

# Drive Anywhere: Generalizable End-to-end Autonomous Driving with Multi-modal Foundation Models

Tsun-Hsuan Wang<sup>1</sup>, Alaa Maalouf<sup>1</sup>, Wei Xiao<sup>1</sup>, Yutong Ban<sup>2</sup>,  
 Alexander Amini<sup>1</sup>, Guy Rosman<sup>3</sup>, Sertac Karaman<sup>4</sup>, Daniela Rus<sup>1</sup>

<https://drive-anywhere.github.io>

**Abstract**—As autonomous driving technology matures, end-to-end methodologies have emerged as a leading strategy, promising seamless integration from perception to control via deep learning. However, existing systems grapple with challenges such as unexpected open set environments and the complexity of black-box models. At the same time, the evolution of deep learning introduces larger, multimodal foundational models, offering multi-modal visual and textual understanding. In this paper, we harness these multimodal foundation models to enhance the robustness and adaptability of autonomous driving systems. We introduce a method to extract nuanced spatial features from transformers and the incorporation of latent space simulation for improved training and policy debugging. We use pixel/patch-aligned feature descriptors to expand foundational model capabilities to create an end-to-end multimodal driving model, demonstrating unparalleled results in diverse tests. Our solution combines language with visual perception and achieves significantly greater robustness on out-of-distribution situations.

## I. INTRODUCTION

The rapid technological advancement in autonomous driving has emerged as a pivotal innovation that shifts control from human hands to AI and sensors, promising safer roads, enhanced mobility, and unparalleled efficiency. In the pursuit of autonomous driving, end-to-end methodology offers a paradigm shift toward a holistic construction of the system that encompasses everything from perception to control. Such an approach has (i) more flexibility with minimal assumptions related to the design or functioning of sub-components, and (ii) better integrality toward an ultimate, unified goal in system performance evaluation and targeted optimization. Notably, the ongoing advancement in establishing end-to-end autonomous systems is propelled by the amalgamation of deep learning techniques by training models on extensively annotated datasets. However, prevalent systems exhibit the following prominent limitations.

**(i) Open set environments:** self-driving vehicles operate in extremely diverse scenarios that are impractical to fully capture within training datasets. When these systems encounter situations that deviate from what they’ve learned (i.e., out-of-distribution (OOD) data), performance deteriorates.

**(ii) Black-/gray-box models:** the ubiquitous use of complex, advanced machine learning models complicates the task of pinpointing the root causes of failures in autonomous systems. Unraveling the intricate interactions and identifying which learned concepts, objects, or even individual pixels contribute to incorrect behavior can be a daunting task.

\*This work is supported by Toyota Research Institute (TRI) and Capgemini Engineering. It, however, reflects solely the opinions and conclusions of its authors and not TRI or any other Toyota entity.

<sup>1</sup>MIT CSAIL, <sup>2</sup>Shanghai Jiao Tong University, <sup>3</sup>TRI, <sup>4</sup>MIT LIDS.

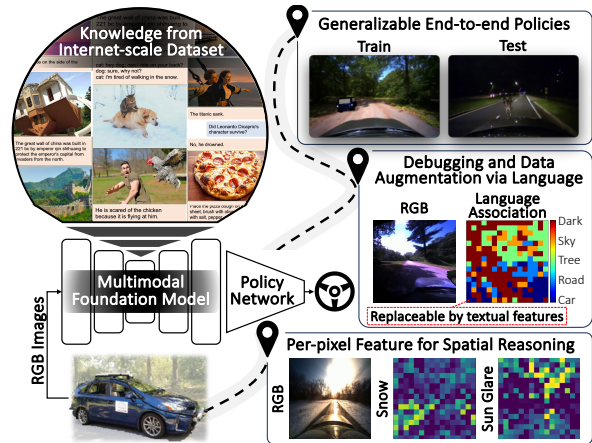


Fig. 1: We harness the power of multimodal foundation models in end-to-end driving to enhance generalization and leverage language for data augmentation and debugging.

On the positive side, deep learning is undergoing a transformative phase of significant advancement, characterized by the emergence of even larger and multimodal models [1], [2]. Trained on immense datasets that encompass billions of images, text segments, and audio clips, these models leverage knowledge gleaned from internet-scale resources to edge closer to achieving common-sense understanding. They have demonstrated exceptional efficacy in adapting to dynamic, open-set environments [3], [4]. In addition, the incorporation of language modality serves a dual purpose: not only does it offer an interface that is straightforwardly comprehensible by human users, but it also furnishes a concise yet rich representation of information that may sufficiently describe the underlying decision-making of autonomous systems.

In this work, we aim to harness the power of multimodal foundation models to enhance the generalization and reliability of end-to-end autonomous driving systems. Importantly, rather than relying on explicitly-defined data formats like scene descriptions or segmentation maps, we exploit the latent features during model inference to preserve all information pertinent to the models’ reasoning process.

**Foundation models as feature extractors.** While these models exhibit certain favorable characteristics for attaining open-set, multimodal representations, they do not translate seamlessly to autonomous driving. The significant constraint arises from the fact that these models are primarily designed for image input consumption, resulting in the generation of a singular vector representation for the entire image within an embedding space. However, decision-making in autonomous

