

Diverse Probabilistic Trajectory Forecasting with Admissibility Constraints

Laura Calem^{*†}, Hedi Ben-Younes[†], Patrick Pérez[†] and Nicolas Thome^{*}

^{*}Conservatoire National des Arts et Métiers (CNAM), Paris, France

[†]Valeo.ai, Paris, France

Email: firstname.lastname@lecnam.net, firstname.lastname@valeo.com

Abstract—Predicting multiple trajectories for road users is important for automated driving systems: ego-vehicle motion planning indeed requires a clear view of the possible motions of the surrounding agents. However, the generative models used for multiple-trajectory forecasting suffer from a lack of diversity in their proposals. To avoid this form of collapse, we propose a novel method for structured prediction of diverse trajectories. To this end, we complement an underlying pretrained generative model with a diversity component, based on a determinantal point process (DPP). We balance and structure this diversity with the inclusion of knowledge-based quality constraints, independent from the underlying generative model. We combine these two novel components with a gating operation, ensuring that the predictions are both diverse and within the drivable area. We demonstrate on the nuScenes driving dataset the relevance of our compound approach, which yields significant improvements in the diversity and the quality of the generated trajectories.

I. INTRODUCTION

In trajectory forecasting, future prediction is inherently stochastic since the human or automated driver has only access to very partial information about other road users’ intents. It is also often multi-modal, since several admissible, yet very different driving actions can be taken at any instant by each agent. Intuitively, ignoring part of these possible future trajectories can hinder an autonomous or assisted driving system. It has been shown in [1] that a diverse future generation improves performance in planning tasks, especially mitigating the issue of conservative driving in autonomous vehicles. For these reasons, multiple-output forecasting models have emerged. The main challenge is to predict a limited number of future trajectories that capture well the available driving options for the near future. A crucial aspect is thus to control the diversity of the proposed trajectory set.

Recent works have built on generative autoencoders to sample multiple future trajectories [2]–[4]. However, the output distribution that such models provide sticks by construction to the one in the training data, which is mostly unimodal if real driving recordings are used: only a single future exists for a given past trajectory. At a higher level, some types of trajectories, such as turning rather than driving straight at an intersection, are severely under-represented.

Many generative models used for trajectory prediction [2], [4], [5], *e.g.*, based on generative adversarial networks (GANs) [6] or variational auto-encoders (VAEs) [7], [8], have no explicit control on the diversity beyond the one of the data distribution. Therefore, the dominant mode will be sampled

every time, a problem exacerbated in our context where we aim at sampling a few trajectories only that summarize well the possible futures. This observation, illustrated in Figure 1, motivates our approach for designing a probabilistic model based on a more structured diversity.

In this paper, we introduce a new method of DIVERse trajectory prediction with Admissibility constraints (DIVA) for probabilistic forecasting of road users. In particular, our approach allows the sampling of the main relevant modes of the future trajectory distribution, as illustrated in Figure 1. To achieve this goal, our contributions are:

- We introduce a diversity sampling function (DSF) based on a DPP [9]. The diversity is explicitly controlled through the definition of the DPP kernel. In particular, we introduce a new kernel adapted to the task at hand, which enforces trajectories’ end-points to be far away in the *lateral* direction (amounting to steering diversity) rather than in the *longitudinal* one (amounting to speeding diversity).
- We also control the “quality” of the sampled trajectories via a loss that penalizes violations of the driving area’s topology. We learn quality and diversity-based latent codes which we merge with a gating fusion mechanism. This enables the quality loss to filter out irrelevant trajectories predicted outside of the drivable area.
- We evaluate the performance of our system on a real-world dataset (nuScenes [11]) with a broad selection of metrics, demonstrating that trajectories that are both diverse and admissible are well produced.

II. RELATED WORK

Diversity. A growing body of research [5], [12]–[15] involves predicting a distribution of future trajectories rather than a univocal future. Many of these methods build upon an encoder-decoder architecture with sampling in the latent space, either with a traditional cVAE [2] or with more elaborated techniques [3], [4], [16]. Ramashinghe *et al.* [17] provide a mechanism for modeling the latent space as a continuous multimodal space, but assume that a distribution of admissible ground truths for each training example is available. This is often not the case in real-world driving datasets.

Several strategies have been applied to overcome this limitation. In MTP [18], a multi-output architecture is proposed, trained to encourage each mode to specialize for a distinct behavior. In recent work [19], the lane information is used




Fig. 1: **Effect of different methods on diversity.** Given a vehicle’s known past trajectory (blue) and the road layout (black and white map), multiple futures are predicted (red). (Left) In real datasets, a single future trajectory is available in training, making a standard generative model such as a conditional VAE (cVAE) unable to sample admissible options far away from the supervision. (Middle) Although Determinantal Point Processes (DPPs) [9] are appealing for sampling diverse predictions, using a standard ℓ_2 kernel as in [10] induces mostly *longitudinal* variations and may overshoot in non-drivable areas. (Right) In the proposed method, DIVA, the designed DPP kernel also considers the *lateral* deviation at destination between two trajectories, and explicitly penalizes predictions outside the drivable area when training the diversity model. Consequently, DIVA samples driving options that are diverse, including steering-wise, and admissible.

