
MULTIPATH++: EFFICIENT INFORMATION FUSION AND TRAJECTORY AGGREGATION FOR BEHAVIOR PREDICTION

Balakrishnan Varadarajan

Ahmed Hefny

Avikalp Srivastava

Khaled S. Refaat

Nigamaa Nayakanti

Andre Cornman

Kan Chen

Bertrand Douillard

Chi Pang Lam

Dragomir Anguelov

Benjamin Sapp*

ABSTRACT

Predicting the future behavior of road users is one of the most challenging and important problems in autonomous driving. Applying deep learning to this problem requires fusing heterogeneous world state in the form of rich perception signals and map information, and inferring highly multi-modal distributions over possible futures. In this paper, we present MultiPath++, a future prediction model that achieves state-of-the-art performance on popular benchmarks. MultiPath++ improves the MultiPath architecture [45] by revisiting many design choices. The first key design difference is a departure from dense image-based encoding of the input world state in favor of a sparse encoding of heterogeneous scene elements: MultiPath++ consumes compact and efficient polylines to describe road features, and raw agent state information directly (e.g., position, velocity, acceleration). We propose a context-aware fusion of these elements and develop a reusable *multi-context gating* fusion component. Second, we reconsider the choice of pre-defined, static anchors, and develop a way to learn latent anchor embeddings end-to-end in the model. Lastly, we explore ensembling and output aggregation techniques—common in other ML domains—and find effective variants for our probabilistic multimodal output representation. We perform an extensive ablation on these design choices, and show that our proposed model achieves state-of-the-art performance on the Argoverse Motion Forecasting Competition [12] and the Waymo Open Dataset Motion Prediction Challenge [18].

1 Introduction

Modeling and predicting the future behavior of human agents is a fundamental problem in many real-world robotics domains. For example, accurately forecasting the future state of other vehicles, cyclists and pedestrians is critical for safe, comfortable, and human-like autonomous driving. However, behavior prediction in an autonomous vehicle (AV) driving setting poses a number of unique modeling challenges:

1) *Multimodal output space*: The problem is inherently stochastic; it is impossible to truly know the future state of the environment. This is exacerbated by the fact that other agents’ intentions are not observable, and leads to a highly multimodal distribution over possible outcomes (e.g., a car could turn left or right at an intersection). Effective models must be able to represent such a rich output space with high precision and recall matching the underlying distribution.

2) *Heterogenous, interrelated input space*: The driving environment representation can contain a *highly heterogeneous* mix of static and dynamic inputs: road network information (lane geometry and connectivity, stop lines, crosswalks), traffic light state information, and motion history of agents. Driving situations are often *highly interactive*, and can involve many agents at once (e.g. negotiating a 4-way stop with crosswalks). This requires careful modeling choices, as explicitly modeling joint future distributions over multiple agents is exponential in the number of agents. Effective models must capture not only the interactions between the agents in the scene, but also the relationships between the road elements and the behavior of agents given the road context.

*Contact: {balakrishnanv, bensapp}@waymo.com. Waymo, 1600 Amphitheatre Pkwy, Mountain View, California, USA.

The novel challenges and high impact of this problem have naturally garnered much interest in recent years. There has been a rich body of work on how to model agents’ futures, their interactions, and the environment. However, there is little consensus to date on the best modeling choices for each component, and in popular benchmark challenge datasets [6, 12, 18, 53], there is a surprisingly diverse set of solutions to this problem; for details see Section 2 and Table 2.

The MultiPath framework [45] addresses the multimodal output space challenge above by modeling the highly multimodal output distributions via a Gaussian Mixture Model. It handles a common issue of *mode collapse* during learning by using static trajectory anchors, an external input to the model. This practical solution gives practitioners a straightforward way of ensuring diversity and an extra level of modeler control via the design of such anchors. The choice of a GMM representation proved to be an extremely popular, appearing in many works—see Table 2, “Trajectory Distribution”, where “Weighted set” is a special case of GMMs where only means and mixture weights are modeled.

The MultiPath input representation and backbone draws heavily upon the computer vision literature. By rasterizing all world state in a top-down orthographic view, MultiPath and others [10, 14, 27, 30, 33, 37, 46, 52] leverage powerful, established CNN architectures like ResNet [25], which offer solutions to the heterogeneous interrelated input space: the heterogeneous world state is mapped to a common pixel format, and interactions occur via local information sharing via convolution operations. While convenient and established, there are downsides to such rasterization: (1) There is an uneasy trade-off between resolution of the spatial grid, field of view, and compute requirements. (2) Rasterizing is a form of manual feature engineering, and some features may be inherently difficult to represent in such a framework (e.g. radial velocity). (3) It is difficult to capture long range interactions via convolutions with small receptive fields. (4) The information content is spatially very sparse, making a dense representation a potentially computationally wasteful choice.

In this paper, we introduce **MultiPath++**, which builds upon MultiPath, taking its output GMM representation and concept of anchors, but reconsidering how to represent and combine highly heterogeneous world state inputs and model interactions between state elements. MultiPath++ introduces a number of key upgrades:

- We eschew the rasterization-and-CNN approach in favor of modeling sparse world state objects more directly from their compact state description. We represent road elements as polylines, agent history as a sequence of physical state encoded with RNNs, and agent interactions as RNNs over the state of neighbors relative to each ego-agent. These choices avoid lossy rasterization in favor of raw, continuous state, and result in compute complexity that scales with the number of scene elements rather than the size of a spatial grid. Long-range dependencies are effectively and efficiently modeled in our representation.
- Capturing relationships between road elements and agents is critical, and we find that encoding each element independently does not perform as well as modeling their interactions (e.g. , when the road encoding is aware of relevant agents, and vice versa). To address this, we propose a novel form of context awareness we call *multi-context gating* (MCG), in which sets of elements have access to a summary context vector upon which encodings are conditioned. MCG is implemented as a generic neural network component that is applied throughout our model. MCG can be viewed as an efficient form of cross-attention, whose efficiency/quality trade-off depends on the size of the context vector.
- We also explore improvements in trajectory modeling, comparing representations based on kinematic controls, and/or polynomials as a function of continuous future time. We further demonstrate a way to learn latent representations of anchors and show they outperform the original static anchors of MultiPath, while simplifying model creation to a single-step process.
- Finally, we find significant additional gains on public benchmarks by applying ensembling techniques to our models. Unlike models with static anchors, the latent anchors of a MultiPath++ ensemble are not in direct correspondence. Furthermore, a lot of popular behavior prediction benchmarks have introduced metrics such as miss-rate (MR) and mean Average Precision (mAP), which require the ability to model diverse outcomes with few trajectories and differ from pure trajectory distance error capturing the average agent behavior. With the above in mind, we formulate the problem of ensembling the results of several models as one of greedy iterative clustering, which maximizes a probabilistic objective using the popular Expectation Maximization algorithm [3].

As of November 1, 2021, MultiPath++ ranks 1st on the Waymo Open Motion Dataset leaderboard², 4th on the Argoverse Motion Forecasting Competition³. We offer MultiPath++ as a reference set of design choices, empirically validated via ablation studies, that can be adopted, further studied and extended by the behavior modeling community.