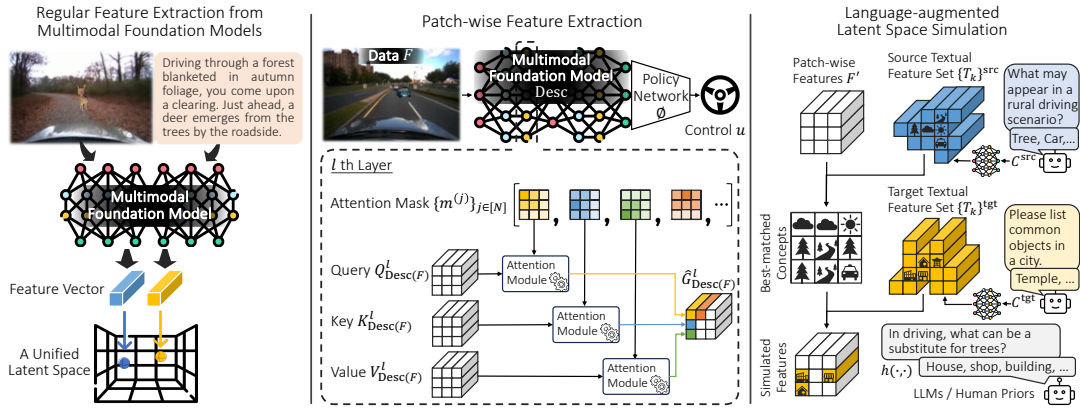


Fig. 2: **Overview.** *Left:* Feature extraction from multimodal foundation models maps data in different modalities (e.g., image, text) to feature vectors in a unified latent space. *Middle:* We introduce a generic method for patch-wise feature extraction that preserves spatial information critical for end-to-end driving; this involves constructing attention masks anchored at each patch location to focus on specific regions (depicted by the coloring) for the attention module. *Right:* The multimodal representations with language modality enable seamless integration with LLMs; this allows to simulate latent features by substituting the original features  $F^l$  with contextually relevant language features (e.g., trees  $\rightarrow$  house, shop, building).

driving demands more than just semantic scene descriptions; it also requires nuanced spatial and geometric information. To address this, we present a generic method to extract per-patch features from transformer-based architectures that is broadly applicable to a wide range of foundation models.

**Simulation using language.** Multimodal representations map data from various modalities into a unified embedding space, offering two key advantages for policies trained to operate in this space: (i) cross-modality feature inspection, and (ii) feature manipulation in modalities distinct from sensor measurements. These capabilities facilitate latent space simulation, allowing target features tied to specific concepts to be swapped out with features from other desired concepts. For instance, one could replace a 'car' feature from images in an urban driving scenario with a 'deer' feature, without requiring actual sensor data synthesis for a deer. This method is valuable for data augmentation and policy debugging. Notably, leveraging the language modality as a conceptual representation enables integration with large language models (LLMs), unlocking enhanced common-sense reasoning capabilities to enrich simulation complexity.

**Our contributions.** We bridge the gap between the robust multimodal open-set capabilities demonstrated by foundation models and the advanced reasoning capabilities expected of futuristic autonomous systems - enabling OOD, end-to-end, multimodal, and more explainable autonomy. We contribute:

- A novel approach to apply end-to-end open-set (any environment/scene) multimodal autonomous driving that achieves SOTA results in both in-distribution and OOD testing, and is able to provide driving decisions from representations queryable by image and text.
- A novel mechanism to extract pixel/patch-aligned features, extending the capabilities of multimodal foundation models that typically yield image-level vectors.
- A latent space simulation technique augmented with language modality for both data augmentation in training and counterfactual reasoning in policy debugging.

- Extensive analysis in photo-realistic simulated environments to demonstrate enhanced generalization of end-to-end driving policies across diverse scenarios; our system can drive seamlessly in environments not seen during training and avoid obstacles not trained on.
- Deployment and validation of our methods on a full-scale autonomous vehicle in real-world environments.

## II. RELATED WORK

**End-to-end driving.** Neural networks trained to process entirely from perception to control in autonomous vehicles have displayed significant potential for various driving abilities [5]–[9]. Nevertheless, these networks encounter challenges in acquiring robust models on a large scale, as they demand extensive training data that is both time-consuming and costly to gather [10]. These challenges not only incur substantial expenses but also pose potential safety risks [11]. Consequently, training and assessment of robotic controllers in simulated environments have emerged as a viable alternative [10], [12]–[14]. However, even these simulated environments can not cover enough scenarios, making trained networks (systems) highly sensitive to scenarios that differ from their training data. Training policies via intermediate visual abstraction is another approach [15]–[17], aligning with our work yet without multimodal foundation models.

**Foundation models in robots.** Recent strides in robotics have embraced foundational models, showcasing their ability to interact adeptly in dynamic open-set scenarios, e.g., for control and planning [18]–[22], for 3D mapping [23], [24], detection and following systems [25]–[28], and 3D scene segmentation and understanding [3], [4]. Moreover, these models have demonstrated their versatility across multiple data modalities [3], [25], [29]–[32] marking a new era of robots that can reason and interact wisely with the environment. Specifically in driving, explainable and language-based representations have been of interest for the ability to introspect and counterfactually reason about event [33]–[37]

**Pixel/patch aligned descriptors.** Several approaches for extracting per-pixel feature descriptors via foundation models were suggested [3], [38], however, they are either (i) not multimodal [38], (ii) trained with a specific focus on aligning foundation features with 2D pixels, and thus, these models tend to lose a substantial number of concepts as part of the fine-tuning process [39], or (iii) relies on the use of a universal segmentation model such as SAM [40], FastSAM [41], or Mask2Former [42] for extracting masks [3], and then applies the foundation models on crops of these masks to extract mask-aligned features, such methods are by definition inefficient for realtime applications and might yield not meaningful features when applying the foundation models on small crops. Finally, they might miss important regions in the image due to the used segmentor limitations.

### III. METHOD

**The end-to-end autonomous driving problem** involves designing a control system  $\phi$  that produces steering and acceleration commands based on a continuous stream of perception data  $F \in \mathbb{R}^{H \times W \times 3}$  (RGB imagery here), acquired through vehicle-mounted sensors  $u = \phi(F)$ . We propose enhancing  $\phi$  by substituting the raw frames  $F$  with a dense feature representation  $F' \in \mathbb{R}^{H' \times W' \times D}$  extracted via a multimodal foundation model  $\text{Desc}$ , where  $(H', W')$  is the resolution of the dense features in the spatial dimensions and  $D$  is the number of channels, i.e.,  $u = \phi(F') = \phi(\text{Desc}(F))$ . We first present to generate these multimodal features, preserving both spatial and semantic informations. We then describe how to employ language modality for latent space simulation, useful for both data augmentation and policy debugging.