as a prior for semantic behavior decision, thus providing feasible and diverse trajectory forecasts. Park *et al.* [3] use a normalizing flow [20] decoder, and approximate the true distribution of future trajectories using the whole drivable area instead of the single ground truth, which encourages sample diversity. CoverNet [21] tackles the issue of diversity by predicting trajectories as distinct classes, where the set of possible categories is chosen to maximize the coverage on a training set. Another line of approaches uses DPPs to increase the diversity in the set of predicted trajectories. DPPs, introduced in [22] in the context of particle physics, are probabilistic models which recently gained the attention of the machine learning community [9], [23]–[25]. They have been explored for various applications such as video subset selection [26], document summarization [27], or time series forecasting [28]. GDPP [29] provides an interesting way to build the DPP kernel by matching the true diversity of the data. However, this method requires access to the ground-truth distribution of the data, which is not available in real-world driving datasets. In the context of trajectory forecasting, DPPs have been used with cVAEs in [10] and with Graph Neural Networks in [30]. In our work, we also use a DPP to improve the diversity of the predicted trajectories. We depart from these previous works by incorporating scene information in the DSF, which guides the sampling towards more admissible regions.

Admissibility. Several works explore using physical constraints to guide trajectory generation. In Neural Motion Planner [31], candidate trajectories are sampled in the space of *clothoids*, which ensures that they are dynamically feasible. CoverNet [21] generates a set of possible future trajectories by integrating the dynamic state of the vehicle. Park *et al.* [3] generate physically-admissible trajectories by setting a low acceleration prior on the predictions. While having no explicit control for admissibility, Salzmann *et al.* [4] constrain the outputs to be admissible under the vehicle’s current dynamic state. Our work differs from these works as we define admissibility with layout constraints in addition to dynamic feasibility.

III. PROPOSED APPROACH

We detail here the DIVA model for diverse trajectory prediction with admissibility constraints. DIVA builds upon a generative model to construct a latent space from which to sample codes representing future trajectories (subsection III-A). We then describe in subsection III-B the proposed method for introducing a structured diversity via a DPP kernel, while controlling the quality of the forecast with respect to the drivable area.

A. Problem formulation

Given the T_p past (and current) 2D positions of an agent and a “map” of its current environment, the multi-output forecasting task amounts to predicting N possible trajectories over the T_f future instants. Denoting $\mathbf{S} = (\mathbf{S}_p, \mathbf{S}_f) \in \mathbb{R}^{(T_p+T_f) \times 2}$ the agent’s trajectory over the whole time interval and $\mathbf{M} \in \mathbb{R}^{H \times W \times 3}$ the environment map centered on agent’s current position $\mathbf{S}_p(T_p)$ (using an RGB encoding of all static and dynamic elements in the scene, see example in Figure 2 and detail in Supplementary A), the forecasting model is trained on example pairs (\mathbf{S}, \mathbf{M}) . At runtime, it must predict for each agent in the scene N trajectory samples, $\hat{\mathbf{S}}_f^{(n)}, n = 1 \dots N$, given $(\mathbf{S}_p, \mathbf{M})$. Following [3], the temporal horizons in our experiments are set to $T_p = 12$ and $T_f = 6$, which amounts to 6 and 3 seconds respectively at 2Hz, and the number of predictions is $N = 12$.

While our method is agnostic to the specific architecture of the underlying generative model, we chose for our experiments a simple conditional variational autoencoder (cVAE), as done in [2] for trajectory prediction, which we adapt to suit our specific needs, as explained next.

Encoding. At a given instant and for a given agent in the scene, the encoding block takes $(\mathbf{S}_p, \mathbf{M})$ as input. The past trajectory is encoded by a gated recurrent unit (GRU) network [32], as $\mathbf{h} = \text{GRU}(\mathbf{S}_p)$, where $\mathbf{h} \in \mathbb{R}^{d_h}$ is the last hidden state of the recurrent network. The map of the agent’s environment is processed by a convolutional neural network to produce an embedding $\mathbf{m} = \text{CNN}(\mathbf{M})$ used as local physical constraints.

Sampling and decoding. Both embeddings \mathbf{m} and \mathbf{h} are concatenated and used to predict the parameters μ and σ of the Gaussian distribution over latent codes $\mathbf{z} \in \mathbb{R}^{d_z}$. A sampled latent code is then concatenated with \mathbf{m} and \mathbf{h} to produce the initialization for the hidden units of the decoder recurrent network. Finally, the output of this RNN decoder is passed through a series of fully-connected layers to produce the final trajectory $\hat{\mathbf{S}}_f$. In effect, N latent codes are sampled for a given (\mathbf{m}, \mathbf{h}) , yielding N distinct future trajectories.

Training the generative model. To train the underlying generative model, we use the VAE loss introduced in [33], adapted to include both inputs \mathbf{S}_p and \mathbf{M} and to reflect the predictive nature of the task rather than an autoencoding one:

$$L_{\text{cvae}}(\phi, \theta) = \mathbb{E}_{q_\phi(\mathbf{z}|\mathbf{S}_p, \mathbf{M})} [\log p_\theta(\hat{\mathbf{S}}_f|\mathbf{z}, \mathbf{S}_p, \mathbf{M})] - \text{KL}(q_\phi(\mathbf{z}|\mathbf{S}_p, \mathbf{M})\|p(\mathbf{z})), \quad (1)$$

where ϕ and θ are the parameters of the encoder and decoder respectively. The first term is the likelihood of the predicted trajectory and can be seen as a reconstruction quality term; the second term is the Kullback-Leibler divergence between the learned latent distribution q_ϕ and a prior $p(\mathbf{z})$, generally chosen to be Gaussian [2], [34] for ease of sampling from this prior. A generative model alone usually suffers from mode collapse, as no incentive is provided to produce diverse samples. In that case, the trajectories generated by the model concentrate around the main mode from the underlying trajectory distribution, as illustrated in Figure 3.

