorded timestamp. Each sample contained $H = 10$ future points at $\Delta t = 0.5$ s over a 5 s horizon. Construction rejected non-finite coordinates, non-monotone timestamps, inconsistent origins, and trajectories exceeding the configured $30\,\text{m}\,\text{s}^{-1}$ sanity limit. LiDAR observations and poses were placed in the same planning-ego frame through the verified Ouster-to-ego calibration.

GT served three roles: the supervised trajectory target, the behavioural reference for displacement errors, and the source of the common local goal, defined as the recorded endpoint at 5 s. It was not treated as the only feasible future or as an absolute safety optimum. Consequently, displacement errors quantify agreement with recorded behaviour, whereas the terrain–vehicle kinematic metrics quantify relative operating cost under the proposed representation. Neither metric family alone establishes safety (Bhavsar et al., 2025; Kurse et al., 2025).

Accordingly, the recorded GT supplies supervised behaviour and its endpoint supplies the goal condition, but neither defines a safety-optimal reference. Agreement with GT and low TVK potential are evaluated separately and are interpreted as complementary rather than interchangeable criteria.

4.3 Planner input and terrain representation

The effective planning condition comprised the local metric goal $\mathbf{g}_t \in \mathbb{R}^3$ and a three-channel terrain bird's-eye-view (BEV) tensor $\mathbf{M}_t = \{T, O, \bar{Z}\}$. The BEV covered 0–24 m in the ego-forward direction and $-12 - 12$ m laterally at 0.25 m resolution. Its channels were traversable-point fraction, obstacle-point density, and normalised mean elevation. Point-wise semantic labels were aggregated during offline BEV construction; no separate semantic image was passed to the neural planner. The software interface retained six ego-history states sampled at 0.1 s, but the frozen model did not use measured history as an effective condition. The scene encoder produced a 192-dimensional terrain feature, which was concatenated with the three-dimensional goal to form a 195-dimensional condition vector. Deterministic regression and both Flow models used matched scene-encoder capacity.

The planner-derived static potential combined occupancy, non-traversability, slope, roughness, and obstacle-proximity terms. Vehicle conditioning increased the slope, roughness, and proximity contributions continuously with candidate-induced speed and used a reliability-weighted heading look-ahead. Curvature and lateral-acceleration soft penalties completed the unified TVK potential. The nominal limits were $0.35\,\text{m}^{-1}$ and $2.5\,\text{m}\,\text{s}^{-2}$, with weights 0.25 for both terms. These settings are planning hyperparameters, not calibrated limits for a particular vehicle, because the dataset does not provide reliable tyre–terrain friction, steering, mass, suspension, or roll-stability parameters.

4.4 Sequence-level evaluation protocol

Random frame splitting was avoided because adjacent frames within a driving sequence are temporally correlated. Sequence 00003 was fixed as the development validation sequence and was used to select the final guidance settings before outer evaluation. Sequences 00000, 00001, and 00002 were then held out in turn. For each outer fold, the remaining sequences, including sequence 00004 but excluding sequence 00003, formed the training set. The resulting folds were

- test 00000: train 00001, 00002, and 00004 (7,175 training, 2,034 validation, and 1,797 test scenes);
- test 00001: train 00000, 00002, and 00004 (7,253 training, 2,034 validation, and 1,719 test scenes);
- test 00002: train 00000, 00001, and 00004 (5,425 training, 2,034 validation, and 3,547 test scenes).

The union of the outer test folds contained 7,063 planning scenes. This count describes repeated frame-level observations; the independent unit for cross-sequence interpretation was the held-out sequence ($n = 3$). No outer test sequence was used to select the final guidance coefficient, schedule, smoothing kernel, or endpoint projection.

To make the separation among training, development validation, outer-holdout testing, and the independent diagnostic explicit, Table 2 summarises the role of each RELLIS-3D sequence. Sequence 00004 contributed to training and the independent diagnostic but was not used for final model or guidance selection.

where H is the number of discrete future points within the planning horizon. All trajectories are expressed in the current planning-ego coordinate frame. The current ego position is mapped to the origin, with the x-axis pointing forward, the y-axis pointing to the left of the vehicle, and the z-axis pointing upward. Although the software interface retains an ego-history input, the frozen implementation evaluated in this study does not use measured ego history as an effective condition. The planning input is therefore restricted to the local goal and terrain representation.

Just Accepted

Table 2. Roles of RELLIS-3D sequences in training, development validation, outer-holdout testing, and independent diagnostic evaluation.

Sequence	Role in outer evaluation	Use for model or guidance selection	Use in independent diagnostic
00000	Outer test in fold 1; training in folds 2 and 3	No	No
00001	Outer test in fold 2; training in folds 1 and 3	No	No
00002	Outer test in fold 3; training in folds 1 and 2	No	No
00003	Development validation in every fold	Yes; guidance and q80 envelope frozen here	No
00004	Training in every outer fold	No; not used for model or guidance selection	Yes; 1,909-trajectory diagnostic

4.5 Benchmark methods

Five planners were compared under the same sequence folds, horizon, local goal, terrain representation, candidate evaluation code, and TVK parameters:

- **Constant Velocity (CV)** estimated recent ego velocity and extrapolated one trajectory ($K = 1$).
- **A* terrain planner** searched the local cost grid and resampled the path to $H = 10$ points ($K = 1$). Occupied cells were forbidden. Goals outside the forward-only BEV were retained as explicit failures rather than removed: 200 of the 7,063 test scenes used a stationary failure fallback, and 225 additional scenes used the recorded soft-cost fallback for a forbidden goal or disconnected grid.
- **Deterministic regression** used the matched scene encoder and an MLP decoder to produce one trajectory ($K = 1$).
- **Flow Matching** used the same goal-anchored conditional Flow backbone as VTF-Flow, but omitted TVK regularisation during training and feasibility guidance during sampling. It generated $K = 8$ candidates.
- **VTF-Flow** used the complete bounded TVK training regulariser and the final endpoint-preserving in-flow TVK guidance, generating $K = 8$ candidates.

4.6 Training and sampling settings

All learning-based models were trained independently for each outer fold and random seed. Seeds 0, 1, and 2 were treated as technical model-training replicates. The regression model used Smooth L1 loss with learning rate 10^{-3}. Both Flow models used the unchanged Flow Matching velocity loss with learning rate 5×10^{-4}. All three learning-based planners were trained for 40 epochs with batch size 128, weight decay 10^{-4}, and gradient-norm clipping at 1.0. The conditional Flow network used three hidden layers of width 192. VTF-Flow used $\lambda_{\mathrm{TVK}} = 1$ and no additional smoothness loss.

Both Flow planners used $K = 8$ initial Gaussian states and 16 fixed-step Euler updates. The final VTF-Flow guidance was selected on sequence 00003 and frozen before outer testing: $\eta = 0.075$, a late-strong schedule with exponent 1.0, the three-point kernel [0.25,0.50,0.25], per-candidate RMS gradient normalisation, low-gradient suppression, trajectory-level norm clipping, and a terminal projection that sets the feasibility correction at the final waypoint to zero. Flow Matching and VTF-Flow used identical initial Gaussian candidates within each paired evaluation run. All experiments were implemented in Python 3.9.23 and PyTorch 2.8.0 with CUDA 12.9, and were executed on an NVIDIA GeForce RTX 5060 GPU with 8 GB of memory. These identifiers describe the local experimental environment rather than a hardware-normalised latency protocol; latency was not used as a primary comparison criterion.