#### A. Patch-wise Feature Extraction

Given a foundation model  $\text{Desc} : \mathbb{R}^{H \times W \times 3} \rightarrow \mathbb{R}^D$  of  $L$  layers, an input image/frame  $F \in \mathbb{R}^{H \times W \times 3}$ , and a desired resolution  $H' \leq H, W' \leq W$ , the goal is to extract a feature descriptors tensor  $F' \in \mathbb{R}^{H' \times W' \times D}$ , such that  $F'$  encapsulates all the semantic information of  $F$  and maintains its location in the scene. For simplicity, we set  $H', W'$  to be equal to the number of (non-overlapping) patches used to divide the input image  $F$  when applying  $\text{Desc}$  on it (in what follows we will demonstrate how  $H', W'$  can be any number smaller than  $H, W$ ) and  $N = H'W'$ . For an integer  $i > 1$ , we use  $[i]$  to denote the set  $\{1, \dots, i\}$ , and for every layer  $\ell \in [L]$ , we use  $Q_{\text{Desc}(F)}^\ell, K_{\text{Desc}(F)}^\ell \in \mathbb{R}^{N \times D_k}, V_{\text{Desc}(F)}^\ell \in \mathbb{R}^{N \times D}$  to denote the resulted query, key, and value matrices in the  $\ell$ th attention layer, when applying  $\text{Desc}$  on  $F$ . We provide a mechanism to extract features for a specific patch (or area in the image)  $F'^{(j)}$ , where  $j \in [N]$ . Notably, this mechanism can be applied at any layer of most transformer-based foundation models such as CLIP [1], DINO [43], [44], and BLIP [2], [45].

**Masking for patch- or region-aligned feature extraction.** When extracting  $F'^{(j)}$ , we introduce an attention mask  $m^{(j)} = (m_1^{(j)}, \dots, m_N^{(j)}) \in \mathbb{R}^N$ . Each element  $m_i^{(j)} \in [0, 1]$  of this vector determines how much the  $i$ th patch should contribute to the desired patch feature  $F'^{(j)}$ . For example, if we want

to completely ignore patch number  $i$ , simply set  $m_i = 0$  and  $m_k = 1, \forall k \in [N] \setminus i$ . We utilize  $m$  to extract features as follows: (1) Set  $r \in (-\infty, 0)$  as the parameter to control the strength of the masking; the larger  $|r|$ , the higher effects of the masking. (2) Define the matrix  $G_{\text{Desc}(F)}^\ell$  as the matrix multiplication of the key and query matrices at the  $\ell$ th attention layer:

$$G_{\text{Desc}(F)}^\ell := Q_{\text{Desc}(F)}^\ell (K_{\text{Desc}(F)}^\ell)^T.$$

(3) Given the matrix  $M^{(j)} = [m^{(j)}, \dots, m^{(j)}]^T \in \mathbb{R}^{N \times N}$ , we obtain a masked version of  $G_{\text{Desc}(F)}^\ell$  as:

$$\hat{G}_{\text{Desc}(F)}^{\ell, (j)} = G_{\text{Desc}(F)}^\ell + (\mathbf{1} - M^{(j)}) \cdot r,$$

where  $\mathbf{1} \in \mathbb{R}^{N \times N}$  is an all-ones matrix. This operation sets the attention scores (in the matrix  $\hat{G}_{\text{Desc}(F)}^{\ell, (j)}$ ) for "non-contributing" patches (where there corresponding  $m_i$  close to 0) to be close to  $r$  (low value), effectively masking them out. The  $\mathbf{1} - M^{(j)}$  term ensures that the patches with a corresponding attention mask equal to 1 have an added softmax'ed score of 0 (no modification), and a very low value (effectively  $r$ ) if the corresponding attention mask is near 0. (4) With the modified attention scores, we obtain the final attention weights with the softmax function as:

$$F'^{(j)} := \text{Desc}^{\ell \rightarrow} \left( \text{SoftMax}(\hat{G}_{\text{Desc}(F)}^{\ell, (j)}) (V_{\text{Desc}(F)}^\ell)^T \right),$$

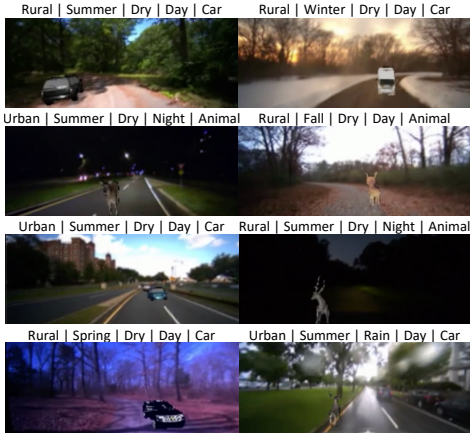
where  $\text{Desc}^{\ell \rightarrow}$  is the rest of the foundation model after the  $\ell$ th layer. Notably, this technique can be directly extended to any region-wise feature extraction by generalizing the definition of patches to arbitrarily-shaped regions.

**Setting the masks.** We define the  $i$ th entry of  $m^{(j)}$  to correspond to "how much patch  $i$  contributes to the semantic information of patch  $j$ ". Analogous to convolutional kernels, "close" neighbors (patches) may contribute more than far ones. Let  $(x_i, y_i)$  be the row-stacked ordering of the image grid after patching. Let  $\text{dist}(i, j)$  denotes the distance between patch  $i$  and  $j$  as  $\text{dist}(i, j) := \|(x_i, y_i) - (x_j, y_j)\|_z$ , where  $z \geq 1$  defines the norm. We set  $m_i^{(j)} := f(\text{dist}(i, j))$ ; where  $f$  can be  $\begin{cases} 0, & \text{if } \text{dist}(0, 1) > \alpha \\ 1, & \text{otherwise} \end{cases}, 1/2^{\text{dist}(i, j)}$ , or  $1/\text{dist}(i, j)$ , etc.

**Spatial resolution/number of patches  $N$ .** Increasing the spatial resolution enhances the foundation models' spatial features, as higher resolution allows for more granular, non-overlapping patches. For our applications, this granularity is beneficial. We adapt ViTs to extract overlapping patches during inference as in [25], [38], interpolating their positional encoding accordingly. This yields multimodal features with a finer spatial resolution notably without requiring additional training. In our empirical experiments, we have observed that this modification consistently performs well.