B. Structured diversity with physical constraints

Given a trained generative model, we propose to replace the sequential random sampling from the prior $p(\mathbf{z})$ with a diversity sampling function (DSF) trained to predict multiple $\hat{\mathbf{S}}_f$'s jointly. As illustrated in the lower part of Figure 2, the DSF is implemented as a small two-branch feed-forward neural network. In contrast with the generative model sampling, where the N latent codes are sampled independently in \mathbb{R}^{d_z} , the network is designed to output all the latent codes at once, producing an output in $\mathbb{R}^{N \times d_z}$.

In order to structure the diversity of the proposed trajectories, we split the DSF between diversity and quality, with each branch controlling a partial latent code. The diversity branch takes the representation of the past trajectory, \mathbf{h} , and produces N partial latent codes $\mathbf{z}_p^{(n)}$, whereas the quality branch takes the map representation \mathbf{m} and gives N partial latent codes $\mathbf{z}_m^{(n)}$. The two associated partial codes are then combined using an element-wise product to produce a final latent code $\mathbf{z}^{(n)}$. Through this gating mechanism, the map-specific constraints are imposed to the diverse set of trajectories. The corresponding training loss,

$$\mathcal{L}_{\text{dsf}} = \lambda \mathcal{L}_{\text{dpp}} + (1 - \lambda) \mathcal{L}_{\text{layout}}, \quad (2)$$

is comprised of two terms. The first term, \mathcal{L}_{dpp} , favors the diversity through an adapted DPP kernel and the second one, $\mathcal{L}_{\text{layout}}$, injects the quality constraints; $\lambda \in (0, 1)$ is a parameter

controlling the tradeoff between the two losses, as discussed in greater detail in subsection IV-D. Next, we detail these two loss terms.

1) **Diversity with a DPP kernel:** In the following section, we first provide some background on Determinantal Point Processes (DPPs), summarizing from [9] and [10], in order to give enough context for our proposed method. DPPs are probabilistic set models that allow for an explicit handling of negative correlations among sets' elements. DPPs were first used in the context of particle physics for their ability to model the repulsion between particles: contrary to sampling a uniform distribution in space, which results in some level of clumping, sampling a DPP results in a more uniform spread.

Formally, given a countable ground set Y of "items", a DPP is a distribution over the power set of Y , giving the probability to draw any part of Y . It thus defines a random set, \mathbf{A} . We focus here on the class of DPPs that are defined through a positive semi-definite kernel L as follows: for any finite subset B of Y , $\mathbb{P}[\mathbf{A} = B] \propto \det(L_B)$, where L_B is the matrix defined by L over B and the normalization constant has a closed form. It can also be shown that the probability $\mathbb{P}[\mathbf{A} \supset B]$ that the random set includes B is exactly $\det(K_B)$, where $K := (L + \text{Id})^{-1}L$. In particular, $\mathbb{P}[\mathbf{A} \supset \{a\}] = K(a, a)$ for any $a \in Y$ and $\mathbb{P}[\mathbf{A} \supset \{a, b\}] = K(a, a)K(b, b) - K(a, b)^2$ for any pair. The latter provides insight into the repulsive behavior captured by the DPP: the more similar two items according to K (and to L , as K derives from L through a rescaling of L 's eigenvalues), the more unlikely they are to be jointly included in \mathbf{A} .

Building upon this base DPP definition, we now explain how DPPs are integrated in our context. The goal is to produce a maximally diverse set of N future trajectories. To this end, we follow [10] and define the ground set Y as the finite set of the N predicted trajectories. Intuitively, the overall diversity defined by L over Y reflects into the expected cardinality of the associated DPP. As this expectation reads $\mathbb{E}[|\mathbf{A}|] = \text{trace}[\text{Id} - (L_Y + \text{Id})^{-1}]$, see [9], the expression in the r.h.s. can be used to define the diversity loss for the DSF. This yields:

$$\mathcal{L}_{\text{dpp}}(\hat{\mathbf{S}}_f^{(1:N)}; L) = -\text{trace}[\text{Id} - (L_Y + \text{Id})^{-1}], \quad (3)$$

where $Y = \{\hat{\mathbf{S}}_f^{(1)}, \dots, \hat{\mathbf{S}}_f^{(N)}\}$ and L is a kernel to be defined on trajectories. Given two future trajectories $\hat{\mathbf{S}}_f^{(i)}$ and $\hat{\mathbf{S}}_f^{(j)}$ predicted from a same past and present, the trajectory kernel can be simply defined as a spherical Gaussian kernel. This, however, proves insufficient to promote *directional* diversity among the generated trajectories. Hence, we also include in the kernel the angular deviation between the final points of the two trajectories:

$$L(\hat{\mathbf{S}}_f^{(i)}, \hat{\mathbf{S}}_f^{(j)}) = \exp -\alpha(\theta_{ij} + \|\hat{\mathbf{S}}_f^{(i)} - \hat{\mathbf{S}}_f^{(j)}\|_F^2), \quad (4)$$

where $\alpha > 0$ is a parameter, $\theta_{ij} \in [0, \pi]$ is the un-oriented angle between segments $(\mathbf{S}_p(T_p), \hat{\mathbf{S}}_f^{(i)}(T_f))$ and $(\mathbf{S}_p(T_p), \hat{\mathbf{S}}_f^{(j)}(T_f))$, and $\|\cdot\|_F$ denotes the Frobenius norm.