4.7 Evaluation metrics and statistical analysis

Let $\hat{\mathbf{p}}_{b,k,i}$ denote waypoint i of candidate k in scene b, and let $\mathbf{p}_{b,i}$ denote the corresponding GT waypoint. The point-wise displacement error is

$$e_{b,k,i} = \|\hat{\mathbf{p}}_{b,k,i} - \mathbf{p}_{b,i}\|_2. \tag{12}$$

Behavioural fidelity and candidate coverage were measured by

$$\mathrm{ADE\text{-}0} = \frac{1}{BH}\sum_{b=1}^{B}\sum_{i=1}^{H} e_{b,1,i}, \tag{13}$$

$$\mathrm{minADE@}K = \frac{1}{B}\sum_{b=1}^{B}\min_{1\le k\le K}\left(\frac{1}{H}\sum_{i=1}^{H} e_{b,k,i}\right), \tag{14}$$

$$\mathrm{minFDE@}K = \frac{1}{B}\sum_{b=1}^{B}\min_{1\le k\le K} e_{b,k,H}. \tag{15}$$

ADE-0 evaluates the first non-oracle candidate, whereas minADE@K reports retrospective best-of-K coverage of the recorded future. Because the common goal was the GT endpoint at 5 s, minFDE@K primarily diagnoses endpoint anchoring and was not used as primary evidence of method superiority.

Candidate diversity was the mean pairwise waypoint distance over the complete horizon,

$$\mathrm{Div} = \frac{1}{B\binom{K}{2}H}\sum_{b=1}^{B}\sum_{k<\ell}\sum_{i=1}^{H}\|\hat{\mathbf{p}}_{b,k,i} - \hat{\mathbf{p}}_{b,\ell,i}\|_2, \tag{16}$$

and was reported as not applicable for $K = 1$. Smoothness was measured by the mean norm of the second finite difference,

$$S(\hat{\boldsymbol{\tau}}) = \frac{1}{H-2}\sum_{i=2}^{H-1}\|\hat{\mathbf{p}}_{i+1} - 2\hat{\mathbf{p}}_i + \hat{\mathbf{p}}_{i-1}\|_2. \tag{17}$$

Unified feasibility was evaluated with the same post hoc objective for every method,

$$J_{\mathrm{TVK}}(\boldsymbol{\tau}) = \frac{1}{H}\sum_{i=1}^{H} C_{\mathrm{TVK},i}. \tag{18}$$

Terrain, occupancy, slope, curvature, and lateral-acceleration violation rates were the proportions of evaluated waypoints exceeding their predefined limits. For a generic quantity q_i,

$$\mathrm{VR}_q(\boldsymbol{\tau}) = \frac{1}{H}\sum_{i=1}^{H}\mathbb{I}\,[q_i > q_{\mathrm{lim}}]. \tag{19}$$

The value J_{TVK} is a relative planning potential, not a collision probability or binary safety criterion.

Just Accepted

To evaluate candidate usability without relying on the weighted TVK objective, we additionally froze a component-wise q80 demonstration envelope on validation GT. A candidate was counted as compliant only if it reached the goal within 0.25 m and simultaneously satisfied the validation-derived limits for occupancy, non-traversability, slope, roughness, clearance, curvature, and lateral acceleration. The compliant-candidate rate is the fraction of the K candidates inside this envelope. Goal-conditioned candidate coverage at K, denoted GCCR@K, is the fraction of scenes containing at least one compliant candidate. The envelope is a relative validation reference and not a calibrated safety boundary.

For each outer fold, the q80 envelope was fitted exclusively to GT trajectories from development sequence 00003 and frozen before outer testing. The upper 80th percentile was used for occupancy exposure, non-traversable exposure, slope exposure, mean roughness, curvature violation, and lateral-acceleration violation. The lower 20th percentile was used for the fifth-percentile clearance metric, while goal error used a fixed tolerance of 0.25 m. A candidate is compliant only when it satisfies all component limits simultaneously. The q80 compliant-candidate rate is the fraction of candidates satisfying this conjunction, whereas GCCR@8 is the fraction of scenes containing at least one compliant candidate among eight. These limits describe a validation-demonstration envelope and are not calibrated safety thresholds.

The reported metrics answer complementary questions. ADE-0 evaluates fixed candidate 0 and provides the closest open-loop proxy to a non-oracle single output. minADE@8 retrospectively selects the candidate closest to the recorded GT and therefore measures candidate-set coverage rather than deployable online selection. J_{TVK} summarizes the weighted terrain-vehicle kinematic potential under a common post hoc evaluator. Component-level violation rates show which terrain or kinematic factor contributes to the aggregate score and prevent a lower J_TVK from masking deterioration in a single component. Displacement and TVK metrics are therefore interpreted jointly.

Three selection regimes are distinguished. ADE-0 and the qualitative figures use fixed candidate 0 and require no GT-based selection. minADE@8 uses a retrospective GT oracle only to measure candidate-set coverage. A deployable selector could rank candidates using the predicted J_{TVK} together with goal and controller constraints, but this online rule and its closed-loop consequences are not evaluated here. We therefore do not interpret minADE@8 as deployable single-trajectory performance and identify online candidate selection and trajectory tracking as future work.

Metrics were first averaged across frames within each test sequence. For each learning-based method, the three training seeds were then averaged within the sequence. The paper reports the unweighted macro mean and sample standard deviation across the three held-out sequences. CV and A* were evaluated once per fold. The strongest controlled contrast, VTF-Flow versus Flow Matching, also reports the direction of the paired sequence-level effect and the number of improving seed–sequence technical pairs. No frame-wise significance test was used, because adjacent frames are not independent. With only three independent held-out sequences, the analysis is descriptive and does not claim external- dataset, cross-season, or cross-vehicle generalisation.

5 Results and analysis

This section evaluates VTF-Flow at three levels: the validity of the continuous feasibility potential, quantitative performance under the unified planning benchmark, and trajectory responses in real off-road scenes. We first assess whether the potential discriminates controlled trajectory perturbations and provides stable gradients and meaningful local descent directions. We then compare the planning methods under a common prediction horizon, goal condition, and post hoc evaluator. Finally, terrain-attribute decompositions and aligned camera–BEV views are used to examine the scene dependence of the resulting trajectory corrections.

5.1 Preliminary validation of the continuous feasibility potential

Before using the continuous feasibility potential for training regularisation and generation-time guidance, we tested whether it could distinguish recorded trajectories from controlled degradations and whether its numerical gradients were sufficiently stable and directionally meaningful for local correction. This independent diagnostic used 1,909 recorded GT trajectories from sequence 00004. A longer horizon of $H = 30$ was used to increase terrain exposure; these experiments were not used for model selection, hyperparameter tuning, or method ranking in the subsequent $H = 10$ outer-holdout benchmark. The score F is only a relative continuous feasibility score; in particular, $F < 0.5$ is not a hazard threshold.

Figure 2 shows that the component fields provide complementary spatial structure. The occupancy potential responds locally to obstacle returns, whereas the obstacle-proximity potential propagates their influence into the surrounding neighbourhood. The non-traversability potential represents insufficient traversable support, while the slope and roughness potentials capture local surface geometry. Their fusion does not collapse to a binary mask: C_T preserves low-potential corridors while providing continuous lateral and longitudinal gradients. This single scene illustrates how the component fields form the potential, but does not by itself establish population-level discrimination. Population-level discrimination was therefore evaluated using 1,909 paired perturbation tests, with the results summarised in Table 3. Numerical gradient behaviour was examined separately in the subsequent diagnostic.