#### B. Language-augmented Latent Space Simulation

Each patch feature  $F'^{(j)}$  incorporates language modality, enabling seamless integration with LLMs. We exploit this property by conducting latent space simulations, where we replace  $F'^{(j)}$  with alternative textual features to simulate different scenarios. We opt for feature replacement over



Setting	Scenarios					Methods			
	Scene	Season	Weather	Time	Actor	No-FM	MF	I-ViT	Ours
ID	Rural	Summer	Dry	Day	Car	1.00	0.72	1.00	1.00
OOD	Rural	Spring	Dry	Day	Car <sup>†</sup>	0.84	0.42	0.86	0.96
		Summer	Dry	Night	Car <sup>†</sup>	0.30	0.35	0.80	0.89
		Fall	Dry	Day	Car <sup>†</sup>	0.90	0.74	0.95	0.91
		Winter	Snow	Day	Car <sup>†</sup>	0.14	0.42	0.88	0.96
		Spring	Dry	Day	Animal	0.85	0.39	0.89	0.95
	Urban	Summer	Dry	Night	Animal	0.29	0.39	0.59	0.85
		Fall	Dry	Day	Animal	0.87	0.71	0.95	0.88
		Winter	Snow	Day	Animal	0.15	0.45	0.87	0.95
		Summer	Dry	Day	Car <sup>†</sup>	0.55	0.50	0.77	0.62
		Summer	Rain	Day	Car <sup>†</sup>	0.69	0.43	0.81	0.81
Urban	Summer	Dry	Night	Car <sup>†</sup>	0.45	0.42	0.81	0.78	
	Summer	Dry	Day	Animal	0.58	0.50	0.80	0.64	
	Summer	Rain	Day	Animal	0.66	0.43	0.83	0.78	
	Summer	Dry	Night	Animal	0.45	0.36	0.86	0.81	

TABLE I: **OOD generalization.** The left figures are illustrations for the scenarios. <sup>†</sup> indicates car types different from training. *ID* is in distribution (evaluated on data not in the training dataset yet of the same scenario). *OOD* is out-of-distribution.

arithmetic operations, as the latent space may not necessarily adhere to a Euclidean metric structure where standard arithmetic would be applicable. The procedure is as follows: (1) Obtain a set of concepts in natural language that may be relevant to autonomous driving from LLMs and compute their corresponding textual feature,  $T_k = \text{Desc}(c_k)$ , where  $c_k \in \mathcal{C}^{\text{src}/\text{tgt}} = \text{LLM}(\langle \text{questions} \rangle)$  where  $\mathcal{C}^{\text{src}}$  refers to the set that may appear in the image feature and  $\mathcal{C}^{\text{tgt}}$  is the set of the desired substitutes.

(2) Find the best match of the patch feature via search,  $T_{F'(j)} = \arg \max_{k \in \|\mathcal{C}^{\text{src}}\|} g(F'(j), T_k)$  where  $g(\cdot, \cdot)$  is the similarity measure, e.g., dot product. This can be improved by more advanced optimization techniques like [46] [47].

(3) Manipulate the dense feature descriptor  $F'$  by replacing  $F'(j)$  with sensible textual features  $h(T_{F'(j)}, \{T_k\}_{k \in \|\mathcal{C}^{\text{tgt}}\|})$  under conditions like similarity above certain threshold or stochasticity. The function  $h$  can be human prior or LLMs that conceptually answer the question of *what may be a plausible substitute from  $\mathcal{C}^{\text{tgt}}$  under the current context.*

This technique is useful for counterfactual testing in policy debugging or data augmentation during training (more concrete examples shown in Sec. IV-B and Sec. IV-C).

#### IV. RESULTS

**Hardware setup.** We collected data and deployed learned policies on a full-scale vehicle (2019 Lexus RX 450H) retrofitted for autonomous driving. The car is equipped with an NVIDIA 4070 Ti GPU and an AMD Ryzen 7 3800X 8-Core Processor. For perception, we employ a 30Hz BFS-PGE-23S3C-CS camera offering a 130° horizontal FoV at a resolution of 960 x 600 pixels. Also, the car also features inertial measurement units (IMUs), wheel encoders, and an OxTS d-GPS system for precise odometry estimation.

**Tasks and evaluation metrics.** We focus on a generic driving task involving both lane-following and obstacle avoidance. Failure conditions are defined as: (i) veering off the lane boundary, (ii) colliding with objects (or approaching them too closely in the real world for safety reasons), and (iii) deviating from the lane’s direction by more than 30°. In simulations, we utilize a “soft” success rate as a performance metric, which gauges the duration the car can operate

without encountering any failure conditions, normalized by a predefined time horizon for each trial. We conducted 100 trials, each with an approximate 20-second time horizon. For real-world testing, we tally the number of interventions made by the safety driver, using criteria that align with the aforementioned failure conditions. Unless stated otherwise, all experiments adhere to a closed-loop control setting.

**Data and learning.** We obtained data from a data-driven simulator VISTA [10] to augment a real-world dataset with diverse synthetic data in a closed-loop control setting. This simulation relies on approximately two hours of real-world driving data, gathered under varying conditions including different times of day, weather conditions, scenes, and seasonal variations. We employ a training approach known as Guided Policy Learning [10], [48], which takes advantage of privileged information within the simulator to guide the learning of the image-based policies. Ground-truth control signals for training are produced using a Proportional-Integral-Derivative (PID) controller for lane-following tasks and Control Barrier Functions (CBFs) [49] for obstacle avoidance. All models are trained with L2 loss via Adam optimizer at a learning rate of  $10^{-3}$ , employing a plateau scheduler with a factor of 1 and a patience of 10, for  $10^6$  iterations. We uses BLIP2 [2] as it is SOTA in various benchmarks like zero-shot VQA, image-text retrieval, etc.