2) **Quality with a layout loss:** In order to explicitly control the quality of the forecast trajectories, we leverage the physical constraints given by the drivable area: We introduce a loss term,

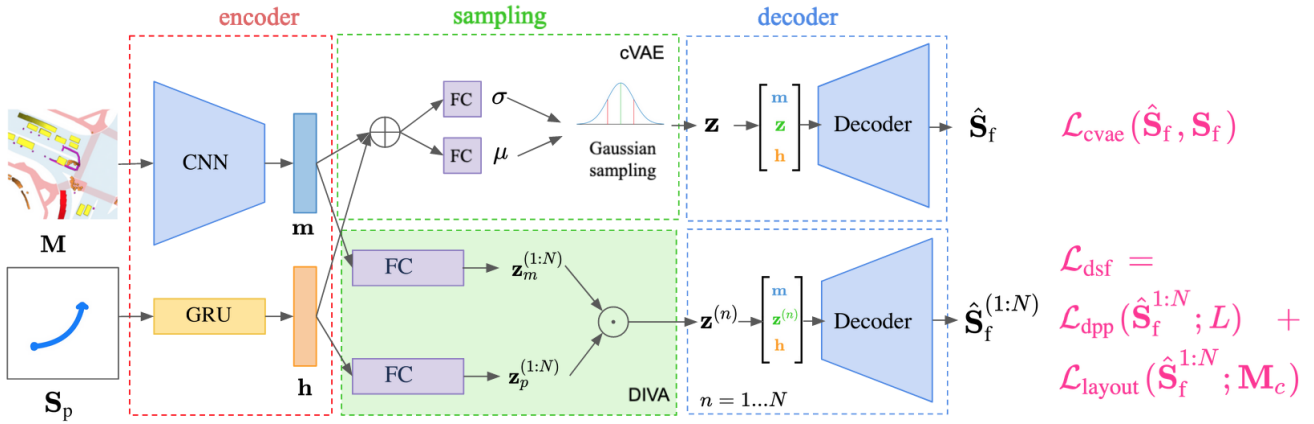


Fig. 2: **General architecture of the proposed trajectory prediction method in DIVA.** The upper part of the figure describes the underlying generative model, here a cVAE adapted to include layout information M . The lower part of the figure shows the proposed diversity sampling function that replaces the sampling part. \oplus and \odot denote concatenation and element-wise product, respectively.

$\mathcal{L}_{\text{layout}}$, to penalize trajectories predicted out of the drivable area. This binary information is part of the environment map M : we extract it and apply a Chamfer distance transform on it [35]. The resulting soft map $M_c \in [0, 1]^{H \times W}$ allows us to define a differentiable objective with respect to the coordinates of a given trajectory. Formally, given $\{\hat{S}_f^{(i)}\}_{i=1 \dots N}$ the N future trajectories predicted by our model from an input pair (S_p, M) , our layout loss is defined as:

$$\mathcal{L}_{\text{layout}}(\hat{S}_f^{(1:N)}; M_c) = \sum_{n=1}^N \sum_{t=1}^{T_f} M_c(\hat{S}_f^{(n)}(t)). \quad (5)$$

As demonstrated in [36] in the context of imitation learning, providing strong learning signals related to driving rules, such as penalizing off-road driving and collisions, does not work if the predictions are trained to match real-world driving recordings which actually do not include off-road examples. As such, when we generate trajectories that maximize their likelihood under the training data, we cannot make use of such a layout loss because it does not generate enough learning signal. Pairing the inclusion of physical constraints with a diversity-generating mechanism via a gating operation allows us to strike a good balance between diversity and quality.

IV. EXPERIMENTS

Given our proposed architecture and training scheme, we conducted experiments aimed at answering the following questions: (1) How does the addition of a DPP-based training scheme improve the overall diversity? (2) How is the quality of the generated diversity impacted by the layout loss? (3) Is the overall accuracy of the model with respect to the ground truth conserved in experiments on a real-world driving dataset?

A. Metrics

As real driving datasets cannot provide ground-truth distribution of possible futures, we need tools to measure the

diversity of the generated trajectories while ensuring that the single ground-truth future is among predictions.

Diversity and admissibility metrics. Measuring diversity is not straightforward since it is not a well-defined concept. As a consequence, several diversity metrics have been proposed. The ratio of average Final Distance Error (FDE) to min FDE, $rF = \frac{\text{avgFDE}}{\text{minFDE}}$ [3], is a measure of the spread of the proposed trajectories relative to the ground truth: A high value indicates a high avgFDE, meaning some predictions are far away from the ground truth, and a small mFDE, i.e., one of these predictions is close to the ground truth. To measure the spread of the predicted set independently from the ground truth, we use the Average Self Distance (ASD) and Final Self Distance (FSD), introduced in [10]. We also use two additional metrics to add a qualitative assessment of the diversity: the Drivable Area Occupancy (DAO) [3], which measures the diversity in predictions that are in the drivable area and the Drivable Area Count (DAC) [37] defined as $\text{DAC} = \frac{N-m}{N}$, where m is the number of predictions that exit the drivable area (DA).