Just Accepted

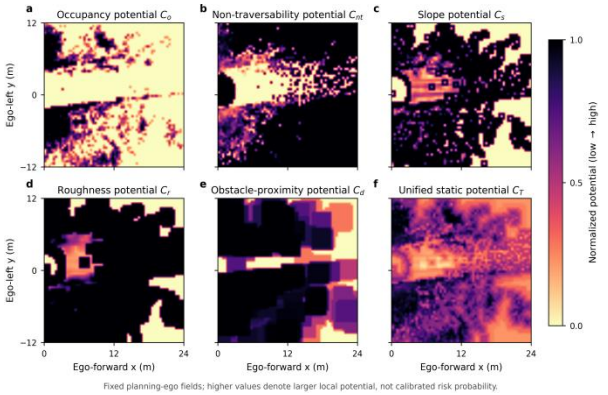


Fig. 2. Decomposition of the planner-used static terrain-feasibility potential. Panels (a)–(e) show the occupancy potential C_o, non-traversability potential C_{nt}, slope potential C_s, roughness potential C_r, and obstacle-proximity potential C_d, respectively; panel (f) shows the static terrain-feasibility potential C_T obtained with the frozen fusion weights. All panels correspond to RELLIS-3D sequence 00004, frame 001124, and use the same planning-ego extent, a common $[0,1]$ normalised scale, and bilinear rendering consistent with continuous field queries. Larger values denote higher local potential, not calibrated safety probabilities. Only terrain components that define fixed two-dimensional spatial fields are shown. Speed, heading reliability, curvature, and lateral acceleration depend on a specific candidate and therefore cannot be interpreted as pre-existing two-dimensional heatmaps.

Table 3. Paired feasibility comparison between recorded GT and controlled trajectory perturbations on sequence 00004 (n=1,909). A positive paired difference indicates higher relative feasibility for the recorded trajectory. The quantity d_z is the paired standardised effect size.

Perturbation type	Mean F_GT	Mean F_pert	Paired difference (95% CI)	P(F_GT > F_pert)	d_z
Smooth spatial perturbation	0.5665	0.5557	0.0108 [0.0095, 0.0121]	69.57%	0.377
Local rotation	0.5665	0.5653	0.0011 [0.0008, 0.0014]	57.62%	0.174
Controlled lateral offset	0.5665	0.5570	0.0095 [0.0074, 0.0115]	62.28%	0.209
Obstacle-directed perturbation	0.5665	0.5531	0.0134 [0.0126, 0.0143]	77.58%	0.701

Across the comparisons reported in Table 3, all four perturbation classes reduced mean feasibility relative to their paired GT trajectories, and none of the confidence intervals crossed zero. The obstacle-directed perturbation was most discriminative, with a paired win rate of 77.58% and an effect size of 0.701. By contrast, local rotations produced a win rate of 57.62%, indicating limited sensitivity to small angular changes on locally homogeneous terrain. The diagnostic therefore supports measurable discrimination of terrain-relevant degradation, but not uniformly strong classification of every geometric perturbation.

Gradient evaluation succeeded for all 512 valid diagnostic samples. Both the zero-gradient ratio and extreme-gradient ratio were zero. Near occupied cells, 93.44–96.49% of negative gradients pointed away from occupancy, and a 1 cm numerical perturbation of the trajectory yielded a gradient direction cosine similarity of 0.9963–0.9976. The mean contributions of the semantic, slope, roughness, occupancy, clearance, unknown-region, and speed-conditioned terms were 0.1565, 0.0767, 0.0137, 0.0299, 0.1777, 0.0950, and 0.0733, respectively. Thus, neither slope noise nor roughness noise dominated the total potential.

Finally, five feasibility-gradient descent steps applied directly to perturbed trajectories increased mean F from 0.5581 to 0.5926, improving 99.80% of samples. Occupancy exposure decreased from 0.924% to 0.417%, the mean trajectory displacement was 0.221 m, and the maximum displacement was bounded at 0.5 m. The smoothness cost nevertheless increased to 1.42 × its initial value. Extending the procedure to 10 and 20 steps continued to reduce some terrain potentials but increased the smoothness cost to 2.04 × and 3.10 × its initial value, respectively. These results empirically establish numerical differentiability, meaningful local descent, and the utility of short-range guidance. They also motivate trajectory-wise gradient smoothing, low-gradient suppression, norm clipping, and progressive coupling to the Flow dynamics instead of treating the potential gradient as an indefinitely iterated standalone trajectory optimiser.

5.2 Unified sequence-holdout benchmark

To distinguish fixed-candidate performance, candidate-set coverage, and relative feasibility, Table 4 reports ADE-0, minADE@K, and TVK and component-level metrics side by side under a common evaluation protocol. Candidate diversity is not applicable when $K = 1$. The zero curvature and zero second finite difference of Constant Velocity arise from its straight-line constant-speed construction and do not imply superior overall planning quality.

Just Accepted

Table 4. Unified outer-holdout benchmark. Values are the unweighted mean ± sample standard deviation across held-out sequences 00000, 00001, and 00002; each cell reports the mean and standard deviation as mean ± SD on a single line. For readability, the ten metrics are arranged in two vertically stacked blocks that repeat the method names. Constant Velocity, A*, and deterministic regression generate one trajectory, whereas both Flow-based methods generate eight candidates. Downward and upward arrows indicate whether lower or higher values are preferred. Boldface marks the best mean in each column; the formal zeros for Constant Velocity are retained but interpreted in the text.

Method	ADE-0 (m) ↓	minADE@K (m) ↓	Diversity (m)	TVK potential ↓	Terrain VR ↓
Constant Velocity	0.5783 ± 0.0213	0.5783 ± 0.0213	–	0.7060 ± 0.1322	0.5832 ± 0.0340
A* terrain planner	1.6191 ± 0.0186	1.6191 ± 0.0186	–	0.9283 ± 0.1566	**0.5505 ± 0.0313**
Deterministic regression	**0.1618 ± 0.0068**	0.1618 ± 0.0068	–	0.6833 ± 0.1217	0.5704 ± 0.0295
Flow Matching	0.1684 ± 0.0049	0.1402 ± 0.0095	0.0524 ± 0.0077	0.6856 ± 0.1210	0.5707 ± 0.0275
VTF-Flow	0.1662 ± 0.0039	**0.1387 ± 0.0099**	0.0506 ± 0.0093	**0.6788 ± 0.1213**	0.5673 ± 0.0281

Method	Occupancy exposure ↓	Curvature VR ↓	Smoothness (m) ↓	q80 compliant rate ↑	GCCR@K ↑
Constant Velocity	0.0236 ± 0.0129	–	–	0.0106 ± 0.0025	0.0106 ± 0.0025
A* terrain planner	**0.0051 ± 0.0037**	0.3393 ± 0.0126	0.6668 ± 0.0434	0.0104 ± 0.0054	0.0104 ± 0.0054
Deterministic regression	0.0135 ± 0.0083	0.0102 ± 0.0084	**0.0201 ± 0.0061**	0.3744 ± 0.1984	0.3744 ± 0.1984
Flow Matching	0.0140 ± 0.0085	0.0113 ± 0.0089	0.0245 ± 0.0026	0.3753 ± 0.1991	0.4040 ± 0.2020
VTF-Flow	0.0137 ± 0.0085	**0.0093 ± 0.0085**	0.0241 ± 0.0025	**0.3845 ± 0.1962**	**0.4105 ± 0.1993**

VTF-Flow provided the most balanced overall performance in the unified outer-holdout benchmark. With $K = 8$, it attained a minADE@8 of 0.1387 m, a unified TVK potential of 0.6788, a q80 compliant-candidate rate of 0.3845, and a GCCR@8 of 0.4105, the best values in Table 4 for all four metrics. Relative to unguided Flow Matching with the same generative backbone, VTF-Flow reduced ADE-0 and minADE@8 by 1.31% and 1.07%, respectively. It also reduced the unified TVK potential by 0.99%, terrain violation rate by 0.34 percentage points, occupancy exposure by 2.14% relative, curvature violation rate by 17.70%. The q80 compliant-candidate rate and GCCR@8 increased by 0.92 and 0.65 percentage points. These results indicate that bounded TVK training regularisation and endpoint-preserving in-flow TVK guidance reduced unified feasibility cost while preserving, and slightly improving, coverage of the recorded future under a common goal and matched initial random candidates.