##### A. Out-of-Distribution and Cross-modality Generalization

In Tab. I, we evaluate the out-of-distribution (OOD) generalization capabilities of end-to-end policies employing various feature extractors. Our baseline comparisons include: (i) *No Foundation Model (No-FM)* [8], [10], which utilizes a CNN-based model (transformer-based architectures yielded comparable results) trained from scratch without leveraging foundation models; (ii) *Mask-based Features (MF)* [25], [41], which initially applies segmentation [40], then extract global feature vectors [1] for each masked/cropped image, and assign the vector to each pixel within the corresponding region (an approach adapted from similarity measure technique); (iii) *Inherent ViT Features (I-ViT)* [38], which suggests using ViT models [50] interlayers outputs (per-patch corresponding) of the key, value, and query matrices

Thresh.	null	0.03	0.05	0.08	0.10	0.20	$\infty$
Perform.	0.657	0.658	0.661	0.751	0.826	0.984	1.000

TABLE II: **Cross modality generalization.** We test the generalization of training in image modality and deploying in language modality by replacing sufficiently similar image features (determined by the threshold) with textual features.

Debugging Concepts	Perform.	Failure at		
		Lane Stable	Avoidance	Recovery
Image Feature	1.000	0.000	0.000	0.000
Car, Road, Tree, Sky	0.267	0.149	0.453	0.117
+ Car & Road + Dark	0.536	0.028	0.381	0.056
+ Car Exterior Parts	0.657	0.040	0.237	0.074

TABLE III: **A debugging tool.** We use LLMs to propose potentially relevant concepts for language-augmented latent space simulation to inspect the decision-making of a policy.

as "inherent" per-pixel/patch features. We also tried CLIP [1] and BLIP2 [2] in *I-ViT* and found them giving worse similarity map [38] than DINO [50], similar findings for CLIP to [25]. This may result from CLIP and BLIP2 being trained for image-text alignment unlike DINO for visual features only. Lastly, we found feature extractors pre-trained on ImageNet give comparable results to training from scratch as in *No-FM*, thus not included in Tab. I. All feature extractors are frozen and followed by a policy network utilizing a consistent transformer-based architecture (a simple ViT with patch size  $1 \times 1$ , depth 3, transformer dimension 128, heads 4, head dimension 32, MLP dimension 256). Firstly, we note that *MF* underperforms in both in-distribution and out-of-distribution settings. We propose two possible explanations for this: (1) the process of masking out non-target regions may inadvertently eliminate valuable contextual information, and (2) while masking is generally effective for objects, it creates ambiguous image crops for more abstract categories, often referred to as "stuffs" [51]. For instance, masking out all elements except the road in a rural setting results in an indistinct, yellowish region. Additionally, *I-ViT* and *Ours* surpass *No-FM* in OOD settings, scenarios, highlighting the benefits of utilizing foundation models as feature extractors to enhance generalization. Notably, *I-ViT* works with visual features only unlike our method for multiple modalities; we envision the extraction of inherent pixel-aligned multimodal features to be a promising future direction.

In Tab. II, we evaluate the ability of our policy to generalize from image to language modalities under the ID setting in Tab. I. Notably, our policy is trained exclusively on image data. To generate cross-modality features, we employ the following procedure: (i) calculate features (from the LLM encoder in BLIP2) for a predefined set of natural language concepts (e.g., road, car); (ii) identify the best-matching textual feature for each image feature pixel, along with the degree of similarity (we manually check a diverse subset of images to verify that the best matches are reasonable); (iii) replace the image feature at pixels where the similarity exceeds a certain threshold with the corresponding textual feature. A null threshold implies driving solely based on textual features, while an infinity threshold denotes reliance exclusively on image features. Our observations indicate that

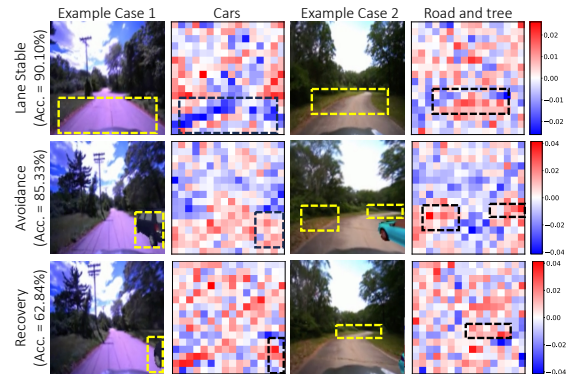


Fig. 3: **Linear classification on driving maneuver.** We report accuracy (the leftmost labels) and coefficients of linear classifier associated with a concept (2nd, 4th columns) along with typical examples of each maneuver (1st, 3rd columns).

the policy performs reasonably well in cross-modality settings, both qualitatively and quantitatively. These results offer promising empirical evidence for the viability of language-augmented latent space simulations.

### B. Debugging and Inspecting Policies With Language

In Tab. III, we present a case study focused on policy debugging through language-augmented latent space simulation. Our procedure is as follows: (i) we consult Large Language Models (LLMs) to generate a base set of natural language concepts relevant to, for example, a rural driving scenario; (ii) we collect driving policy rollouts along with intermediate features, filtering out less pertinent concepts based on similarity statistics, and potentially human judgment; (iii) we then evaluate the policy by replacing image pixel features with textual features drawn from various subsets of these relevant concepts; (iv) lastly, we pinpoint specific concepts whose presence across subsets leads to significant performance changes. It's worth noting that the third step involves combinatorial complexity and may benefit from some degree of human intervention for enhanced efficiency. We highlight two major discoveries. Firstly, the concepts *Car & Road* and *Dark* play critical roles in both recovery and stable lane maneuver. Considering *Car* and *Road* in conjunction is essential, as the policy uses the end of the road as a navigational reference to maintain lane position. Typically, this reference point in the image is relatively small and encompasses both concepts simultaneously. Moreover, this combined feature serves as a natural differentiator, setting itself apart from standalone features like *Road* or *Tree* that may appear elsewhere in an image. As for the concept of *Dark*, it exposes a loophole where the policy takes advantage of simulation artifacts. Specifically, when the car deviates significantly from the lane center, the simulation produces dead or black pixels due to the absence of content to render.