²<https://waymo.com/open/challenges/2021/motion-prediction/>

³<https://eval.ai/web/challenges/challenge-page/454/leaderboard/1279>

Method	Road Enc.	Motion Enc.	Interactions	Decoder	Output	Trajectory Distribution
Jean [34]	–	LSTM	attention	LSTM	states	GMM
TNT [54]	polyline	polyline	maxpool, attention	MLP	states	Weighted set
LaneGCN [32]	GNN	1D conv	GNN	MLP	states	Weighted set
WIMP [28]	polyline	LSTM	GNN+attention	LSTM	states	GMM
VectorNet [20]	polyline	polyline	maxpool, attention	MLP	states	Single traj.
SceneTransformer [35]	polyline	attention	attention	attention	states	Weighted set
GOHOME [21]	GNN	1D conv + GRU	GNN	MLP	states	heatmap
MP3 [11]	raster	raster	conv	conv	cost function	Weighted samples
CoverNet [37]	raster	raster	conv	lookup	states	GMM w/ dynamic anchors
DESIRE [30]	raster	GRU	spatial pooling	GRU	states	Samples
RoadRules [27]	raster	raster	conv	LSTM	states	GMM
SocialLSTM [1]	–	LSTM	spatial pooling	LSTM	states	Samples
SocialGan [24]	–	LSTM	maxpool	LSTM	states	Samples
MFP [46]	raster	GRU	RNNs+attention	GRU	states	Samples
MANTRA [33]	raster	GRU	–	GRU	states	Samples
PRANK [2]	raster	raster	conv	lookup	states	Weighted set
IntentNet [10]	raster	raster	conv	conv	states	Single traj.
SpaGNN [9]	raster	raster	GNN	MLP	state	Single traj.
Multimodal [14]	raster	raster	conv	conv	states	Weighted set
PLOP [5]	raster	LSTM	conv	MLP	state poly	GMM
Precog [41]	raster	GRU	multi-agent sim.	GRU	motion	Samples
R2P2 [40]	raster	GRU	–	GRU	motion	Samples
HYU_ACE [36]	raster	LSTM	attn	LSTM	motion	Samples
Trajectron++[44]	raster	LSTM	RNNs+attention	GRU	controls	GMM
DKM [13]	raster	raster	conv	conv	controls	Weighted set
MultiPath [45]	raster	raster	conv	MLP	states	GMM w/ static anchors
MultiPath++	polyline	LSTM	RNNs+maxpool	MLP	control poly	GMM

Table 1: A survey of recent work in behavior prediction, categorized by choice of road encoding, motion history encoding, agent interaction encoding, trajectory decoding, intrinsic output representation, and distribution over future trajectories.

2 Related Work

We focus on architectural design choices for behavior prediction in driving environments—what representations to use to encode road information, agent motion, agent interactions, output trajectories, and output distributions. Table 2 is a summary of past work, which we go over here with additional context.

For **road encoding**, there is a dichotomy of representations. The *raster* approach encodes the world as a stack of images, from a top-down orthographic (or “bird’s-eye”) view. Rasterizing the world state has the benefit of simplicity—all the various types of input information (road configuration, agent state history, spatial relationships) are unified via rendering as a multi-channel image, enabling one to leverage powerful off-the-shelf Convolutional Neural Network (CNN) techniques. However, this one-size-fits-all approach has significant downsides: difficulty in modeling long-range interactions, constrained field of view, and difficulty in representing continuous physical states. As an alternative, the *polyline* approach describes curves (e.g., lanes, crosswalks, boundaries) as piecewise linear segments. This is a significantly more compact form due to the sparse nature of road networks. Previous works typically process a set-of-polylines description of the world in a per-agent, agent-centric coordinate system. LaneGCN [32] stands apart by treating road lanes as nodes in a graph neural network, leveraging road network connectivity structure.

To model **motion history**, one popular choice is to encode the sequence of past observed states via a recurrent net (GRU, LSTM) or temporal (1D) convolution. As an alternative, in the *raster* framework, the state sequence is typically rendered as a stack of binary mask images depicting agent oriented bounding boxes, or rendered in the same image, with the corresponding time information rendered separately [39].

To model **agent interactions**, one must deal with a dynamic set of neighboring agents around each modeled agent. This is typically done by aggregating neighbor motion history with a permutation-invariant set operator: pooling or soft attention. Notably, Precog [41] jointly rolls out agent policies in a step-wise simulation. Raster approaches rely on convolution over the 2D spatial grid to implicitly capture interactions; long-term interactions are dependent on the network receptive fields.

Agent **trajectory decoding** choices are similar to choices for encoding motion history, with the exception of methods that do lookup on a fixed or learned trajectory database [2, 37].

The most **popular output trajectory representation** is a sequence of *states* (or state differences). A few works [40, 42] instead model Newton’s laws of *motion* in a discrete time-step aggregation capturing Verlet integration. Other works [13, 44] explicitly model *controls* which parameterize a kinematically-feasible model for vehicles and bicycles. With any of these representations, the spacetime trajectory can be intrinsically represented as a sequence of sample points or a continuous polynomial representation [5]. In our experimental results, we explore the effectiveness of states and kinematic controls, with and without an underlying polynomial basis. Notably unique are (1) HOME [22] and GOHOME [21] which first predict a heatmap, and then decode trajectories after sampling, and (2) MP3 [11] and NMP [52] which learn a cost function evaluator of trajectories, and the trajectories are enumerated heuristically rather than generated by a learned model.

Nearly all work assumes an independent, per-agent output space, in which agent interactions cannot be explicitly captured. A few works are notable in describing joint interactions as output, either in an asymmetric [28, 47] or symmetric way [18, 35, 41].

The choice of output **trajectory distribution** has ramifications on downstream applications. An intrinsic property of the driving setting is that a vehicle or a pedestrian can follow one of a diverse set of possible trajectories. It is thus essential to capture the multimodal nature of the problem. Gaussian Mixture Models (*GMMs*) are a popular choice for this purpose due to their compact parameterized form; mode collapse is addressed through training tricks [14, 50] or the use of trajectory anchors [45]. Other approaches model a discrete distribution over a set of trajectories (learned or fixed a priori) [2, 14, 32, 54], or via a collection of trajectory samples drawn from a latent distribution and decoded by the model [1, 30, 33, 40, 41].

3 Model Architecture

Figure 1 depicts the proposed MultiPath++ model architecture, which on a high level is similar to that of MultiPath [45]; the model consists of an encoding step and a predictor head which conditions on anchors and outputs a Gaussian Mixture Model (GMM) [3] distribution for the possible agent position at each future time step.

MultiPath used a common, top-down image based representation for all input modalities (e.g., agents’ tracked state, road network information), and a CNN encoder. In contrast, MultiPath++ has encoders processing each input modality and converting it to a compact and sparse representation; the different modality encodings are later fused using a multi-context gating (MCG) mechanism.

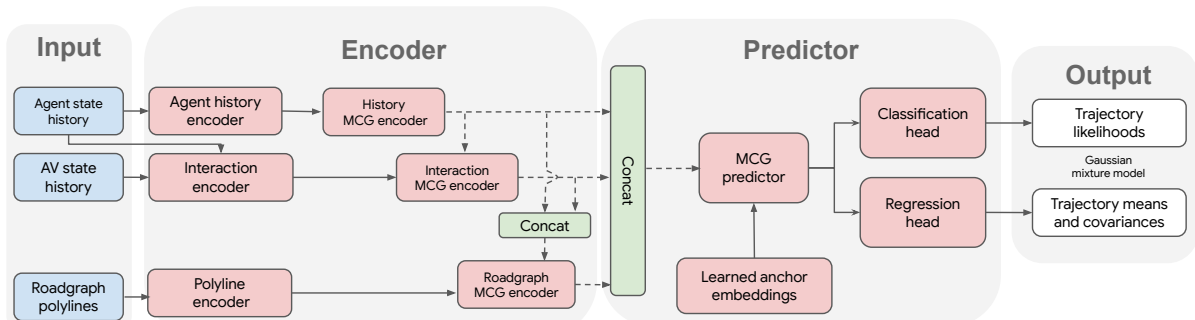


Figure 1: MultiPath++ Model Architecture. MCG denotes Multi-Context Gating, described in Section 3. Blocks in red highlight portions of the model with learned parameters. Dotted inputs to the MCG denotes context features. Each of the encoder MCG outputs aggregated embeddings (one per agent) as shown by dotted arrows. On the other hand, the predictor MCG outputs one embedding per trajectory per agent

3.1 Input Representation

MultiPath++ makes predictions based on the following input modalities:

- *Agent state history*: a state sequence describing the agent trajectory for a fixed number of past steps. In the Waymo Open Motion dataset [18], this state information includes position, velocity, 3D bounding box size, heading angle and object type; for Argoverse [12] only position information is provided. The state is