As discussed in [3], these metrics are complementary: rF , ASD and FSD quantify the diversity in terms of mere spread; DAC only assesses the admissibility of the trajectories in the set; DAO captures a mix of diversity and admissibility, by measuring the spread among admissible trajectories only. DAO and DAC, being directly related to the drivable area, they provide valuable insights into how the proposed trajectories resemble real trajectories. We were not able to reproduce [3], so we do not report the FSD and ASD metrics as they were not originally evaluated in the paper. This does not impair the results as DAO and rF provide a good diversity assessment.

Ground-truth metrics. In addition to diversity metrics, we evaluate the accuracy of our method with respect to the dataset's unique ground-truth trajectory, traditionally assessed with an Euclidean distance. Following the existing trajectory forecasting literature [3], [10], [21], [36], [38], [39], we use the minimum Average Distance Error (mADE) and Final Distance Error

(mFDE), computing the error on respectively all the points of the trajectory or only the final one.

B. Experimental setup

Dataset. nuScenes [11] is a real-world driving dataset consisting of around 850 driving scenes of 20 seconds each. These scenarios were recorded in Boston and Singapore, respectively left and right-hand traffic regions, and include complex maneuvers and layouts. Annotations are very detailed and allow the use of “bird’s-eye-view” (BEV) maps containing information such as drivable area and pedestrian crossings. However, as a real-world dataset, it offers only a single ground-truth future trajectory for each past trajectory, which makes learning multiple-output prediction difficult.

DIVA setup. Our experiments are conducted with the following parameters: The loss balancing coefficient λ is set to 0.5, the latent dimension d_z to 16 and the past-embedding dimension d_h to 128. More details can be found in Supplementary A and in our associated code (<https://github.com/lcalem/DIVA>).

C. Results and discussion

Results. We compare our method to three baselines using a generative backbone: for methods using a cVAE backbone, we compare with [2] as a reference for cVAE generative models without any explicit diversification mechanism, and with [10], a diversity method also using DPPs. Originally tested on a toy dataset for trajectory prediction, we report here the results of [10] when tested on the real-world dataset nuScenes. We also include a comparison with CAM-NF [3], a recent method involving a diversification mechanism built upon a Normalizing Flow (NF) attentional backbone with the whole drivable area as an equiprobable ground-truth distribution for possible futures. As for other diversity methods, we measure diversity and admissibility for 3s predictions with 6s of past history.

Model	Backbone	mADE / mFDE ↓	DAO ↑	DAC ↑	rF ↑
DESIRE [2]	cVAE	1.079 / 1.844	16.29	0.776	1.717
L2 DPP [10]*	cVAE	1.148 / 2.272	13.31	0.975	1.891
DIVA	cVAE	0.942 / 1.449	34.99	0.972	4.907
CAM-NF [3]	NF-A	0.639 / 1.171	22.62	0.918	2.558

TABLE I: **Prediction assessment on nuScenes.** Evaluation of quality, diversity and admissibility metrics (computed on $N = 12$ predictions) for 3s forecast by our best model and cVAE-backbone baselines. We also include CAM-NF [3] for the sake of completeness, even though it has a different backbone, preventing comparisons. *: Our implementation.

Results in Table I indicate that our best model, including the DPP loss with a combined angle and Gaussian kernel, has the best performance, improving the diversity both in quantity, as measured with the spread relative to the ground truth (rF), and also in quality. For completeness, we included in our results the mADE and mFDE metrics which measure the precision of the best prediction compared to the ground truth, although the focus of this work is on diversity. These metrics depend mostly

on the generative model used during the initial training, which explains the better precision on these metrics of [3] which has a backbone relying on attention mechanisms and normalizing flows. We use a cVAE backbone and obtain results similar to those of DESIRE and L2 DPP on these metrics, as expected due to the generative backbone being the same. All diversity metrics (DAO, rF and DAC) show a marked increase compared to [3] despite the simpler backbone, showing the significance of our contribution on diversity.

Ablation study. To analyze the contributions of our architecture and losses, we perform an ablation study, the results of which can be seen in Table II. As a baseline, we start by training a very simple cVAE backbone with the loss given by Equation 1, and assess its performance on the predicted trajectories decoded from \mathbf{h} , \mathbf{m} and a \mathbf{z} component sampled from the Gaussian prior.

Model	mADE/mFDE ↓	DAO	DAC	rF	ASD	FSD↑
cVAE	1.374 / 2.682	10.91	0.975	1.246	0.120	0.165
DSF 1B D	1.152 / 2.275	13.04	0.975	1.881	0.642	0.872
DSF 1B L	1.383 / 2.723	5.81	0.977	1.046	0.048	0.058
DSF 2B D	1.018 / 1.594	35.08	0.917	4.948	2.319	3.033
DSF 2B (D+L)	0.942 / 1.449	34.99	0.972	4.907	2.142	2.842

TABLE II: **Impact of each component.** Evaluation of the contribution of each component to the quality and diversity of $N = 12$ predictions over 3s on nuScenes.