In Fig. 3, we display linear classification results for various maneuver types using low-dimensional projected features. This evaluation diverges from a closed-loop control setting, as "features-to-control" requires more complex models than linear functions. Our research question aims to explore the existence of simple decision boundaries that differentiate


<b>RSDDC</b>	<b>RSDNC</b>	<b>RFDDC</b>	<b>RWSDC</b>	<b>RSDDA</b>
+3.43%	-2.49%	+8.32%	+3.12%	-0.48%
<b>RSDNA</b>	<b>RFDDA</b>	<b>RWSDA</b>	<b>USDDC</b>	<b>USRDC</b>
-5.09%	+9.83%	+1.02%	+12.49%	+14.44%
<b>USDNC</b>	<b>USDDA</b>	<b>USRDA</b>	<b>USDNA</b>	<b>All</b>
+10.80%	-0.65%	+13.08%	+12.75%	<b>+5.47%</b>

TABLE IV: **Improved generalization from data augmentation.** We augment training with unseen yet potentially relevant concepts from LLMs via language-augmented latent space simulation to improve performance. The labels are the scenarios in Tab. I (*RSDDC* is Rural, Spring, Dry, Day, Car).

higher-level maneuvers (e.g., avoidance vs. 0.5 steering angle) and examine how feature spatial distribution within an image influences classification. We follow this procedure: (i) apply K-means clustering (with 4 clusters) to feature vectors from policy rollouts and project them into a lower-dimensional space based on distances to cluster centers; (ii) perform linear support vector classification on these projected features of all  $N$  patches; (iii) anchor the cluster centers by matching with the most relevant textual features proposed by LLMs. In Fig. 3, the cluster centers pertain to *Cars*, *Road & Tree*, *Forests and Jungles*, and *Campgrounds*. Due to space limitations, we focus on the first two. Training on 10 trajectories and testing on 90 distinct ones yield high accuracy (in the leftmost labels), affirming simple decision boundaries. Visualizations of typical maneuvers and classifier coefficients show more structured coefficients in the image’s lower part, closely linked to driving maneuvers. For *Car*, negative coefficients (in blue) appear in the road’s center during lane-stable maneuvers, indicating that no cars should obstruct the path. Positive coefficients (in red) are observed throughout the road during avoidance, suggesting that cars can appear anywhere. During recovery, these coefficients are mainly positive at the road’s edges, as the ego car initiates recovery only after passing other vehicles. For *Road & Tree*, positive coefficients clutter at the image’s edges during avoidance, reflecting the ego car’s heading deviation from the road to evade obstacles. In lane-stable and recovery, these coefficients are variably distributed in the middle, aligning with the ego car’s road-oriented direction.

### C. Data Augmentation using Language

In Table IV, we showcase the performance improvements achieved through data augmentation using language-augmented latent space simulation. Our procedure is as follows: (i) We first identify a set of target concepts likely to appear in the training data that are candidates for replacement, selecting *Tree* and *Dark* for this experiment; (ii) We then consult LLMs to suggest possible replacement concepts; in this experiments, they are broadly defined as *any non-drivable objects or entities likely to appear in a driving scenario*; (iii) Finally, we randomly swap image pixel features—those exhibiting high similarity to the target concepts—with the textual features corresponding to these suggested replacement concepts. We note performance improvements in most OOD scenarios, with the exceptions of *RSDNA* and *RSDNC*, where we observe a non-marginal decline in performance. Both represent rural, nighttime con-



Crashes	Lane Stable	Obstacle Avoidance				
		Pedestrian	Roadblock	Cone	Chair	House
No-FM	7.3 /km	8/10	9/10	10/10	9/10	10/10
Ours	0 /km	0/10	0/10	0/10	0/10	1/10

TABLE V: **Real Car Test.** We verify the generalization capability of our method on a real autonomous car.

ditions, characterized by extremely low light, as depicted in row 3, column 2 of Tab. I. These low-light environments make data augmentation particularly error-prone when targeting the concept of *Dark* for replacement.

### D. Real Car Deployment

In Tab. V, we present the outcomes of tests conducted on a full-scale autonomous vehicle within a rural test track. These tests were carried out during the summer season and spanned various times of the day. Importantly, the evaluation took place on different road segments and occurred two years subsequent to the summer data used in the training set, allowing us to assess performance amid noticeable changes in the environment. We assess the system’s proficiency in lane-following and its ability to avoid a variety of objects not encountered during training. These objects include pedestrians, roadblocks, traffic cones, chairs, and even a toy house. Some of these test scenarios are illustrated in the top row of Tab. V. Our approach, which utilizes foundation models, yields near-flawless driving performance, further validating its effectiveness in generalizing to real-world robotic systems. However, it’s worth noting that the inference speed is somewhat limited, averaging around 3 fps, compared to non-foundation-model-based policies, which achieve between 10 to 30 fps depending on the architecture. That said, we have not yet focused on optimizing runtime performance, which could be improved through techniques such as quantization.

## V. CONCLUSION

This exploration into enhancing autonomous driving through multimodal foundation models has offered several lessons learned. The incorporation of these models improves the system’s adaptability in unpredictable open-set environments, emphasizing their role in advancing real-world applicability of autonomous vehicles. Our model’s blend of visual and textual understanding provides insights into the often murky decision-making processes inherent to autonomous systems, suggesting a promising trajectory for future models that prioritize both performance and transparency. The development of pixel/patch-aligned feature descriptors and latent space simulation, enriched with language modality, suggests potential for optimizing the training and debugging processes for end-to-end learning based control. Moreover, the successful deployment and performance of our methods on a real-world, full-scale autonomous vehicle provides encouraging first steps toward autonomous driving solutions that integrate multi-modal foundational models with perception.