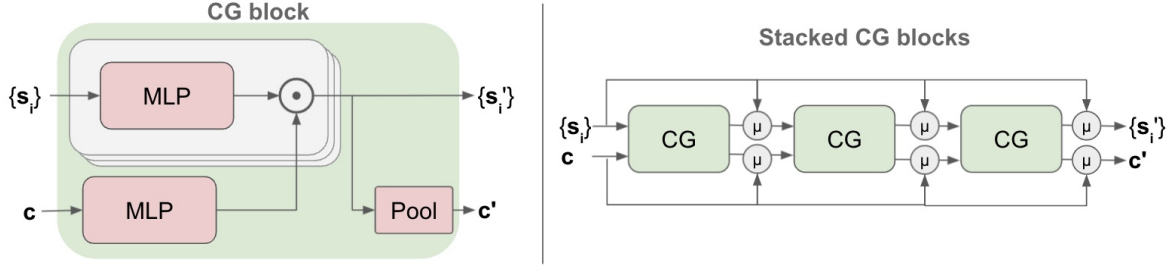


Figure 2: **Left:** Context gating (CG) block diagram. **Right:** 3 CG blocks stacked together, with running-average skip-connections (shown as components labeled “ μ ”). See Section 3.2 for details.

transformed to an agent-centric coordinate system, such that the most recent agent pose is located at the origin and heading east.⁴

- *Road network:* Road network elements such as lane lines, cross walks, and stop lines are often represented as parametric curves like clothoids [52], which can be sampled to produce point collections that are easily stored in multi-dimensional array format, as is done in many public datasets [12, 18]. We further summarize this information by approximating point sequences for each road element as a set of piecewise linear segments, or polylines, similar to [20, 26, 32].
- *Agent interactions:* For each modeled agent, we consider all neighboring agents. For each neighboring agent, we extract features in the modeled agent’s coordinate frame, such as relative orientation, distance, history and speed.
- *AV-relative features:* Similar to the interaction features, we extract features of the autonomous vehicle / sensing vehicle (AV) relative to each other agent. We model the AV separately from the other agents. We hypothesize this is a helpful distinction for the model because: (a) The AV is the center of sensors’ field of view. Tracking errors due to distance and occlusion are relative to this center. (b) The behavior of the AV can be unlike the other road users, which to a good approximation can be assumed to all be humans.

Details on how these features are encoded and fused are described next. These steps comprise the “Encoder” block of Figure 1, whose output is an encoding per agent, in each agent’s coordinate frame.

3.2 Multi Context Gating for fusing modalities

In this section we focus on how to combine the different input modality encodings in an effective way. Other works use a common rasterized format [45, 52], a simple concatenation of encodings [30, 41, 44], or employ attention [20, 32, 35, 46]. We propose an efficient mechanism for fusing information we term *multi-context gating* (MCG), and use MCG blocks throughout the MultiPath++ architecture.

Given a set of elements $s_{1:n}$ and an input context vector c , a CG block assigns an output $s'_{1:n}$ to each element in the set, and computes an output context vector c' . The output does not depend on the ordering of input elements. Mathematically, let $CG(\cdot, \cdot)$ be the function implemented by the CG block, and π be any permutation operation on a sequence of n elements. The following equations hold for CG:

$$\begin{aligned} (s'_{1:n}, c') &= CG(s_{1:n}, c) \\ (s'^*_1, c^*) &= CG(\pi(s_{1:n}), c) \end{aligned} \quad (1)$$

which imply that we have

$$\begin{aligned} c^* &= c' && \text{(permutation-invariance)} \\ s'^*_1 &= \pi(s'_{1:n}) && \text{(permutation-equivariance).} \end{aligned}$$

The size of the set n can vary across calls to $CG(\cdot, \cdot)$.

CG’s set function properties—permutation invariance/equivariance and ability to process arbitrarily sized sets—are naturally motivated by the need to encode a variable, unordered set of road network elements and agent relationships. A number of set functions have been proposed in the literature such as DeepSets [51], PointNet [38] and SetTransformers [29].

⁴Since the explicit heading is missing in Argoverse data, we use the last two time steps to get the current orientation.




Figure 3: A comparison of the *element relationship graph* for cross-attention and CG. In cross-attention, each element s_i aggregates information from $c_{1:m}$. In CG, $c_{1:m}$ summarized with a single context vector c .

A single CG block is implemented via

$$\tilde{s}_i = \text{MLP}(s_i) \quad (2)$$

$$\tilde{c} = \text{MLP}(c) \quad (3)$$

$$s'_i = \tilde{s}_i \odot \tilde{c} \quad (4)$$

$$c' = \text{Pool}(s'_{1:n}), \quad (5)$$

where \odot denotes element-wise product and Pool is a permutation-invariant pooling layer such as max or average pooling. These operations are illustrated in Figure 2. In the absence of an input context, we simply set \tilde{c} to an all-ones vector in the first context gating block. Note that both s'_i and c' depend on all inputs. It can be shown that c' is permutation-invariant w.r.t the input embeddings. It can also be shown that s'_i are permutation-equivariant.

We stack multiple CG blocks by incorporating running-average skip-connections, as is done residual networks [25]:

$$\bar{s}_{1:n}^k = \frac{1}{k} \sum_{j=1}^k s_{1:n}^j \quad (6)$$

$$\bar{c}^k = \frac{1}{k} \sum_{j=1}^k c^j \quad (7)$$

$$(s_{1:n}^{k+1}, c^{k+1}) = \text{CG}(\bar{s}_{1:n}^k, \bar{c}^k). \quad (8)$$

We denote such multi-layer CG blocks as $\text{MCG}_N(\cdot, \cdot)$ for a stack of N CG blocks.

Comparison with attention. Attention is a popular mechanism in domains such as NLP [48] and computer vision [15, 17], in which the encoding for each element of a set is updated via a combination of encodings of all other elements. For a set of size n , this intrinsically requires $O(n^2)$ operations. In models of human behavior in driving scenarios, *self* attention has been employed to update encodings for, e.g., road lanes, by attending to neighboring lanes, or to update encodings per agent based on the other agents in the scene. *Cross* attention has also been used to condition one input type (e.g. agent encodings) on another (e.g. road lanes) [20, 32, 35]. Without loss of generality, if there are n agents and m road elements, this cross attention scales as $O(nm)$ to aggregate road information for each agent.

CG can be viewed as an approximation to cross-attention. Rather than each of n elements attending to all m elements of the latter set, CG summarizes the latter set with the *single* context vector c , as shown in Figure 3. Thus the dimensionality of c needs to be great enough to capture all the useful information contained in the original m encodings. If the dimensionality of elements is d , and the dimensionality of c is d^c , then if $d^c = md$, CG can be reduced to some form of cross-attention by setting $c = \text{Concat}(\{c_j\}_{j=1}^m)$. When $d^c < md$, we are trading off the representational power of full cross-attention with computational efficiency.

3.3 Encoders

In this section we detail the specific encoders shown in Figure 1.

Agent history encoding. The agent history encoding is obtained by concatenating the output of three sources:

1. A LSTM on the history features from H time steps ago to the present time: $(\mathbf{x}_t, \mathbf{y}_t)_{t=-H:0}$.
2. A LSTM on the *difference* in the history features $(\mathbf{x}_t - \mathbf{x}_{t-1}, \mathbf{y}_t - \mathbf{y}_{t-1})_{t=-H+1:0}$.
3. MCG blocks applied to the set of history elements. Each element in the set consists of a historical position and time offset in seconds relative to the present time. The context input here is an all-ones vector with an identity

context MLP. Additionally we also encode the history frame id as a one hot vector to further disambiguate the history steps.

We denote the final embedding, which concatenates these three state history encodings, as $\phi^{(\text{state})}$.

Agent interaction encoding. For each modeled agent, we build an interaction encoding by considering each neighboring agent ν 's past state observations: $\{\mathbf{x}_{-H}^\nu, \dots, \mathbf{x}_{-1}^\nu, \mathbf{x}_0^\nu\}$. We transform ν 's state into the modeled agent's coordinate frame, and embed it with a LSTM to obtain an embedding $\phi_\nu^{(\text{interaction})}$. Note this is similar to the ego-agent history embedding but instead applied to the *relative coordinates of another agent*.