Deterministic regression retained a small advantage in first-candidate behavioural fidelity: its ADE-0 was 0.0044 m below that of VTF-Flow. However, VTF-Flow's minADE@8 was 14.28% below the single-trajectory ADE of the deterministic model; its unified TVK potential was 0.66% lower, and its q80 compliant-candidate rate and GCCR@8 were higher by 1.01 and 3.61 percentage points, respectively. Because minADE@8 retrospectively selects the best candidate, it measures candidate-set coverage of the recorded future rather than online selection performance. VTF-Flow's curvature violation rate of 0.0093 was effectively equal to the regression model's 0.0092. Candidate diversity decreased modestly from 0.0524 m for unguided Flow Matching to 0.0506 m for VTF-Flow, a reduction of 3.44%. The guidance therefore did not collapse the candidate set, although improved low-potential, goal-consistent coverage was accompanied by limited contraction.

Relative to the conventional planners, the principal advantage of VTF-Flow was coordination among terrain, goal attainment, and trajectory kinematics, rather than the absolute minimisation of any single violation component. The A· terrain planner achieved the lowest terrain violation rate and occupancy exposure, but its ADE-0, unified TVK potential, curvature violation rate, and smoothness cost were 1.6191 m, 0.9283, 0.3393, and 0.6668, respectively. VTF-Flow reduced these four quantities by 89.74%, 26.88%, 97.26%, and 96.39%, while consistently producing goal-directed candidate sets. The 200 stationary failure fallbacks and 225 soft-cost fallbacks of A· were retained in the statistics. The result illustrates that discrete search on a forward local grid can emphasise obstacle avoidance but does not simultaneously preserve continuous motion coordination and consistency with recorded driving behaviour. Similarly, the zero curvature and zero second finite difference of Constant Velocity are mathematical consequences of linear extrapolation. Its ADE-0 of 0.5783 m and q80 compliant-candidate rate of 0.0106 show that formal geometric smoothness alone is not equivalent to scene-adaptive planning.

The improvements over Flow Matching were also directionally consistent across technical repetitions. Across the nine seed–held-out-sequence pairs, VTF-Flow improved the unified TVK potential, terrain violation rate, curvature violation rate, and q80 compliant-candidate rate in 9/9 pairs; occupancy exposure and GCCR@8 improved in 8/9 pairs. ADE-0, minADE@8, and smoothness improved in 7/9, 6/9, and 6/9 pairs, respectively. The most consistent evidence therefore concerns unified potential, terrain exposure, kinematic coordination, and candidate compliance, whereas the gains in displacement error are smaller.

5.3 Component ablation of the hierarchical feasibility potential

To identify the contributions of the hierarchical feasibility representation, we performed component ablations under the common $H = 10$, 5 s, and $K = 8$ protocol. The ablation used sequences 00000–00002 for training, sequence 00003 for

Just Accepted

validation, and sequence 00004 for testing, yielding 1,909 test scenes. Results were averaged over three random seeds. This fixed-split experiment probes internal component effects and does not replace the main method comparison across the three outer held-out sequences in Section 5.2.

To make the contribution of each component directly auditable, Table 5 separates the ablation configuration from the behavioural, feasibility, and component-level constraint results. The complete VTF-Flow configuration is placed in the final row of each block.

Table 5. Component ablation of the hierarchical VTF-Flow feasibility potential on sequence 00004. Values are averages over three random seeds.

Method	Vehicle conditioning	Kinematic terms	Training regularisation	Inference-time guidance
Flow Matching	No	No	No	No
VTF-Flow w/o kinematic potential	Yes	No	Yes (VT)	Yes (VT)
VTF-Flow w/o training regularisation†	Yes	Yes	No	Yes (TVK)
VTF-Flow (full)	Yes	Yes	Yes (TVK)	Yes (TVK)
Method	**ADE-0 ↓**	**minADE@8 ↓**	**Diversity**	**TVK potential ↓**
Flow Matching	0.1383	0.1070	0.0570	0.5435
VTF-Flow w/o kinematic potential	0.1435	0.1104	0.0626	0.5316
VTF-Flow w/o training regularisation†	0.1375	0.1069	0.0567	0.5403
VTF-Flow (full)	0.1388	0.1085	0.0571	0.5301

Method	Terrain VR ↓	Curvature VR ↓	Smoothness ↓
Flow Matching	0.5954	0.0178	0.0255
VTF-Flow w/o kinematic potential	0.5902	0.0200	0.0260
VTF-Flow w/o training regularisation†	0.5945	0.0169	0.0254
VTF-Flow (full)	0.5887	0.0124	0.0239

Note. VT denotes the vehicle-state-conditioned terrain potential; TVK denotes the unified terrain–vehicle kinematic potential, including the curvature and lateral-acceleration terms. The † variant used a separately optimised inference-guidance configuration and is reported as an accuracy–feasibility diagnostic rather than a strict one-factor ablation of training regularisation. Within this table only, VTF-Flow (full) denotes the complete configuration and is placed after all reduced variants.

Introducing only the vehicle-state-conditioned terrain potential reduced the TVK potential from 0.5435 to 0.5316 and the terrain violation rate from 0.5954 to 0.5902, but increased ADE-0, minADE@8, curvature violation, and smoothness cost. The terrain potential can therefore move candidates towards lower-cost regions, but cannot by itself constrain trajectory bending induced by local spatial gradients. After adding the curvature and lateral-acceleration terms, VTF-Flow (full) reduced curvature violation and smoothness cost by 37.91% and 8.14%, respectively, relative to the variant without kinematic potential. It also reduced ADE-0 and minADE@8 by 3.23% and 1.72%. Diversity decreased from 0.0626 to 0.0571, effectively matching the value of 0.0570 for unguided Flow Matching. These results indicate that the principal role of the kinematic potential is not to increase candidate dispersion, but to coordinate the terrain-descent direction with trajectory geometry.

Relative to unguided Flow Matching on this fixed test sequence, VTF-Flow (full) reduced the unified TVK potential by 2.46%, terrain violation rate by 0.68 percentage points, curvature violation rate by 30.54%, and smoothness cost by 6.42%, while limiting the relative changes in ADE-0 and minADE@8 to 0.36% and 1.39%. The diagnostic variant without training regularisation obtained lower ADE-0 and minADE@8, but its TVK potential, terrain violation rate, curvature violation rate, and smoothness cost were respectively 1.91%, 0.58 percentage points, 35.96%, and 6.69% higher than those of VTF-Flow (full). This comparison suggests that bounded TVK training regularisation changes the optimisation emphasis from fitting the recorded future alone to jointly accounting for terrain and kinematic quality. Because the diagnostic variant also used separately optimised inference-guidance parameters, however, the observed differences cannot be attributed exclusively to training regularisation. Strict separation of training-time and inference-time effects will require an otherwise identical guidance configuration.