As expected, the results of the cVAE baseline are relatively mediocre on the quality metrics mADE and mFDE, although consistent with the results of DESIRE, which makes use of a “rank-and-refine” module in addition to the cVAE. Low scores of DAO, rF, ASD and FSD are also expected, since the model fails to diversify the predictions and essentially predicts stacked trajectories that go straight. This outcome also explains well the very high DAC measure, as the prediction almost exits the drivable area. By replacing the Gaussian sampling by a “weak” DSF composed of only one branch (‘DSF 1B’), improvements on diversity are seen when training the DSF with the diversity loss (‘D’) but not when training with the layout loss only (‘L’); this is expected as the layout loss does not enforce any diversity constraints. When using our two-branch DSF architecture (‘DSF 2B’) with an element-wise product to combine \mathbf{z}_m and \mathbf{z}_p , significant improvements in diversity occur. If the DSF is trained using the diversity loss only (‘D’), diversity scores are at their maximum, at the expense of the DAC metric which shows that some (8.3%) trajectories exit the drivable area as a result of the diversification. Adding the layout loss (‘D+L’) improves the quality again to levels comparable to the non-diverse baseline cVAE, at the expense of a slight drop in raw spread as indicated by the decrease in ASD and FSD (−3.67% and −5.67% respectively).

Fusion. As discussed subsection III-B, the fusion between layout and diversity encodings \mathbf{z}_m and \mathbf{z}_p is a crucial feature of our model. We compare various fusions in Table III and

highlight that the best results are obtained by the element-wise product, validating the gating hypothesis.

Model	mADE / mFDE ↓	DAO	DAC	rF	ASD	FSD ↑
concat	1.146 / 2.270	12.293	0.972	1.765	0.633	0.858
sum \oplus	1.007 / 1.833	28.175	0.935	3.083	1.474	1.881
product \odot	0.942 / 1.449	34.992	0.972	4.907	2.142	2.842

TABLE III: **Ablation of the fusion between layout and diversity encodings.** Metrics computed on $N = 12$ predictions over 3s in nuScenes, with three ways to combine \mathbf{z}_m and \mathbf{z}_p .

D. Model analysis

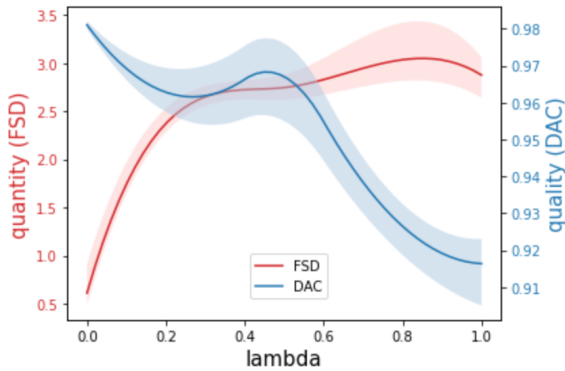


Fig. 4: **Impact of the balance between quality and quantity losses.** Influence of the weighting parameter in the loss $\mathcal{L}_{\text{dst}} = \lambda \mathcal{L}_{\text{dpp}} + (1 - \lambda) \mathcal{L}_{\text{layout}}$, measured by FSD (red, left axis) and DAC (blue, right axis). For $\lambda = 0$, diversity is suppressed and trajectories stay in the drivable area but have low spread; at $\lambda = 1$, the diversity is maximal at the expense of admissibility.

Loss balancing. In Figure 4, we show the effects of varying the balance between the $\mathcal{L}_{\text{layout}}$ and $\mathcal{L}_{\text{past}}$ loss terms in Equation 2. The axes are chosen to be FSD as a measure of diversity quantity (as it measures the spread of the proposed trajectories) and DAC as a measure of diversity quality (as it measures the percentage of proposed trajectories that stay in the drivable area), to best show the tradeoff between quantity of diversity and quality of this diversity when varying λ . At one extreme $\lambda = 0$, we suppress the DPP loss entirely, resulting in a very low diversity and a very high quality. The other extreme, $\lambda = 1$, zeroes out the layout loss (although we still have both DSF branches in the architecture), yielding, as expected, results similar to the second row of Table II.

Qualitative results. In Figure 3, we highlight the diversity improvements from our model on a variety of scenes, including intersections, straight lines and parking. Note how the basic cVAE model exhibits a mode collapse, and how the diversity of our model is influenced by both the past trajectory and the layout of the scene. Additional qualitative results are also available in Supplementary B.

V. CONCLUSION

In this paper, we introduce DIVA, a multi-output forecasting method for predicting diverse yet admissible trajectories. We use a DPP probabilistic model for diversity, and introduce a specific DPP kernel for predicting diverse driving options, leveraging the variety of settings present in the training data. The compatibility of the proposed diverse set with the drivable area is controlled by the inclusion of an admissibility loss independent from the underlying generative model. Quantitative and qualitative experiments on real-world dataset nuScenes confirm the benefit on diversity of the proposed architecture and training scheme.