By doing this for n neighboring agents we obtain a set of interaction embeddings $\phi_{i=1:n}^{(\text{interaction})}$. We fuse neighbor information with stacked MCG blocks as follows

$$(\phi'^{(\text{state})}, \phi'_{1:n}^{(\text{interaction})}) = \text{MCG}_N(\phi_{1:n}^{(\text{interaction})}, [\phi^{(\text{state})}, \phi_{\text{AV}}^{(\text{interaction})}]) \quad (9)$$

where the second argument is the input context vector to MCG, which in this case is a concatenation of the modeled agent's history embedding, and the AV's interaction embedding. In this way we emphasize the AV's representation as a unique entity in the context for all interactions; see Section 3.1 for motivation.

Road network encoding. We use the polyline road element representation discussed in Section 3.1 as input. Each line segment is parameterized by its start point, end point and the road element semantic type Υ (e.g., *Crosswalk*, *SolidDoubleYellow*, etc). For each agent of interest, we transform the closest $P = 128$ polylines into their frame of reference and call the transformed segment $\mathbf{p} = (\mathbf{a}, \mathbf{b})$. Let \mathbf{r} be the closest point from the agent to the segment, and \mathbf{a}^\perp be the unit tangent vector at \mathbf{a} on the original curve. Then we represent the agent's spatial relationship to the segment via the vector $[\|\mathbf{r}\|_2, \mathbf{r}/\|\mathbf{r}\|_2, (\mathbf{b} - \mathbf{a})/\|\mathbf{b} - \mathbf{a}\|_2, \|\mathbf{b} - \mathbf{a}\|_2, \|\mathbf{b} - \mathbf{r}\|_2, \mathbf{a}^\perp, \text{OneHotEncoding}(\Upsilon)]$. These feature vectors are each processed with a shared MLP, resulting in a set of agent-specific embeddings per road segment, which we denote by $\phi_{1:P}^{(\text{road})}$. We then fuse road element embeddings with the agent history embedding $\phi^{(\text{state})}$ using stacked MCG blocks

$$(\phi'_{1:P}^{(\text{road})}, \phi'^{(\text{state})}) = \text{MCG}_N(\phi_{1:P}^{(\text{road})}, \phi^{(\text{state})}) \quad (10)$$

and thus enrich the road embeddings with dynamic state information.

3.4 Output representation

MultiPath++ predicts a distribution of future behavior parameterized as a Gaussian Mixture Model (GMM), as is done in MultiPath [45] and other works [4, 34, 37]. For efficient long-term prediction, the distribution is conditionally independent over time steps across mixture components, thus each mode at each time step is represented as a Gaussian distribution over (x, y) with a mean $\mu^t \in \mathbb{R}^2$ and covariance $\Sigma^t \in \mathbb{R}^{2 \times 2}$. The M mode likelihoods $p_{1:M}$ are tied over time. MAP inference per mode is equivalent to taking the sequence $\mu^{1:T}$ as state waypoints defining a possible future trajectory for the agent. The full output distribution is

$$p(s) = \sum_{i=1}^M p_i \prod_{t=1}^T \mathcal{N}(s_t - \mu_i^t, \Sigma_i^t) \quad (11)$$

where $s = s^{1:T}$ represents a trajectory; $s^t \in \mathbb{R}^2$.

The classification head of Figure 1 predicts the p_i as a softmax distribution over mixture components. The regression head outputs the parameters of the Gaussians μ and Σ for M modes and T time steps.

Training objective. We follow the original MultiPath approach and maximize the likelihood of the groundtruth trajectory under our model's predicted distribution. We make a hard-assignment labeling of a "correct" mixture component by choosing the one with the smallest Euclidean distance to the groundtruth trajectory.

The average log loss over the entire training set is optimized using Adam. We use an initial learning rate of 0.0002 and a batch size of 64, with decay rate of 0.5 every 2 epochs. The final model is chosen after training for 800,000 steps.

3.5 Prediction architecture with learned anchor embeddings

The goal of the Predictor module (Figure 1) is to predict the parameters of the GMM described in Section 3.4, namely M trajectories, with likelihoods and uncertainties around each waypoint.

In applications related to future prediction, capturing the highly uncertain and multimodal set of outcomes is a key challenge and the focus of much work [31, 37, 40, 45, 50]. One of MultiPath's key innovations was to use a static set of

anchor trajectories as pre-defined modes that applied to all scenes. One major downside to this is that most modes are not a good fit to any particular scene, thus requiring a large amount modes to be considered, with most obtaining a low-likelihood and getting discarded. Another downside is the added complexity and effort stemming from a 2-phase learning process (first estimating the modes from data, then training the network).

In this work, we learn anchor embeddings as part of the overall model training. We interpret these embeddings $\mathbf{e}_{1:M}$ as anchors in latent space, and construct our architecture to have a one-to-one correspondence with these embeddings and the output trajectory modes of our GMM. The vectors $\mathbf{e}_{1:M}$ are trainable model parameters that are independent of the input. This has connections to Detection Transformers (DETR) [8] which propose a way to learn anchors rather than hand-design them for object detection. This is also similar in spirit to MANTRA [33], a trajectory prediction network, which has an explicit learned memory network which consists of a database of embeddings that can be retrieved and decoded into trajectories.

We concatenate the embeddings $\phi^{(\text{interaction})}$, $\phi^{(\text{road})}$ and $\phi^{(\text{state})}$ obtained from the output of the respective MCG blocks to obtain a fixed-length feature vector $\phi^{(\text{combined})}$ for each modeled agent. We then use this as context in stacked MCG blocks that operate on the set of anchor embeddings $\mathbf{e}_{1:M}$, with a final MLP that predicts all parameters of the output GMM:

$$(\mu, \sigma_x, \sigma_y, \rho_{xy}, \log q)_{1:M} = \text{MLP}(\text{MCG}(\mathbf{e}_{1:M}, \phi^{(\text{combined})})),$$

where Σ is formed from $(\sigma_x, \sigma_y, \rho_{xy})$.

3.6 Internal Trajectory Representation

We model the future position and heading of agents, along with agent-relative longitudinal and lateral Gaussian uncertainties. We parameterize the trajectory using $x(t), y(t), \theta(t), \sigma_{\text{lng}}(t), \sigma_{\text{lat}}(t)$ —position, heading, and standard deviation for longitudinal and lateral uncertainty.

The most popular approach in the literature is to directly predict a sequence of such states at uniform time-discretization. Here we also consider two non-mutually exclusive variants.

1. We can represent functions over time as polynomials, which add an inductive bias that ensures a smooth trajectory. It gives us a compact, interpretable representation of each predicted signal.
2. Instead of directly predicting $\{x(t), y(t), \theta(t)\}$, we can predict the underlying kinematic control signals, which can then be integrated to evaluate the output state. In this work, we experiment with predicting the acceleration $a(t)$ and heading change rate $\dot{\theta}(t)$ and integrating them to recover the trajectory as follows:

$$\begin{aligned} v(t) &= v(0) + \int_0^t a(\tau) d\tau \\ \theta(t) &= \theta(0) + \int_0^t \dot{\theta}(\tau) d\tau \\ x(t) &= x(0) + \int_0^t v(\tau) \cos(\theta(\tau)) d\tau \\ y(t) &= y(0) + \int_0^t v(\tau) \sin(\theta(\tau)) d\tau. \end{aligned}$$

These representations add inductive bias encouraging natural and realistic trajectories that are based on realistic kinematics and consistent with the current state of the predicted agent. For the polynomial representation, it is also possible to specify a *soft constraint* by regularizing the polynomial’s constant term, which determines the shift of the predicted signal from its current value.

Algorithm 1 demonstrates the conversion from control signals to output positions. Note that this operation is differentiable, permitting end-to-end optimization. It is a numerical approximation of Equation 12 with additional technical considerations: (1) When computing the next position $(x(t), y(t))$, we use the midpoint approximation of the speed and heading $\tilde{\theta}(t) \equiv \theta(t - \delta_t) + \dot{\theta}(t - \delta_t) \cdot (\delta_t/2)$. (2) Given vehicle dimensions, we cap the heading change rate to match a predetermined maximum feasible curvature. (3) These equations are applied to the rear-axle of the vehicle rather than the center position. We use the rear-end position of the vehicle as an approximation of the rear-axle position.