5.4 Trajectory correction and cross-view interpretation in real terrain

This subsection presents three complementary forms of qualitative evidence. The first uses a predeclared improvement criterion to visualise the magnitude and direction of representative VTF-Flow corrections. The second is independent of method error or gain and uses the three outer held-out sequences to show how measured terrain attributes form the planner-used static potential. The third places synchronised camera images alongside the BEV potential to connect the real environment with the metric planning representation. Every figure uses fixed candidate 0 rather than a candidate selected retrospectively using GT.

Just Accepted

5.4.1 Representative trajectory correction.

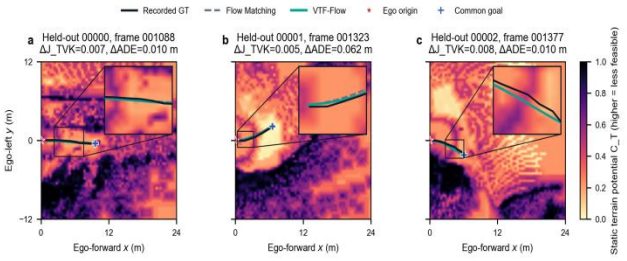


Fig. 3. Representative final VTF-Flow trajectory corrections on the three outer held-out sequences. The background is the static terrain-feasibility potential C_T actually derived from the planner's three-channel BEV. The recorded GT is black; unguided Flow Matching is grey and dashed; VTF-Flow is teal. Thin lines show all eight candidates, whereas thick lines identify fixed candidate 0; no candidate is selected retrospectively using GT. Each panel contains a metrically scaled inset. A common online rule locates the internal waypoint with the greatest separation among Flow candidate 0, VTF-Flow candidate 0, and recorded GT, and displays that waypoint with two neighbours on either side. Scenes were selected using a rule fixed before plotting: within each sequence, scenes showing simultaneous ADE-0 and TVK improvement were restricted to the upper quartile of recorded path length, after which the scene closest to the median relative TVK improvement was chosen. The figure therefore does not show the single largest-gain sample. The panels show sequences 00000, 00001, and 00002 at frames 001090, 001516, and 003776, respectively. Each BEV spans 0-24 m longitudinally and -12 to 12 m laterally in the planning-ego frame. Larger colour-bar values denote higher relative planning potential, not collision probability or an absolute safety level.

Across the three scenes in Fig. 3, the VTF-Flow and original Flow candidates remained largely coincident instead of being displaced substantially from the recorded driving distribution. Their TVK potentials decreased by 0.0073, 0.0049, and 0.0084, while ADE-0 decreased by 0.0104 m, 0.0623 m, and 0.0098 m, respectively. The visually modest differences are consistent with the small macro-averaged effects: $\eta = 0.075$, trajectory-wise gradient smoothing, RMS normalisation, low-gradient suppression, norm clipping, and terminal projection jointly constrain the correction magnitude. The figure therefore supports limited low-potential correction while preserving generative structure and the goal condition; it does not demonstrate large-scale replanning or a strict obstacle-avoidance guarantee.

Because the scene-selection rule for Fig. 3 explicitly required simultaneous ADE-0 and TVK improvement, the figure shows how correction occurs but cannot estimate the prevalence of improving scenes. The population-level cross-sequence effect must instead be judged from Table 4 and the paired results in Section 5.2. Figures 4– 9 use selection rules independent of method gain and provide more neutral real-data interpretation.

5.4.2 Terrain-attribute decomposition across held-out sequences.

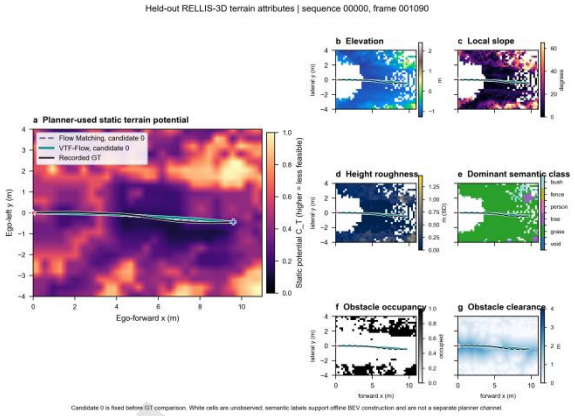


Fig. 4. Terrain-attribute decomposition for held-out sequence 00000, frame 001090. Panel (a) shows the static terrain-feasibility potential C_T queried by the final planner, overlaid with Flow Matching candidate 0, VTF-Flow candidate 0, and recorded GT. Panels (b)–(g) show elevation, local slope, roughness, dominant semantic class, occupancy, and obstacle clearance in the same planning-ego coordinate frame. White cells are unobserved. Candidate indices were fixed before inference and were not selected using GT. The displayed BEV spans 0-24 m longitudinally and -12 to 12 m laterally. Larger colour-bar values denote higher relative planning potential, not collision probability or an absolute safety level.

In Fig. 4, the low-potential corridor extends primarily in the forward direction and coincides with greater obstacle clearance and lower local occupancy. Both generated trajectories remain near the recorded trajectory, and VTF-Flow applies only a limited lateral correction. Online potential descent therefore does not draw the candidate towards an isolated low-potential region far from the behavioural prior.

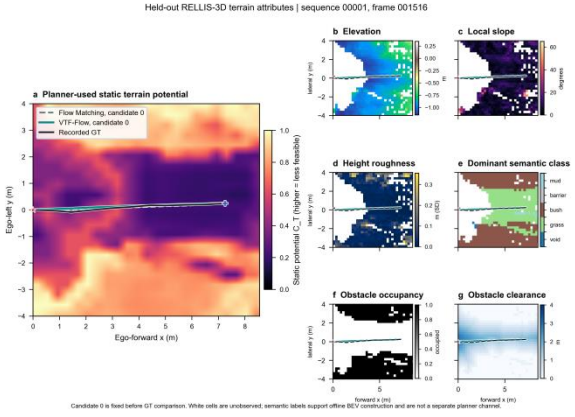


Fig. 5. Terrain-attribute decomposition for held-out sequence 00001, frame 001516. Panel definitions and trajectory-display rules are identical to Fig. 4. The scene contains a comparatively well-defined forward low-potential corridor bounded by higher occupancy, reduced clearance, and locally rough terrain. The displayed BEV spans 0-24 m longitudinally and -12 to 12 m laterally. Larger colour-bar values denote higher relative planning potential, not collision probability or an absolute safety level.

Compared with Fig. 4, the feasible corridor in Fig. 5 is more concentrated. Its low-potential band is jointly formed by occupancy, clearance, and local surface attributes rather than by the slope term alone. Fixed candidate 0 extends along this corridor, indicating that the fused potential preserves corridor geometry rather than merely seeking a local minimum of one component.