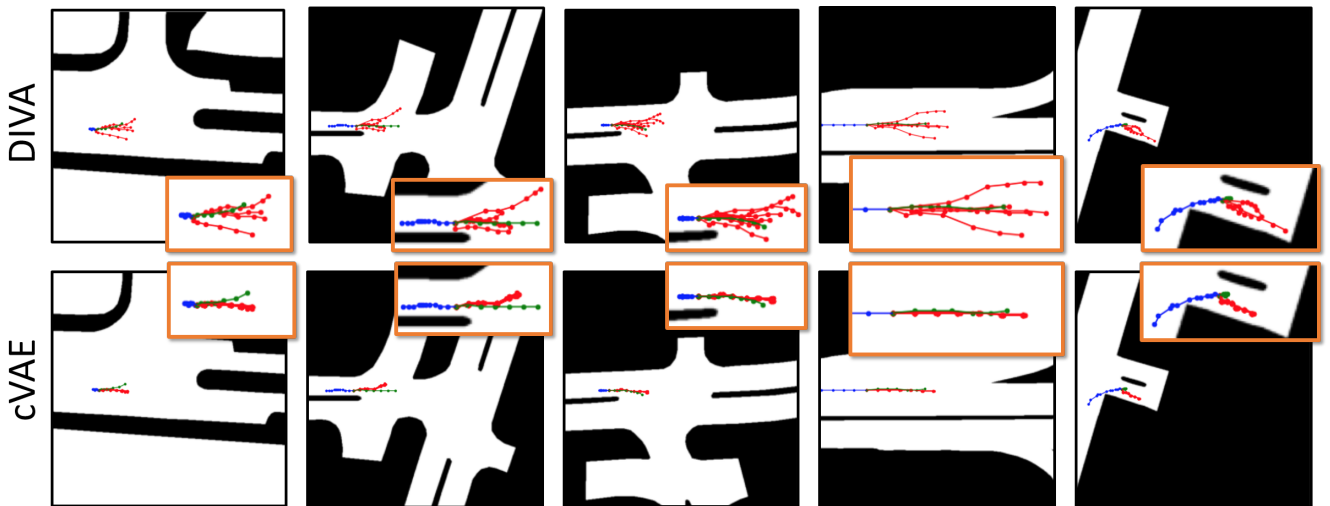


Fig. 3: **Qualitative results for various scene layouts in nuScenes.** (Top) Results from proposed DIVA model. (Bottom) Results on the same scenes with a simple cVAE, showing a focus on longitudinal diversity (speed) at the great expense of lateral diversity (direction). In each scene: past and future ground-truth trajectories are in blue and green, resp., while predicted future trajectories are in red (best viewed in color).