Note that Algorithm 1 can be viewed as a special type of recurrent network, without learned parameters. This decoding stage then mirrors other works which use a learned RNN (LSTM or GRU cells) to decode an embedding vector into a trajectory [27, 28, 34, 44, 46]. In our case, the recurrent network state consists of $x(t), y(t), v(t)$ and $\theta(t)$, and the input consists of $\dot{\theta}(t)$ and $a(t)$. Encoding an inductive bias derived from kinematic modeling spares the network the need to explicitly learn these properties makes the predicted state compact. This promotes data efficiency and generalization power, but can be more sensitive to perception errors in the current state estimate.

Data: $a(t)$ and $\dot{\theta}(t)$ at $t = 0, 0.1, 0.2, \dots, 10.0$, $x(0), y(0), \theta(0), v(0)$

Result: $x(t), y(t)$ and $\theta(t)$ at $t = 0.1, 0.2, \dots, 10.0$

$\delta_t = 0.1$

for $t = 0.1, 0.2, \dots, 10.0$ **do**

$\tilde{v} \leftarrow v(t - \delta_t) + a(t) \cdot (\delta_t/2)$
 $\dot{\theta}_{\text{cap}} \leftarrow \text{CapCurvature}(\dot{\theta}(t - \delta_t))$
 $\tilde{\theta} \leftarrow \theta(t - \delta_t) + \dot{\theta}_{\text{cap}} \cdot (\delta_t/2)$
 $x(t) \leftarrow x(t - \delta_t) + \tilde{v} \cos(\tilde{\theta})\delta_t$
 $y(t) \leftarrow y(t - \delta_t) + \tilde{v} \sin(\tilde{\theta})\delta_t$
 $\theta(t) \leftarrow \theta(t - 1) + \dot{\theta}_{\text{cap}}\delta_t$
 $v(t) \leftarrow \theta(t - 1) + a(t - \delta_t)\delta_t$

end

Algorithm 1: Integrating control-signal to positions.

4 Ensembling predictor heads via bootstrap aggregation

Ensembling is a powerful and popular technique in many machine learning applications. For example, ensembling is a critical technique for getting the best performance on ImageNet [25]. By combining multiple models which are to some degree complementary, we can enjoy the benefits of a higher capacity model with lower statistical variance.

We specifically apply bootstrap aggregation (bagging) [19] to our predictor heads by training E such heads together. To encourage models learning complementary information, the weights of the E heads are initialized randomly, and an example is used to update the weights of each head with a 50% probability.

Unlike scalar regression or classification, it is not obvious how to combine output from different heads in our case—each is a Gaussian Mixture Model, with no correspondence of mixture components across ensemble heads. Furthermore, we consider allowing each predictor head to predict a richer output distribution with more modes $L > M$; where M is fixed as a requirement for the task (and is used in benchmark metrics calculations).

Let Ψ denote the union of the predictions from all heads

$$\Psi = \{(\mu_1, \Sigma_1, q_1), \dots, (\mu_{M'}, \Sigma_{M'}, q_{M'})\}, \quad (12)$$

where $M' = L \cdot E$, and the mode likelihoods q are divided by the number of heads E so that they sum up to 1. Then we pose the ensemble combination task as one of converting Ψ to a more compact GMM $\bar{\Psi}$ with M modes:

$$\bar{\Psi} = ((\bar{\mu}_1, \bar{\Sigma}_1, \bar{q}_1), \dots, (\bar{\mu}_M, \bar{\Sigma}_M, \bar{q}_M)), \quad (13)$$

while requiring that $\bar{\Psi}$ best approximates Ψ . In this section we describe the aggregation algorithm we use. Theoretical motivations and derivation can be found in Appendix A.

We find fit $\bar{\Psi}$ to Ψ using an iterative clustering algorithm, like Expectation-Maximization [3], but with hard assignment of cluster membership. This setting lends itself to efficient implementation in a compute graph, and allows us to train this step end-to-end as a final layer in our deep network.

We start by selecting M cluster centroids from $\mu_{1:M'}$ in a greedy fashion. The selection criteria is to maximize the probability that a centroid sampled from Ψ lies within τ distance from at least one selected centroid:

$$\bar{\mu}_{1:M} = \operatorname{argmax}_{\bar{\mu}_{1:M}} \sum_{i=1}^{M'} q_i \max_{\bar{\mu} \in \bar{\mu}_{1:M}} \mathbb{I}(\|\mu_i - \bar{\mu}\|_2 \leq \tau) \quad (14)$$

This is a criterion that explicitly optimizes trajectory diversity, which is a fit for metrics such as miss rate, mAP and minADE, as defined in [12, 18]. Other criteria could also be used depending on the metric of interest. It is interesting to relate this criteria to the ensembling and sampling method employed by GOHOME [21]. In that work, they output an intermediate spatial heatmap representation, which is amenable to ensemble aggregation. Then they greedily sample end-points in a similar fashion.

Since jointly optimizing $\bar{\mu}_{1:M}$ is hard, we select each μ_i greedily for $i = 1, \dots, M$ according to

$$\bar{\mu}_i = \operatorname{argmax}_{\bar{\mu}_i} \sum_{i=1}^{M'} q_i \max_{\bar{\mu} \in \bar{\mu}_{1:i}} \mathbb{I}(\|\mu_i - \bar{\mu}\|_2 \leq \tau) \quad (15)$$

which differs in that the outer argmax is done iteratively over $\bar{\mu}_i$ rather than jointly $\bar{\mu}_{1:M}$.

Starting with the selected centroids, We iteratively update the parameters of $\bar{\Psi}$ using an expectation-maximization-style [16] algorithm, where each iteration consists of the following updates

$$\bar{q}_h \leftarrow \sum_{i=1}^{M'} q_i p(h|\mu_i; \bar{\Psi}) \quad (16)$$

$$\bar{\mu}_h \leftarrow \frac{1}{\bar{q}_h} \sum_{i=1}^{M'} q_i p(h|\mu_i; \bar{\Psi}) \mathbf{x} \quad (17)$$

$$\bar{\Sigma}_h \leftarrow \frac{1}{\bar{q}_h} \sum_{i=1}^{M'} q_i p(h|\mu_i; \bar{\Psi}) [\Sigma_i + (\mu_i - \bar{\mu}_h)(\mu_i - \bar{\mu}_h)^T], \quad (18)$$

where $p(h|\mathbf{x}; \bar{\Psi})$ is the posterior probability that a given sample \mathbf{x} is sampled from the h^{th} component of the mixture model specified by $\bar{\Psi}$, which can be computed as

$$p(h|\mathbf{x}; \bar{\Psi}) = \frac{\bar{q}_h \mathcal{N}(\mathbf{x} - \bar{\mu}_h, \bar{\Sigma}_h)}{\sum_{k=1}^M \bar{q}_k \mathcal{N}(\mathbf{x} - \bar{\mu}_k, \bar{\Sigma}_k)} \quad (19)$$

5 Experiments

5.1 Datasets

The Waymo Open Motion Dataset (WOMD) [18] consists of 1.1M examples time-windowed from 103K 20s scenarios. The dataset is derived from real-world driving in urban and suburban environments. Each example for training and inference consists of 1 second of history state and 8 seconds of future, which we resample at 5Hz. The object-agent state contains attributes such as position, agent dimensions, velocity and acceleration vectors, orientation, angular velocity, and turn signal state. The long (8s) time horizon in this dataset tests the model’s ability to capture a large field of view and scale to an output space of trajectories, which in theory grows exponentially with time.

The Argoverse dataset [12] consists of 333K scenarios containing trajectory histories, context agents, and lane centerline inputs for motion prediction. The trajectories are sampled at 10Hz, with 2 seconds of past history and a 3-second future prediction horizon.

5.2 Metrics

We compare models using competition specific metrics associated with various datasets⁵,

Specifically, we report the following metrics.

minDE_k^t (Minimum Distance Error): The minimum distance, over the top k most-likely trajectories, between a predicted trajectory and the ground truth trajectory at time t .

minADE_k (Minimum Average Distance Error): Similar to minDE_k^t, but the distance is calculated as an average over all timesteps.