Just Accepted

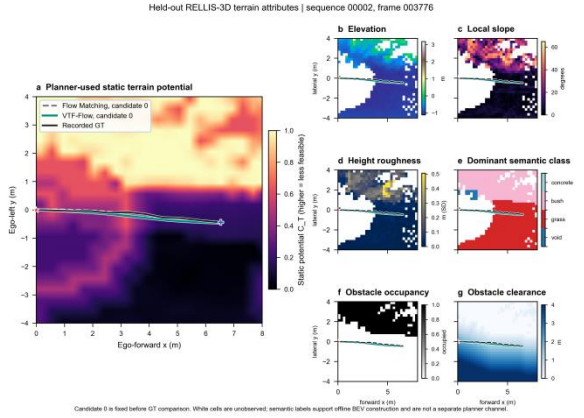


Fig. 6. Terrain-attribute decomposition for held-out sequence 00002, frame 003776. Panel definitions and trajectory-display rules are identical to Fig. 4. The static potential exhibits a pronounced lateral partition, and its low-potential region corresponds to greater obstacle clearance and lower occupancy. The displayed BEV spans 0-24 m longitudinally and -12 to 12 m laterally. Larger colour-bar values denote higher relative planning potential, not collision probability or an absolute safety level.

The terrain distribution in Fig. 6 differs from the first two examples, confirming that C_T varies with the measured geometry and semantic composition. VTF-Flow nevertheless maintains a continuous, goal-directed trajectory without introducing pronounced corners or endpoint drift in response to local high-potential regions. This observation is consistent with the lower curvature violation rate and Flow-like smoothness in Table 4.

Figures 4–6 are not intended to rank the methods again. Instead, they test whether the association among measured terrain attributes, the planner-used static potential, and the final candidate trajectory recurs across sequences. In all three scenes, low-potential corridors coincide with lower occupancy, greater obstacle clearance, and comparatively smooth local surfaces, while the locations and geometries of high-potential regions vary substantially. Thus, C_T is a scene-grounded representation derived from the measured terrain at each frame, not a fixed road template. Fixed candidate 0 from VTF-Flow largely follows these low-potential corridors while preserving the common 5 s goal and avoiding large displacements towards local high-potential regions. Semantic labels contribute only to the offline construction of T and O and are not provided as an independent online planner input channel.

The camera view in Fig. 7 supplies real-world context, including grass, wheel tracks, and vegetation boundaries, whereas the BEV provides the metric space queried by the planner. The forward open region occupied by the candidates is geometrically aligned with the low-potential band in the BEV. Because the camera image is not used by the network, this alignment serves only to interpret the planning result.

5.4.3 Camera–BEV cross-view interpretation.

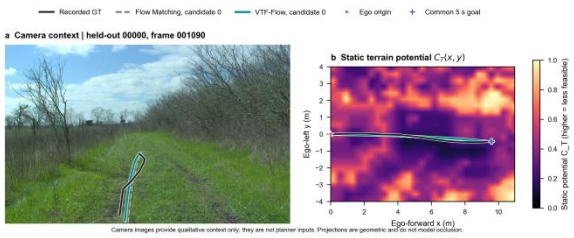


Fig. 7. Camera–BEV interpretation for held-out sequence 00000, frame 001090. Panel (a) projects recorded GT, Flow Matching candidate 0, and VTF-Flow candidate 0 into the synchronised camera image; panel (b) shows the corresponding trajectories on the static terrain-feasibility potential in the same planning-ego coordinate frame. The camera image supplies qualitative environmental context and is not a planner input. The geometric projection does not model occlusion explicitly. The BEV panel spans 0-24 m longitudinally and -12 to 12 m laterally. Candidate 0 was fixed before inference and was not selected with a GT-dependent oracle. Larger colour-bar values denote higher relative planning potential, not collision probability or an absolute safety level.

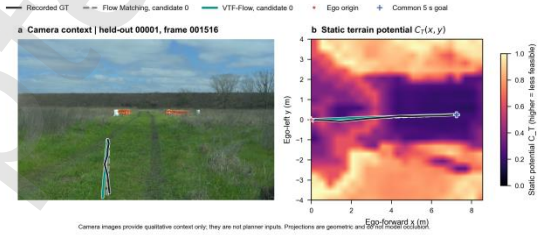


Fig. 8. Camera–BEV interpretation for held-out sequence 00001, frame 001516. Panel definitions are identical to Fig. 7. The grass wheel tracks in the camera view correspond geometrically to the low-potential corridor extending near $y \approx 0$ in the BEV. The BEV panel spans 0-24 m longitudinally and -12 to 12 m laterally. Candidate 0 was fixed before inference and was not selected with a GT-dependent oracle. Larger colour-bar values denote higher relative planning potential, not collision probability or an absolute safety level.

In Fig. 8, the boundaries of the forward corridor are more apparent than in Fig. 7. The fixed candidates from Flow Matching and VTF-Flow both remain within the visible corridor, while the VTF-Flow change remains a local coordination adjustment rather than a post-generation path replacement.

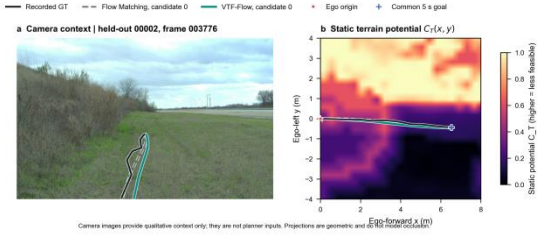


Fig. 9. Camera–BEV interpretation for held-out sequence 00002, frame 003776. Panel definitions are identical to Fig. 7. The trajectories extend through the open grass-covered region and remain close to the low-potential band in the BEV. The BEV panel spans 0-24 m longitudinally and -12 to 12 m laterally. Candidate 0 was fixed before inference and was not selected with a GT-dependent oracle. Larger colour-bar values denote higher relative planning potential, not collision probability or an absolute safety level.

Figure 9 shows a more open off-road region with weaker boundaries. In the absence of regularised lane markings, the BEV potential still provides a continuous relative preference without implying a unique correct route. The example

Just Accepted

therefore illustrates limited correction by a continuous terrain potential under weak road structure, not the detection of a deterministic road centreline from the camera image.

The scenes in Figs. 7–9 were preselected from the upper quartile of recorded path length in each held-out sequence according to geometric projection visibility, without using method error, TVK gain, or a GT-oracle candidate. Camera images provide qualitative environmental context only and are not planner inputs. Because the trajectory projection does not explicitly model occlusion, these figures do not constitute pixel-level collision detection or safety validation.

Together, the cross-view results show an interpretable geometric correspondence between low-potential corridors in the BEV and grass passages, wheel tracks, and local obstacle boundaries in the camera views. VTF-Flow generally applies small lateral adjustments within the visible corridor rather than large deviations from the data-driven prior. This observation is consistent with the modest but directionally coherent improvements in Fig. 3, Table 4, and Section 5.2.

6 Conclusions

This study addresses trajectory planning in unstructured off-road environments, where road boundaries are weakly defined, terrain conditions vary continuously, and multiple local motions may be plausible, by proposing VTF-Flow. Rather than treating terrain feasibility solely as a post-generation screening criterion, VTF-Flow constructs a continuous static terrain-feasibility potential from measured terrain attributes and augments it with candidate-induced speed, heading, curvature, and lateral acceleration to form a unified terrain–vehicle kinematic potential. Through bounded TVK training regularisation and endpoint-preserving in-flow TVK guidance, this potential acts directly on the trajectory evolution of the conditional Flow, allowing the data-driven multi-candidate trajectory prior to undergo continuous low-potential correction while preserving the goal condition and overall trajectory structure.