## REFERENCES

- [1] A. Radford, J. W. Kim, C. Hallacy, A. Ramesh, G. Goh, S. Agarwal, G. Sastry, A. Askell, P. Mishkin, J. Clark, *et al.*, “Learning transferable visual models from natural language supervision,” in *International conference on machine learning*. PMLR, 2021, pp. 8748–8763.
- [2] J. Li, D. Li, S. Savarese, and S. Hoi, “Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models,” *arXiv preprint arXiv:2301.12597*, 2023.
- [3] K. M. Jatavallabhula, A. Kuwajerwala, Q. Gu, M. Omama, T. Chen, S. Li, G. Iyer, S. Saryazdi, N. Keetha, A. Tewari, *et al.*, “Conceptfusion: Open-set multimodal 3d mapping,” *arXiv preprint arXiv:2302.07241*, 2023.
- [4] S. Peng, K. Genova, C. Jiang, A. Tagliasacchi, M. Pollefeys, T. Funkhouser, *et al.*, “Openscene: 3d scene understanding with open vocabularies,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023, pp. 815–824.
- [5] D. A. Pomerleau, “Alvinn: An autonomous land vehicle in a neural network,” *Advances in neural information processing systems*, vol. 1, 1988.
- [6] M. Bojarski, D. Del Testa, D. Dworakowski, B. Firner, B. Flepp, P. Goyal, L. D. Jackel, M. Monfort, U. Muller, J. Zhang, *et al.*, “End to end learning for self-driving cars,” *arXiv preprint arXiv:1604.07316*, 2016.
- [7] A. Amini, G. Rosman, S. Karaman, and D. Rus, “Variational end-to-end navigation and localization,” in *2019 International Conference on Robotics and Automation (ICRA)*. IEEE, 2019, pp. 8958–8964.
- [8] T.-H. Wang, W. Xiao, M. Chahine, A. Amini, R. Hasani, and D. Rus, “Learning stability attention in vision-based end-to-end driving policies,” in *Learning for Dynamics and Control Conference*. PMLR, 2023, pp. 1099–1111.
- [9] W. Xiao, T.-H. Wang, R. Hasani, M. Chahine, A. Amini, X. Li, and D. Rus, “BarrierNet: Differentiable control barrier functions for learning of safe robot control,” *IEEE Transactions on Robotics*, 2023.
- [10] A. Amini, T.-H. Wang, I. Gilitschenski, W. Schwarting, Z. Liu, S. Han, S. Karaman, and D. Rus, “Vista 2.0: An open, data-driven simulator for multimodal sensing and policy learning for autonomous vehicles,” in *2022 International Conference on Robotics and Automation (ICRA)*. IEEE, 2022, pp. 2419–2426.
- [11] A. Kendall, J. Hawke, D. Janz, P. Mazur, D. Reda, J.-M. Allen, V.-D. Lam, A. Bewley, and A. Shah, “Learning to drive in a day,” in *2019 International Conference on Robotics and Automation (ICRA)*. IEEE, 2019, pp. 8248–8254.
- [12] R. Tedrake *et al.*, “Drake: Model-based design and verification for robotics,” 2019.
- [13] A. Dosovitskiy, G. Ros, F. Codevilla, A. Lopez, and V. Koltun, “Carla: An open urban driving simulator,” in *Conference on robot learning*. PMLR, 2017, pp. 1–16.
- [14] S. Shah, D. Dey, C. Lovett, and A. Kapoor, “Airsim: High-fidelity visual and physical simulation for autonomous vehicles,” in *Field and Service Robotics: Results of the 11th International Conference*. Springer, 2018, pp. 621–635.
- [15] M. Müller, A. Dosovitskiy, B. Ghanem, and V. Koltun, “Driving policy transfer via modularity and abstraction,” *arXiv preprint arXiv:1804.09364*, 2018.
- [16] M. Toromanoff, E. Wirbel, and F. Moutarde, “End-to-end model-free reinforcement learning for urban driving using implicit affordances,” in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2020, pp. 7153–7162.
- [17] A. Behl, K. Chitta, A. Prakash, E. Ohn-Bar, and A. Geiger, “Label efficient visual abstractions for autonomous driving,” in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2020, pp. 2338–2345.
- [18] S. Tellex, N. Gopalan, H. Kress-Gazit, and C. Matuszek, “Robots that use language,” *Annual Review of Control, Robotics, and Autonomous Systems*, vol. 3, pp. 25–55, 2020.
- [19] Y. Bisk, A. Holtzman, J. Thomason, J. Andreas, Y. Bengio, J. Chai, M. Lapata, A. Lazaridou, J. May, A. Nisnevich, *et al.*, “Experience grounds language,” *arXiv preprint arXiv:2004.10151*, 2020.
- [20] M. Ahn, A. Brohan, N. Brown, Y. Chebotar, O. Cortes, B. David, C. Finn, C. Fu, K. Gopalakrishnan, K. Hausman, *et al.*, “Do as i can, not as i say: Grounding language in robotic affordances,” *arXiv preprint arXiv:2204.01691*, 2022.
- [21] A. Brohan, N. Brown, J. Carbajal, Y. Chebotar, J. Dabis, C. Finn, K. Gopalakrishnan, K. Hausman, A. Herzog, J. Hsu, *et al.*, “Rt-1: Robotics transformer for real-world control at scale,” *arXiv preprint arXiv:2212.06817*, 2022.
- [22] S. Li, X. Puig, C. Paxton, Y. Du, C. Wang, L. Fan, T. Chen, D.-A. Huang, E. Akyürek, A. Anandkumar, *et al.*, “Pre-trained language models for interactive decision-making,” *Advances in Neural Information Processing Systems*, vol. 35, pp. 31 199–31 212, 2022.
- [23] C. Huang, O. Mees, A. Zeng, and W. Burgard, “Audio visual language maps for robot navigation,” *arXiv preprint arXiv:2303.07522*, 2023.
- [24] R. Ding, J. Yang, C. Xue, W. Zhang, S. Bai, and X. Qi, “Pla: Language-driven open-vocabulary 3d scene understanding,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023, pp. 7010–7019.
- [25] A. Maalouf, N. Jadhav, K. M. Jatavallabhula, M. Chahine, D. M. Vogt, R. J. Wood, A. Torralba, and D. Rus, “Follow anything: Open-set detection, tracking, and following in real-time,” *arXiv preprint arXiv:2308.05737*, 2023.
- [26] S. Liu, Z. Zeng, T. Ren, F. Li, H. Zhang, J. Yang, C. Li, J. Yang, H. Su, J. Zhu, *et al.*, “Grounding dino: Marrying dino with grounded pre-training for open-set object detection,” *arXiv preprint arXiv:2303.05499*, 2023.
- [27] G. Ghiasi, X. Gu, Y. Cui, and T.-Y. Lin, “Scaling open-vocabulary image segmentation with image-level labels,” in *Computer Vision—ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXXVI*. Springer, 2022, pp. 540–557.
- [28] B. Li, K. Q. Weinberger, S. Belongie, V. Koltun, and R. Ranftl, “Language-driven semantic segmentation,” *arXiv preprint arXiv:2201.03546*, 2022.
- [29] A. Ramesh, M. Pavlov, G. Goh, S. Gray, C. Voss, A. Radford, M. Chen, and I. Sutskever, “Zero-shot text-to-image generation,” in *International Conference on Machine Learning*. PMLR, 2021, pp. 8821–8831.
- [30] K. Crowson, S. Biderman, D. Kornis, D. Stander, E. Hallahan, L. Castrioto, and E. Raff, “Vqgan-clip: Open domain image generation and editing with natural language guidance,” in *European Conference on Computer Vision*. Springer, 2022, pp. 88–105.
- [31] O. Patashnik, Z. Wu, E. Shechtman, D. Cohen-Or, and D. Lischinski, “Styleclip: Text-driven manipulation of stylegan imagery,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 2085–2094.
- [32] A. Ramesh, P. Dhariwal, A. Nichol, C. Chu, and M. Chen, “Hierarchical text-conditional image generation with clip latents,” *arXiv preprint arXiv:2204.06125*, 2022.
- [33] J. Kim, T. Misu, Y.-T. Chen, A. Tawari, and J. Canny, “Grounding human-to-vehicle advice for self-driving vehicles,” Nov. 2019.
- [34] D. Omeiza, H. Webb, M. Jirotko, and L. Kunze, “Explanations in autonomous driving: A survey,” *IEEE Transactions on Intelligent Transportation Systems*, vol. 23, no. 8, pp. 10 142–10 162, 2021.
- [35] Y.-L. Kuo, X. Huang, A. Barbu, S. G. McGill, B. Katz, J. J. Leonard, and G. Rosman, “Trajectory prediction with linguistic representations,” in *2022 International Conference on Robotics and Automation (ICRA)*. IEEE, 2022, pp. 2868–2875.
- [36] S. Tan, B. Ivanovic, X. Weng, M. Pavone, and P. Kraehenbuehl, “Language conditioned traffic generation,” *arXiv preprint arXiv:2307.07947*, 2023.
- [37] Z. Zhong, D. Rempe, Y. Chen, B. Ivanovic, Y. Cao, D. Xu, M. Pavone, and B. Ray, “Language-guided traffic simulation via scene-level diffusion,” *arXiv preprint arXiv:2306.06344*, 2023.
- [38] S. Amir, Y. Gandelsman, S. Bagon, and T. Dekel, “Deep ViT features as dense visual descriptors,” *arXiv preprint arXiv:2112.05814*, 2021.
- [39] Y. Ding, L. Liu, C. Tian, J. Yang, and H. Ding, “Don’t stop learning: Towards continual learning for the clip model,” *arXiv preprint arXiv:2207.09248*, 2022.
- [40] A. Kirillov, E. Mintun, N. Ravi, H. Mao, C. Rolland, L. Gustafson, T. Xiao, S. Whitehead, A. C. Berg, W.-Y. Lo, *et al.*, “Segment anything,” *arXiv preprint arXiv:2304.02643*, 2023.
- [41] X. Zhao, W. Ding, Y. An, Y. Du, T. Yu, M. Li, M. Tang, and J. Wang, “Fast segment anything,” 2023.
- [42] B. Cheng, I. Misra, A. G. Schwing, A. Kirillov, and R. Girdhar, “Masked-attention mask transformer for universal image segmentation,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 1290–1299.
- [43] M. Oquab, T. Darcet, T. Moutakanni, H. Vo, M. Szafraniec, V. Khalidov, P. Fernandez, D. Haziza, F. Massa, A. El-Nouby, *et al.*, “Dinov2: Learning robust visual features without supervision,” *arXiv preprint arXiv:2304.07193*, 2023.