MR_k^t@d (Miss Rate): Measures the rate at which minFDE_k^t exceeds d meters. Note that WOMD leaderboard uses a different definition [18].

mAP: For each set of predicted trajectories, we have at most one positive - the one closest to the ground truth and which is within τ distance from the ground truth. The other predicted trajectories are reported as misses. From this, we can compute precision and recall at various thresholds. Following WOMD metrics definition [18] the agents future trajectories are partitioned into behavior buckets, and an area under the precision-recall curve is computed using the possible true positive and false positives per agent, giving us Average Precision per behavior bucket. The total mAP value is a mean over the AP’s for each behavior bucket.

Overlap rate: The fraction of times the most likely trajectory prediction of any agent overlaps with a real future trajectory of another agent (see [18] for details).

⁵For each dataset, we report the results of our model against published results of publicly available models.

Argoverse leaderboard ($k = 6, d = 2m, t = 3s$)					
	Rank [†]	brier-minDE	minFDE	MR	minADE
LaneGCN [32]	50	2.059	1.364	0.163	0.868
DenseTNT [23]	23	1.976	1.282	0.126	0.882
HOME + GOHOME [21]	10	1.860	1.292	0.085	0.890
TPCN++ [49]	5	1.796	1.168	0.116	0.780
MultiPath++ (ours)	4	1.793	1.214	0.132	0.790
QCraft Blue Team	1	1.757	1.214	0.114	0.801

Table 2: Comparison with select state-of-the-art methods on Argoverse leaderboard. k is the number of trajectories. d is the maximum distance for no miss, and t is the trajectory time duration. [†] Rank on the public leaderboard <https://eval.ai/web/challenges/challenge-page/454/leaderboard/1279> as of November 12, 2021, which is sorted by brier-minDE.

TRI: (Turning Radius Infeasibility) We compute the turning radius along the predicted trajectories using two approaches: one that uses the predicted yaw output from the model (**TRI-h**), and the other that doesn’t require yaw predictions and instead uses the circumradius constituting three consecutive waypoints (**TRI-c**). If the radius is less than a certain threshold τ , it is treated as a violation. We set this threshold as the approximate minimum turning radius threshold for a midsize sedan, $\tau = 3.5m$. Note that a model that simply predicts a constant heading can achieve a TRI-h rate of zero, hence we also compute inconsistencies between turning radius suggested by the coordinates and the predicted headings (**TRI-hc**). TRI-hc inconsistency is true when the difference in heading based on circumradius from waypoints and predicted headings is greater than 0.05 radians at any time step in a trajectory.

5.3 MultiPath baseline

As our work evolved from MultiPath, we include a reference MultiPath model where the input and backbone are faithful to the original paper [45] for a point of comparison, with a few minor differences. Specifically, we use a top-down rendering of the scene as before, but now employ a *splat rendering* [55] approach for rasterization, in which we sample points uniformly from scene elements and do an orthographic projection. This is a simpler, sparse form of rendering, which doesn’t employ anti-aliasing, but is efficient and straightforward to implement in TensorFlow and run as part of the model compute graph on hardware accelerators (GPU/TPU).

As in the original paper, we use a grid of 400×400 cells, with grid cell physical dimension of $0.2m \times 0.2m$, thus a total field-of-view of $80m$ centered around the AV sensing vehicle in WOMD, with a ResNet18 backbone [25]. We use 128 static anchors obtained via k-means, which are shared among all agent types (vehicles, pedestrians, cyclists) for simplicity. Figure 10 illustrates this model’s inputs and architecture.

5.4 External benchmark results

On Argoverse, MultiPath++ achieves top-5 performance on most metrics (Table 2). Our technique is ranked 1st on all metrics on Waymo Open Motion Dataset [18] (Table 3).

The tested model is based on the best configuration in Table 4, where the outputs from multiple ensemble heads are aggregated as described in Section 4.

On WOMD, we also see that the original MultiPath model, even with the refinement of learned anchors and ensembling, is outperformed by more recent methods. It is interesting to note that MultiPath is the best performing top-down scene-centric model employing a CNN; every known method which outranks it uses sparse representations.

5.5 Qualitative Examples

Figure 4 shows examples of Multipath++ on WOMD scenes. Figure 5 shows examples of Multipath++ on Argoverse scenes. These examples show the ability of MultiPath++ to handle different road layouts and agent interactions.

5.6 Ablation Study

In this section we evaluate our design choices through an ablation study. Table 4 summarizes ablation results. In the following subsections we discuss how our architecture choices affect the model performance.