An independent diagnostic conducted before training the generative planner evaluated the basic suitability of the feasibility potential as a guidance signal. Across 1,909 paired tests on RELLIS-3D sequence 00004, all four controlled perturbation types reduced mean feasibility relative to the recorded GT, and the obstacle-directed perturbation achieved a paired win rate of 77.58%. Neither zero nor extreme gradients occurred among the valid samples, and negative gradients near occupied regions predominantly pointed away from occupied cells. Only five local gradient-descent steps increased the mean feasibility score from 0.5581 to 0.5926 and reduced occupancy exposure from 0.924% to 0.417%, whereas longer standalone descent progressively degraded trajectory smoothness. These results indicate that the potential can provide stable correction signals with meaningful local directions.

The unified comparison across three outer held-out sequences, comprising 7,063 planning scenes, further showed that VTF-Flow achieved a balanced trade-off among behavioural fidelity, candidate coverage, and terrain–kinematic quality. Among the five planning methods, VTF-Flow obtained the lowest minADE@8 and unified TVK potential, together with the highest q80 compliant-candidate rate and GCCR@8. Relative to unguided Flow Matching with the same backbone and matched initial random candidates, VTF-Flow reduced ADE-0, minADE@8, the unified TVK potential, and curvature violation rate by 1.31%, 1.07%, 0.99%, and 17.70%, respectively, while increasing the q80 compliant-candidate rate and GCCR@8 by 0.92 and 0.65 percentage points. Across the nine seed–held-out-sequence technical pairs, the unified TVK potential, terrain violation rate, curvature violation rate, and q80 compliant-candidate rate all improved in 9/9 pairs. Component ablation further showed that VTF-Flow (full) reduced curvature violation and smoothness cost by 37.91% and 8.14%, respectively, relative to the variant without the kinematic potential. This finding indicates that the curvature and lateral-acceleration terms primarily coordinate terrain-descent directions with trajectory geometry rather than simply increasing candidate dispersion.

RELLIS-3D does not provide reliable vehicle mass, steering-system, suspension, or tyre–ground friction parameters. The current curvature and lateral- acceleration constraints therefore remain kinematic approximations and cannot substitute for a complete analysis of vehicle dynamics and stability. The current experiments also rely primarily on open-loop local prediction and recorded-endpoint conditioning; online candidate selection, closed-loop control error accumulation, and real-vehicle execution have yet to be evaluated. Future work could extend the unified potential using calibrated vehicle parameters and tyre–ground conditions, incorporate time-varying terrain and observation uncertainty, and evaluate closed-loop receding-horizon planning and real-vehicle operation across routes, seasons, and vehicle platforms. VTF-Flow thus provides a reproducible technical route for directly embedding scene-specific terrain information and vehicle operating constraints into multimodal trajectory generation, establishing a basis for feasibility-aware generative planning in unstructured off-road transport environments.

Author contributions

Linheng Li: Conceptualization, Methodology, Validation, Visualization, Writing – original draft, Writing – review & editing, Resources. **Weijie Wang:** Formal analysis, Writing – original draft, Writing – review & editing. **Tiantong Nan:** Formal analysis, Methodology, Writing – original draft, Writing – review & editing. **Jing Gan:** Conceptualization, Writing – review & editing, Resources. **Xu Qu:** Conceptualization, Writing – review & editing, Resources. **Bin Ran:** Supervision, Project administration, Resources.

Acknowledgements

This work was supported by the National Natural Science Foundation of China (Grant No. 52572370), Natural Science

Just Accepted

Foundation of Jiangsu Province (Grant No. BK20250655), Natural Science Foundation of the Jiangsu Higher Education Institution (Grant No. 25KJB580012), and Sichuan Science and Technology Program (Grant No. 24NSFSC6306).

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Declaration of generative Al and AI-assisted technologies in the writing process (if applicable)

During the preparation of this work, the authors used ChatGPT (OpenAI) for language refinement. The authors take full responsibility for the content of the publication.

Replication and data sharing

provided after publication. The data used in this study are from the third-party RELLIS-3D (https://github.com/unmannedlab/RELLIS-3D) dataset and remain subject to the original provider's licence. Source code is available at https://doi.org/10.26599/ETSD.2026.9190096.

References

Bhavsar, Y. M., Zaveri, M. S., Raval, M. S., and Zaveri, S. B. 2025. Evaluating defensive driving behaviour based on safe distance between vehicles: A case study using computer vision on UAV videos at urban roundabout. Multimodal Transportation, 4(3), 100227.

Cai, X., Ancha, S., Sharma, L., Osteen, P. R., Bucher, B., Phillips, S., et al. 2024. EVORA: Deep evidential traversability learning for risk-aware off-road autonomy. IEEE Transactions on Robotics, 40, 3756–3777.

Carvalho, J., Le, A. T., Kicki, P., Koert, D., and Peters, J. 2024. Motion planning diffusion: Learning and adapting robot motion planning with diffusion models. arXiv preprint arXiv:2412.19948.

Chai, Y., Sapp, B., Bansal, M., and Anguelov, D. 2020. MultiPath: Multiple probabilistic anchor trajectory hypotheses for behavior prediction. In: Proceedings of the conference on robot learning, 86–99.

Chen, Z., Sun, C., Nie, S., Min, C., Ning, C., Li, H., et al. 2024. Semantic scene completion based 3D traversability estimation for off-road terrains. arXiv preprint arXiv:2412.08195.

Cheng Z, Zhu J, Feng Z, et al. 2025. Driving safety risk analysis and assessment in a mixed driving environment of connected and non-connected vehicles: A systematic survey. IEEE transactions on intelligent transportation systems, 2025, 26(5): 5747-5781.

Chi, C., Feng, S., Du, Y., Xu, Z., Cousineau, E., Burchfiel, B., et al. 2023. Diffusion policy: Visuomotor policy learning via action diffusion. In: Robotics: Science and systems XIX.

Fan, D. D., Otsu, K., Kubo, Y., Dixit, A., Burdick, J., and Agha-Mohammadi, A.-A. 2021. STEP: Stochastic traversability evaluation and planning for risk-aware off-road navigation. In: Proceedings of robotics: Science and systems.

Farahnakian, F., Nevalainen, P., Farahnakian, F., Vähämäki, T., and Heikkonen, J. 2025. Maritime vessel movement prediction: A temporal convolutional network model with optimal look-back window size determination. Multimodal Transportation, 4(1), 100191.

Frey, J., Patel, M., Atha, D., Nubert, J., Padgett, C., Hutter, M., et al. 2024. RoadRunner: Learning traversability estimation for autonomous off-road driving. IEEE Transactions on Field Robotics, 1, 192–212.

Guan, T., Kothandaraman, D., Chandra, R., Sathyamoorthy, A. J., Weerakoon, K., and Manocha, D. 2022. GA-Nav: Efficient terrain segmentation for robot navigation in unstructured outdoor environments. IEEE Robotics and Automation Letters, 7, 8138–8145.

Hart, P. E., Nilsson, N. J., and Raphael, B. 1968. A formal basis for the heuristic determination of minimum cost paths. IEEE Transactions on Systems Science and Cybernetics, 4, 100–107.

Janner, M., Du, Y., Tenenbaum, J., and Levine, S. 2022. Planning with diffusion for flexible behavior synthesis. In: Proceedings of the 39th international conference on machine learning, 9902–9915.

Jiang, C., Cornman, A., Park, C., Sapp, B., Zhou, Y., and Anguelov, D. 2023. MotionDiffuser: Controllable multi-agent motion prediction using diffusion. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, 9644–9653.

Jiang, J., Hu, Z., Xie, Z., Hao, C., Liu, H., Xu, W., et al. 2025. A risk-aware planning framework of UGVs in off-road environment. IEEE Transactions on Vehicular Technology, 74, 3870–3884.