REFERENCES

- [1] A. Cui, S. Casas, A. Sadat, R. Liao, and R. Urtasun, "Lookout: Diverse multi-future prediction and planning for self-driving," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 16 107–16 116.
- [2] N. Lee, W. Choi, P. Vernaza, C. B. Choy, P. H. Torr, and M. Chandraker, "Desire: Distant future prediction in dynamic scenes with interacting agents," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017.
- [3] S. H. Park, G. Lee, J. Seo, M. Bhat, M. Kang, J. Francis, A. Jadhav, P. P. Liang, and L.-P. Morency, "Diverse and admissible trajectory forecasting through multimodal context understanding," in *ECCV*, 2020.
- [4] T. Salzmann, B. Ivanovic, P. Chakravarty, and M. Pavone, "Trajectron++: Multi-agent generative trajectory forecasting with heterogeneous data for control," in *European Conference on Computer Vision (ECCV)*, 2020.
- [5] Y. Chai, B. Sapp, M. Bansal, and D. Anguelov, "Multipath: Multiple probabilistic anchor trajectory hypotheses for behavior prediction," in *Conference on Robot Learning (CoRL)*, 2019.
- [6] I. J. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, "Generative adversarial networks," in *Advances in Neural Information Processing Systems (NeurIPS)*, 2014.
- [7] A. Razavi, A. van den Oord, B. Poole, and O. Vinyals, "Preventing posterior collapse with delta-VAEs," in *International Conference on Learning Representations*, 2019. [Online]. Available: <https://openreview.net/forum?id=BJe0Gn0cY7>
- [8] J. Lucas, G. Tucker, R. Grosche, and M. Norouzi, "Understanding posterior collapse in generative latent variable models," in *International Conference on Learning Representations*, 2019.
- [9] A. Kulesza and B. Taskar, "Determinantal point processes for machine learning," *Foundations and Trends in Machine Learning*, 2012.
- [10] Y. Yuan and K. Kitani, "Diverse trajectory forecasting with determinantal point processes," in *International Conference on Learning Representations (ICLR)*, 2020.
- [11] H. Caesar, V. Bankiti, A. H. Lang, S. Vora, V. E. Liong, Q. Xu, A. Krishnan, Y. Pan, G. Baldan, and O. Beijbom, "nuscnets: A multimodal dataset for autonomous driving," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020.
- [12] S. Zhang, W. Deng, Q. Zhao, H. Sun, and B. Litkouhi, "Dynamic trajectory planning for vehicle autonomous driving," in *2013 IEEE International Conference on Systems, Man, and Cybernetics*. IEEE, 2013.
- [13] T. Zhao, Y. Xu, M. Monfort, W. Choi, C. Baker, Y. Zhao, Y. Wang, and Y. N. Wu, "Multi-agent tensor fusion for contextual trajectory prediction," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [14] X. Weng, Y. Yuan, and K. Kitani, "End-to-end 3d multi-object tracking and trajectory forecasting," *arXiv preprint arXiv:2008.11598*, 2020.
- [15] A. Robicquet, A. Sadeghian, A. Alahi, and S. Savarese, "Learning social etiquette: Human trajectory understanding in crowded scenes," in *European conference on computer vision (ECCV)*, 2016.
- [16] A. Alahi, K. Goel, V. Ramanathan, A. Robicquet, L. Fei-Fei, and S. Savarese, "Social lstm: Human trajectory prediction in crowded spaces," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [17] S. Ramasinghe, K. Ranasinghe, S. Khan, N. Barnes, and S. Gould, "Conditional generative modeling via learning the latent space," in *International Conference on Learning Representations (ICLR)*, 2021.
- [18] H. Cui, V. Radosavljevic, F. Chou, T. Lin, T. Nguyen, T. Huang, J. Schneider, and N. Djuric, "Multimodal trajectory predictions for autonomous driving using deep convolutional networks," in *International Conference on Robotics and Automation, ICRA 2019, Montreal, QC, Canada, May 20-24, 2019*. IEEE, 2019, pp. 2090–2096. [Online]. Available: <https://doi.org/10.1109/ICRA.2019.8793868>
- [19] B. Kim, S. H. Park, S. Lee, E. Khoshimjonov, D. Kum, J. Kim, J. S. Kim, and J. W. Choi, "Lapred: Lane-aware prediction of multi-modal future trajectories of dynamic agents," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2021, pp. 14 636–14 645.
- [20] D. Rezende and S. Mohamed, "Variational inference with normalizing flows," in *International Conference on Machine Learning (ICML)*, 2015.
- [21] T. Phan-Minh, E. C. Grigore, F. A. Boulton, O. Beijbom, and E. M. Wolff, "Covnet: Multimodal behavior prediction using trajectory sets," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020.
- [22] O. Macchi, "The coincidence approach to stochastic point processes," *Advances in Applied Probability*, 1975.
- [23] Z. Mariet, Y. Ovdadia, and J. Snoek, "Dppnet: Approximating determinantal point processes with deep networks," in *Neural Information Processing Systems (NeurIPS)*, 2019.
- [24] J. Robinson, S. Sra, and S. Jegelka, "Flexible modeling of diversity with strongly log-concave distributions," in *Neural Information Processing Systems (NeurIPS)*, 2019.
- [25] L. E. Celis, A. Deshpande, T. Kathuria, D. Straszak, and N. K. Vishnoi, "On the complexity of constrained determinantal point processes," *arXiv preprint arXiv:1608.00554*, 2016.
- [26] B. Gong, W.-L. Chao, K. Grauman, and F. Sha, "Diverse sequential subset selection for supervised video summarization," *Advances in neural information processing systems*, 2014.
- [27] K. Hong and A. Nenkova, "Improving the estimation of word importance for news multi-document summarization," in *Proceedings of the 14th Conference of the European Chapter of the Association for Computational Linguistics*, 2014.
- [28] V. L. Guen and N. Thome, "Probabilistic time series forecasting with structured shape and temporal diversity," in *Neural Information Processing Systems (NeurIPS)*, 2020.
- [29] M. Elfeki, C. Couprie, M. Riviere, and M. Elhoseiny, "Gdpp: Learning diverse generations using determinantal point processes," in *International Conference on Machine Learning (ICML)*, 2019.
- [30] X. Weng, Y. Yuan, and K. Kitani, "Ptp: Parallelized tracking and prediction with graph neural networks and diversity sampling," in *International Conference on Robotics and Automation (ICRA)*, 2021.
- [31] W. Zeng, W. Luo, S. Suo, A. Sadat, B. Yang, S. Casas, and R. Urtasun, "End-to-end interpretable neural motion planner," in *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2019, Long Beach, CA, USA, June 16-20, 2019*. Computer Vision Foundation / IEEE, 2019, pp. 8660–8669. [Online]. Available: http://openaccess.thecvf.com/content_CVPR_2019/html/Zeng_End-To-End_Interpretable_Neural_Motion_Planner_CVPR_2019_paper.html
- [32] K. Cho, B. Van Merriënboer, C. Gulcehre, D. Bahdanau, F. Bougares, H. Schwenk, and Y. Bengio, "Learning phrase representations using rnn encoder-decoder for statistical machine translation," in *Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2014.
- [33] D. P. Kingma and M. Welling, "Auto-encoding variational bayes," in *International Conference on Learning Representations (ICLR)*, 2014.
- [34] I. Higgins, L. Matthey, A. Pal, C. Burgess, X. Glorot, M. Botvinick, S. Mohamed, and A. Lerchner, "beta-vae: Learning basic visual concepts with a constrained variational framework," in *International Conference on Learning Representations (ICLR)*, 2017.
- [35] G. Borgefors, "Distance transformations in arbitrary dimensions," *Graphical Models graphical Models and Image Processing computer Vision, Graphics, and Image Processing*, vol. 27, pp. 321–345, 1984.
- [36] M. Bansal, A. Krizhevsky, and A. Ogale, "Chauffeurnet: Learning to drive by imitating the best and synthesizing the worst," *Robotics: Science and Systems (RSS)*, 2018.
- [37] M.-F. Chang, J. Lambert, P. Sangkloy, J. Singh, S. Bak, A. Hartnett, D. Wang, P. Carr, S. Lucey, D. Ramanan *et al.*, "Argoverse: 3d tracking and forecasting with rich maps," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019.
- [38] Y. Yuan and K. Kitani, "Dlow: Diversifying latent flows for diverse human motion prediction," in *Proceedings of the European Conference on Computer Vision (ECCV)*, 2020.
- [39] N. Rhinehart, K. Kitani, and P. Vernaza, "R2p2: A reparameterized pushforward policy for diverse, precise generative path forecasting," in *European Conference on Computer Vision (ECCV)*, 2018.