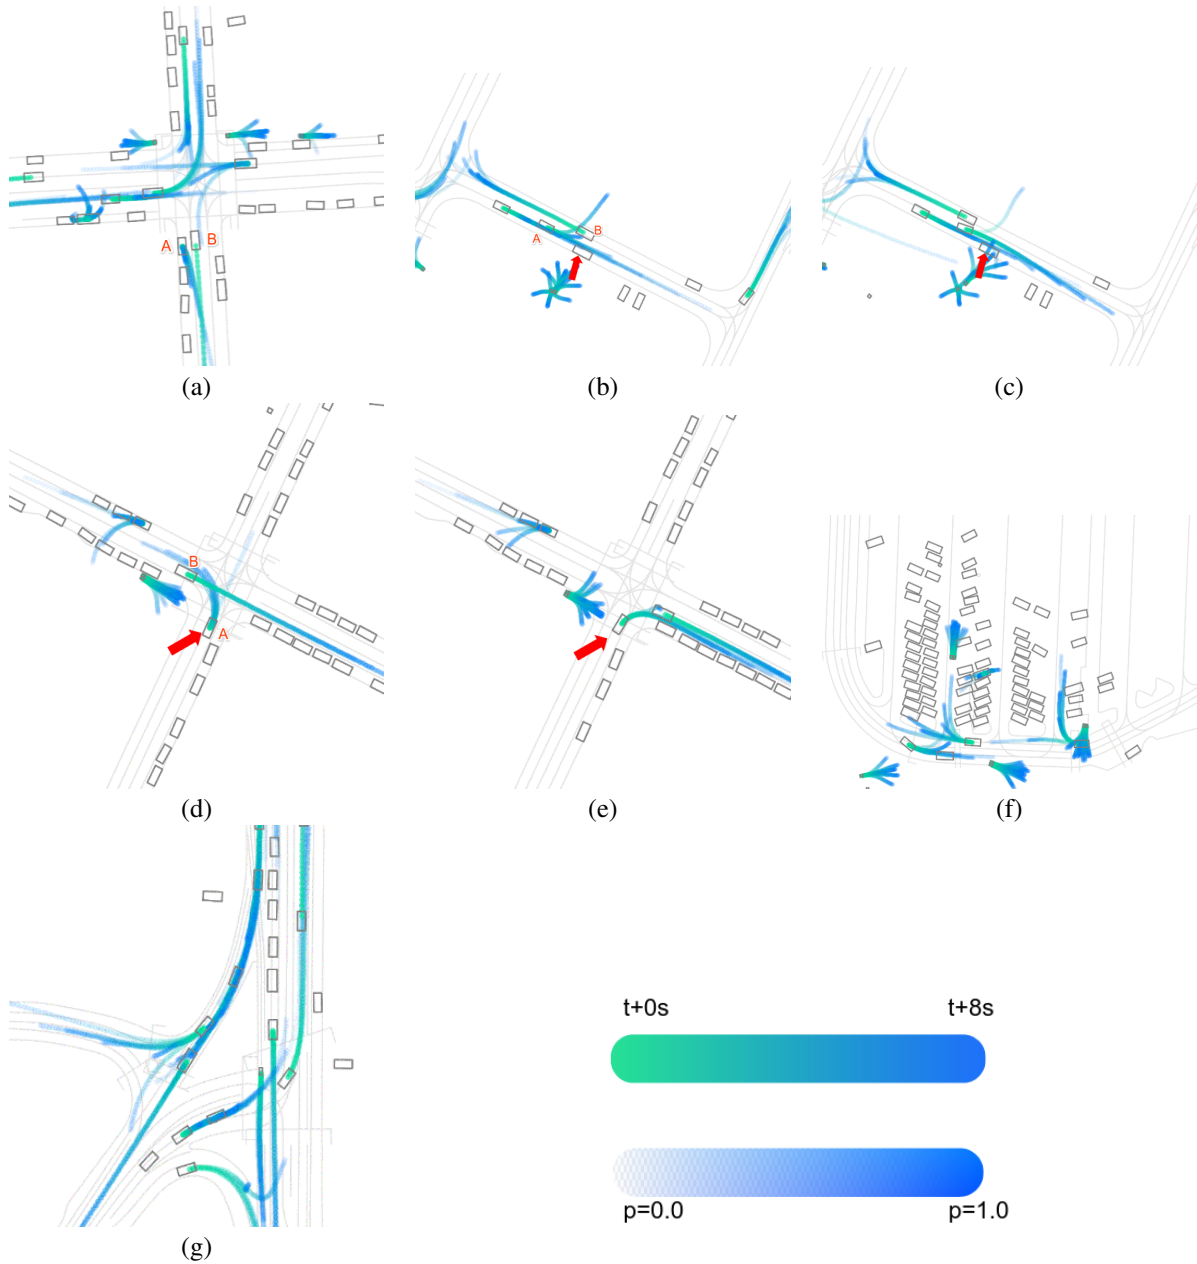


Figure 4: Examples of MultiPath++ predictions for 8 seconds in WOMD scenes. Hue indicates time horizon while transparency indicates predicted probability. Rectangles indicate vehicles while small squares indicate pedestrians. (a): A four-way intersection involving multiple interactions. For example, car A is predicted to yield for car B. (b & c): Narrow road interaction. Car A is predicted to yield for car B and then nudge around the parking car, designated by the arrow. (d & e): Interaction between two vehicles at an intersection where car A is predicted to yield for car B or make a right-turn behind. After car B passes, car A can make a left turn. Also, note the bimodal prediction of the pedestrian that is located at the corner. (f and g): Predictions in a parking lot and atypical roadgraph.

Waymo Open Motion Prediction ($k = 6, t = 8s$)						
	Rank [†]	minDE	minADE	MR	Overlap	mAP
MultiPath [45]	11	2.04	0.880	0.345	0.166	0.409
SceneTransformer [35]	7	1.212	0.612	0.156	0.147	0.279
DenseTNT [23]	5	1.551	1.039	0.157	0.178	0.328
MultiPath++	1	1.158	0.556	0.134	0.131	0.409

Table 3: Comparison with published state-of-the-art methods on WOMD public leaderboard. k is the number of trajectories and t is the trajectory time horizon.
[†] Rank on the public leaderboard <https://waymo.com/open/challenges/2021/motion-prediction/> as of November 12, 2021, which is sorted by mAP.




Figure 5: Examples of MultiPath++ predictions for 8 seconds in Argoverse scenes, showing the ability to follow different lane geometries.

5.6.1 Set Functions

Recall that MultiPath++ uses two types of set functions. Invariant set functions are used to encode a set of elements (e.g. agents, roadgraph segments) into a single feature vector. Equivariant set functions are used to convert the set of learned anchors, together with the encoded feature vector as a context, into a corresponding set of trajectories with likelihoods.

We use multi-context gating to represent both types of functions. We experimented with other representations of set functions:

- **MLP+MaxPool:** In this experiment, we replace the multi-context gating (MCG) road network encoder with a MLP+MaxPool applied on points rather than polylines, inspired by PointNet [38]. We use a 5 layer deep MLP and RELU activations.
- **Equivariant DeepSet [51]:** The equivariant set function is represented as a series of blocks, each involving an element-wise transformation followed by pooling to compute the context. Unlike MCG, it does not use gating (pointwise multiplication) between set elements and the context vector. Instead, a linear transformation of the context is added to each element. We use a DeepSet of 5 blocks in the predictor.
- **Transformers [29]:** We replace the gating mechanism (element-wise multiplication) on polylines with self-attention. For decoding, we used cross attention where the queries are the learned embeddings and the keys are the various encoder features.

5.6.2 Trajectory representation

As mentioned in Section 3.6, we experiment with predicting polynomial coefficients for the trajectory, as well predicting kinematic control signals (acceleration and heading change rate). We found that polynomial representations hurt performance, counter to conclusions made in PLOP [4], where they demonstrated improvements over the then state of the art on PRECOG[43] and nuScenes[7] using polynomials to represent output trajectories. Furthermore, in the PLOP datasets, we need to predict 4s into the future which is much shorter than our prediction horizon of 10s. For such short futures, polynomial representations are more suitable. In our case, we do not see much gains from using the polynomial representation, possibly due to the larger dataset size and longer-term prediction horizon.

The controls-based output works better in distance metrics than a polynomial representations, which suggests it is a more beneficial and domain-specific form of inductive bias. Overall, our results suggest that the simple sequence of raw coordinates trajectory representation works best for distance-based metrics. However, these unconstrained representations have a non-trivial rate of kinematic infeasibility (TRI-x metrics in Table 4). Kinematic feasibility and consistency between headings and positions is crucial in practice when such behavior models are used for planning and controls of a real-world robot, an issue that is not captured by public benchmark metrics.

5.6.3 Ensembling

We explore ensembling, producing an over-complete set of trajectories that is then summarized using the aggregation proposed in Section 4, as well as their combination. The number of ensembles is denoted by E and the number of trajectories per ensemble is denoted by L . Finally we aggregate the $E \cdot L$ trajectories to $M = 6$ which is the required number of trajectories for the WOMD submission.

5.6.4 Anchor representation

We explore learned and kmeans based anchor representation.

5.7 Discussion

First, we remark that MultiPath++ is a significant improvement over its predecessor MultiPath, as seen in Tables 3 and 4. As discussed in this paper, they differ in many design dimensions, the primary being the change from a dense top-down raster representation to a sparse, element-based representation with agent-centric coordinate systems. Other design choices are validated in isolation in the following discussion.

We find that MLP+MaxPool performs the worst among all set function variants as expected due to limited capacity. DeepSet is able to outperform MLP+MaxPool. Also increasing the depth of the MCG gives consistently better results owing to effective increase in capacity and flow of information across skip connections. We get the best performance by increasing the depth of the MCG to 5 layers.

We find that learning anchors (“Learned anchors”) is more effective than using a set of anchors obtained a priori via k-means. This runs counter to the original findings in the MultiPath paper [45] that anchor-free models suffer from mode

	minDE	minADE	MR	AUC	TRI-h (%)	TRI-c (%)	TRI-hc (%)
Original MultiPath	4.752	1.796	0.749	–	–	–	–
Set Function							
MLP+MaxPool	2.693	1.107	0.528	0.367	–	–	–
DeepSet	2.562	1.055	0.5	0.368	–	–	–
Transformer	2.479	1.042	0.479	0.3687	–	–	–
1 MCG block	2.764	1.15	0.55	0.312	–	–	–
5 stacked MCG blocks*	2.305	0.978	0.44	0.393	–	–	–
Trajectory Representation							
Polynomial	2.537	1.041	0.501	0.368	n/a	1.92	n/a
Control	2.319	0.987	0.449	0.386	0.00	1.22	0.00
Raw coordinates	2.305	0.978	0.44	0.393	n/a	1.08	n/a
Raw coordinates w/ heading*	2.311	0.978	0.443	0.395	4.10	1.04	9.92
Ensembling							
$E = 1, L = 6$	2.333	0.982	0.410	0.240	–	–	–
$E = 5, L = 6$	2.18	0.948	0.395	0.297	–	–	–
$E = 1, L = 64$	2.487	1.057	0.473	0.367	–	–	–
$E = 5, L = 64$ *	2.305	0.978	0.44	0.393	–	–	–
Anchors							
Static k-means anchors	2.99	1.22	0.578	0.324	–	–	–
Learned anchors*	2.305	0.978	0.44	0.393	–	–	–

Table 4: Model ablation on WOMB validation set. Metrics are parameterized by $k = 6$, $t = 8s$, and $d = 2m$.

* denotes the reference configuration: road encoding, state history encoding and interaction encoding as described in Section 3. “n/a” denotes a model that does not predict heading.