Jiang, P., Osteen, P., Wigness, M., and Saripalli, S. 2020. RELLIS-3D dataset: Data, benchmarks and analysis. https://arxiv.org/abs/2011.12954.

Kalakrishnan, M., Chitta, S., Theodorou, E., Pastor, P., and Schaal, S. 2011. STOMP: Stochastic trajectory optimization for motion planning. In: 2011 IEEE international conference on robotics and automation, 4569–4574.

Khatib, O. 1986. Real-time obstacle avoidance for manipulators and mobile robots. The international journal of robotics research, 5, 90–98.

Kurse, T. K., Gebresenbet, G., Daba, G. F., and Tefera, N. T. 2025. Experimental determination of factors causing crashes involving automated vehicles. Multimodal Transportation, 4(1), 100186.

Liao, B., Chen, S., Yin, H., Jiang, B., Wang, C., Yan, S., et al. 2025. DiffusionDrive: Truncated diffusion model for end-to-end autonomous driving. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, 12037–12047.

Lipman, Y., Chen, R. T., Ben-Hamu, H., Nickel, M., and Le, M. 2022. Flow matching for generative modeling. arXiv preprint arXiv:2210.02747.

Liu, L., Jia, C., Yu, G., Song, Z., Li, J., Jia, F., et al. 2026. GuideFlow: Constraint-guided flow matching for planning in end-to-

Just Accepted

end autonomous driving. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, 3719–3728.

Liu, X., Gong, C., and Liu, Q. 2023. Flow straight and fast: Learning to generate and transfer data with rectified flow. In: International conference on learning representations.

Meng, X., Hatch, N., Lambert, A., Li, A., Wagener, N., Schmittle, M., et al. 2023. TerrainNet: Visual modeling of complex terrain for high-speed, off-road navigation. In: Proceedings of robotics: Science and systems.

Ndunguru, D. D., Xing, F., Oroni, C. Z., Ani, A., and Li, C. 2025. STG-LSTM: Spatial-temporal graph-based long short-term memory for vehicle trajectory prediction. Multimodal Transportation, 4(3), 100222.

Phan-Minh, T., Grigore, E. C., Boulton, F. A., Beijbom, O., and Wolff, E. M. 2020. CoverNet: Multimodal behavior prediction using trajectory sets. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, 14074–14083.

Ratliff, N., Zucker, M., Bagnell, J. A., and Srinivasa, S. 2009. CHOMP: Gradient optimization techniques for efficient motion planning. In: 2009 IEEE international conference on robotics and automation, 489–494.

Schmid, R., Atha, D., Schöller, F., Dey, S., Fakoorian, S., Otsu, K., et al. 2022. Self-supervised traversability prediction by learning to reconstruct safe terrain. In: 2022 IEEE/RSJ international conference on intelligent robots and systems, 12419–12425.

Schulman, J., Ho, J., Lee, A., Awwal, I., Bradlow, H., and Abbeel, P. 2013. Finding locally optimal, collision-free trajectories with sequential convex optimization. In: Proceedings of robotics: Science and systems.

Tan, T., Zheng, Y., Liang, R., Wang, Z., Zheng, K., Zheng, J., et al. 2025. Flow matching-based autonomous driving planning with advanced interactive behavior modeling. In: Advances in neural information processing systems.

Vidanapathirana, K., Knights, J., Hausler, S., Cox, M., Ramezani, M., Jooste, J., et al. 2025. WildScenes: A benchmark for 2D and 3D semantic segmentation in large-scale natural environments. The International Journal of Robotics Research, 44, 532–549.

Wang, B., Li, Z., Yu, R., Che, Y., Zhang, X., Liu, M., et al. 2026. FeaXDrive: Feasibility-aware trajectory-centric diffusion planning for end-to-end autonomous driving. arXiv preprint arXiv:2604.12656.

Wigness, M., Eum, S., Rogers, J. G., Han, D., and Kwon, H. 2019. A RUGD dataset for autonomous navigation and visual perception in unstructured outdoor environments. In: 2019 IEEE/RSJ international conference on intelligent robots and systems, 5000–5007.

Williams, G., Wagener, N., Goldfain, B., Drews, P., Rehg, J. M., Boots, B., et al. 2017. Information theoretic MPC for model-based reinforcement learning. In: 2017 IEEE international conference on robotics and automation, 1714–1721.

Xing, Z., Zhang, X., Hu, Y., Jiang, B., He, T., Zhang, Q., et al. 2025. GoalFlow: Goal-driven flow matching for multimodal trajectories generation in end-to-end autonomous driving. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, 1602–1611.

Yu, S., Shen, C., Dallas, J., Epureanu, B. I., Jayakumar, P., and Ersal, T. 2025. A real-time terrain-adaptive local trajectory planner for high-speed autonomous off-road navigation on deformable terrains. IEEE Transactions on Intelligent Transportation Systems, 26, 3324–3340.

Zhao, Z., Fu, T., Su, S., Du, Q., Esfahani, E. T., Dantu, K., et al. 2026. Learning when to jump for off-road navigation. In: Proceedings of robotics: Science and systems.

Zheng, Y., Li, J., Yu, D., Yang, Y., Li, S. E., Zhan, X., et al. 2024. Safe offline reinforcement learning with feasibility-guided diffusion model. In: International conference on learning representations.

Zheng, Y., Liang, R., Zheng, K., Zheng, J., Mao, L., Li, J., et al. 2025. Diffusion-based planning for autonomous driving with flexible guidance. In: International conference on learning representations.

Zhu, Z., Li, N., Sun, R., Zhao, H., and Xu, D. 2020. Off-road autonomous vehicles traversability analysis and trajectory planning based on deep inverse reinforcement learning. In: 2020 IEEE intelligent vehicles symposium.

Author biography



Linheng Li received the Ph.D. degree from Southeast University, Nanjing, China. He is currently an Associate Professor with the School of Transportation, Southeast University. His research interests include traffic flow theory in connected and intelligent transportation environments, traffic safety risk assessment, and intelligent decision-making for connected and autonomous vehicles.



Weijie Wang is currently a doctoral student in transportation engineering at the School of Transportation, Southeast University, Nanjing, China. His research interests include intelligent connected vehicles, vehicle trajectory prediction, and data-driven traffic analysis.



Tiantong Nan is with the School of Transportation, Southeast University, Nanjing, China. Research interests include autonomous driving, generative trajectory planning, and off-road navigation.

Just Accepted



Jing Gan received the Ph.D. degree in transportation engineering from Southeast University, Nanjing, China, in 2021. She is currently with the School of Modern Posts & School of Intelligent Transportation, Nanjing University of Posts and Telecommunications. Her research interests include traffic environmental assessment, and intelligent connected transportation safety.



Xu Qu received the Ph.D. degree from Southeast University, Nanjing, China, in 2013. He is currently an Associate Professor with the School of Transportation, Southeast University, where he is the Vice Dean of the Joint Research Institute on Internet of Mobility. His research interests include connected automated vehicle highway system intelligent transportation systems, and traffic control and management.



Bin Ran received the B.S. degree from Tsinghua University, Beijing, China, in 1986, the M.S. degree from the University of Tokyo, Tokyo, Japan, in 1989, and the Ph.D. degree from the University of Illinois Chicago, Chicago, IL, USA, in 1993. He is currently the Vilas Distinguished Achievement Professor and Director of the ITS Program at the University of Wisconsin–Madison. His research interests include connected autonomous mobility, intelligent transportation systems, dynamic transportation network models, traffic simulation, and smart cities.

Just Accepted