- [44] M. Caron, H. Touvron, I. Misra, H. Jégou, J. Mairal, P. Bojanowski, and A. Joulin, “Emerging properties in self-supervised vision transformers,” in *Proceedings of the IEEE/CVF international conference on computer vision*, 2021, pp. 9650–9660.
- [45] J. Li, D. Li, C. Xiong, and S. Hoi, “Blip: Bootstrapping language-image pre-training for unified vision-language understanding and generation,” in *International Conference on Machine Learning*. PMLR, 2022, pp. 12 888–12 900.
- [46] R. Gal, Y. Alaluf, Y. Atzmon, O. Patashnik, A. H. Bermano, G. Chechik, and D. Cohen-Or, “An image is worth one word: Personalizing text-to-image generation using textual inversion,” *arXiv preprint arXiv:2208.01618*, 2022.
- [47] B. Lester, R. Al-Rfou, and N. Constant, “The power of scale for parameter-efficient prompt tuning,” *arXiv preprint arXiv:2104.08691*, 2021.
- [48] S. Levine and V. Koltun, “Guided policy search,” in *International conference on machine learning*. PMLR, 2013, pp. 1–9.
- [49] W. Xiao and C. Belta, “Control barrier functions for systems with high relative degree,” in *Proc. of 58th IEEE Conference on Decision and Control*, Nice, France, 2019, pp. 474–479.
- [50] H. Zhang, F. Li, S. Liu, L. Zhang, H. Su, J. Zhu, L. M. Ni, and H.-Y. Shum, “Dino: Detr with improved denoising anchor boxes for end-to-end object detection,” 2022.
- [51] J. Dai, K. He, and J. Sun, “Convolutional feature masking for joint object and stuff segmentation,” in *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, June 2015.