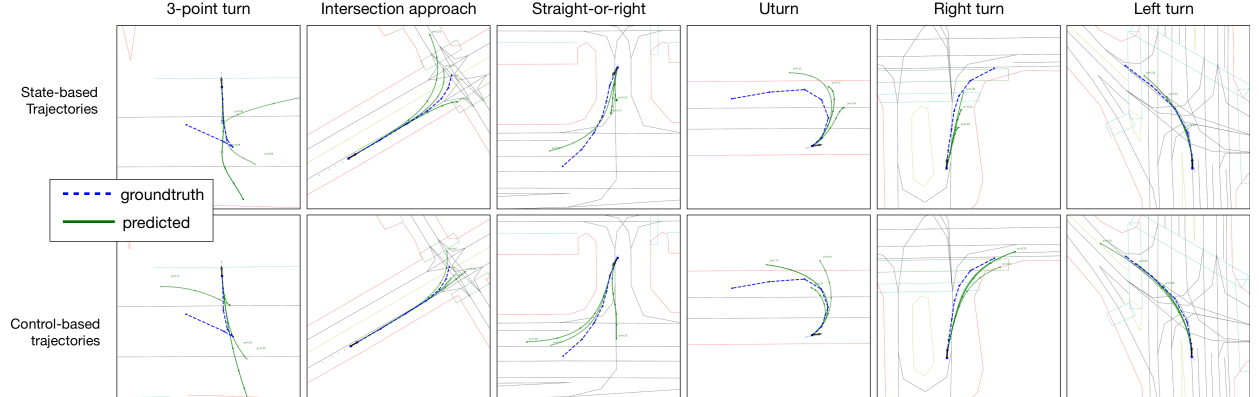


Figure 6: Qualitative results showing the effect of control-based trajectory representation. Top row: MultiPath++ with state-based trajectories. Bottom row: MultiPath++ with control-based trajectories.

collapse. The difference could possibly be due to the richer and more structured inputs, improved model architecture, and larger batch sizes in MultiPath++. We leave more detailed ablations on this issue between the two approaches to future work. We compare the baseline of directly outputting a single head with 6 trajectories ($E = 1, L = 6$), to training 5 ensemble heads ($E = 5, L = 6$). We see that ensembling significantly improves most metrics, and particularly minDE, for which this combination is best. We also train a model with a single head that outputs 64 trajectories, followed by our aggregation method that reduces them to 6 ($E = 1, L = 64$). Compared to our initial baseline, this model significantly improves MR and AUC that require diverse predictions, but regresses the average trajectory distance metrics $minDE$, and even $minADE$ a little bit. This suggests that the different metrics pose different solution requirements. As expected, our aggregation criterion is well suited to preserving diversity, while straight-up ensembling is better at capturing the average distribution. Finally, our experiment ($E = 5, L = 64$) with more ensemble heads and more predictions per ensemble combines the strengths of both techniques, obtaining a strictly superior performance in all metrics compared to the baseline.

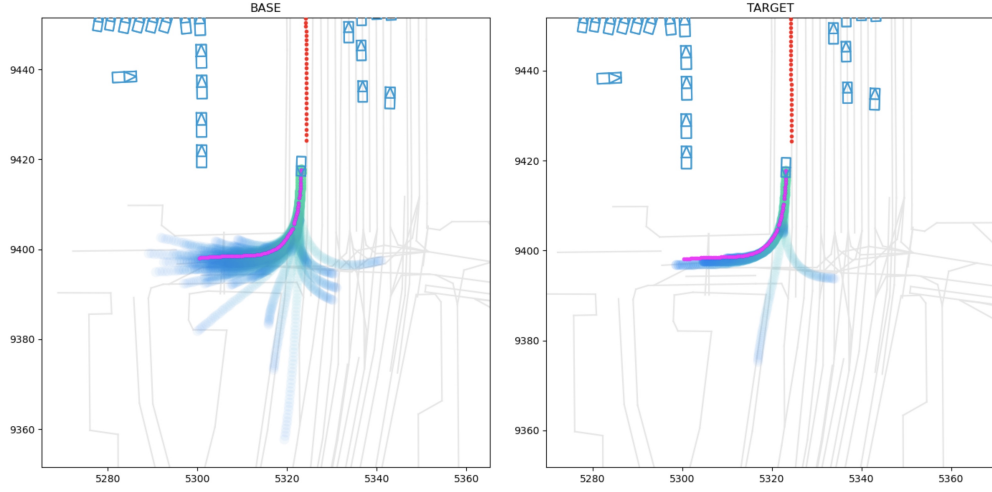


Figure 7: Example of improved diversity due to aggregation.




Figure 8: Illustration of weeding out unrealistic trajectories after aggregation.

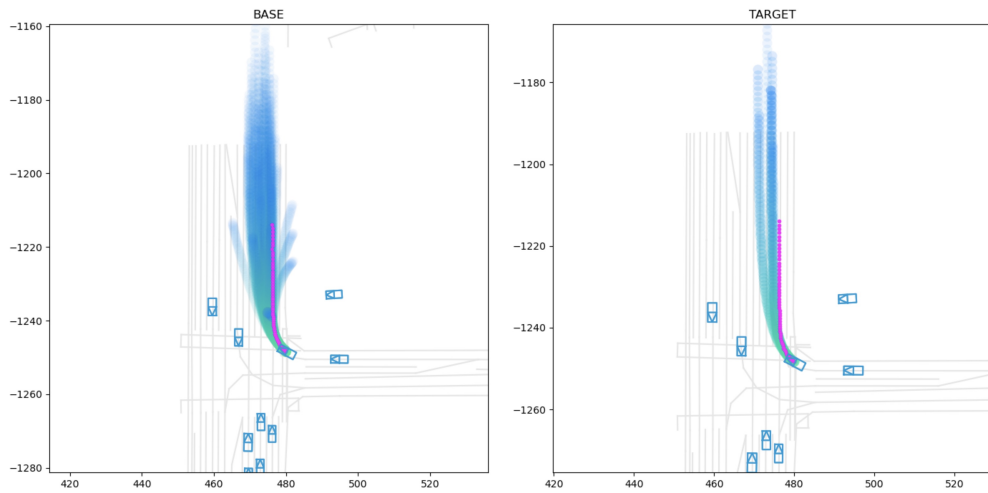


Figure 9: Illustration of improved lane diversity: The model is able to predict the agent going to multiple lanes after aggregation.

5.8 Conclusion

We proposed a novel behavior prediction system, MultiPath++, by carefully considering choices for input representation and encoding, fusing encodings, and representing the output distribution. We demonstrated state-of-the-art results on popular benchmarks for behavior prediction. Furthermore, we surveyed existing methods, analyzed our approach empirically, and provided practical insights for the research community. In particular, we showed the importance of sparse encoding, efficient fusion methods, control-based methods, and learned anchors. Finally, we provided a practical guide for various tricks used for training and inference to improve robustness, increase diversity, handle missing data, and ensure fast convergence during training.

References

- [1] Alexandre "Alahi, Kratarth Goel, Vignesh Ramanathan, Alexandre Robicquet, Li Fei-Fei, and Silvio Savarese. Social LSTM: Human Trajectory Prediction in Crowded Spaces. In *CVPR*, 2016.
- [2] Yuriy Biktairov, Maxim Stebelev, Irina Rudenko, Oleh Shliashko, and Boris Yangel. Prank: motion prediction based on ranking. *arXiv preprint arXiv:2010.12007*, 2020.
- [3] Christopher M Bishop. *Pattern recognition and machine learning*. springer, 2006.
- [4] Thibault Buhet, É. Wirbel, and Xavier Perrotton. Plop: Probabilistic polynomial objects trajectory planning for autonomous driving. *ArXiv*, abs/2003.08744, 2020.
- [5] Thibault Buhet, Emilie Wirbel, and Xavier Perrotton. Plop: Probabilistic polynomial objects trajectory planning for autonomous driving. *arXiv preprint arXiv:2003.08744*, 2020.
- [6] Holger Caesar, Varun Bankiti, Alex H. Lang, Sourabh Vora, Venice Erin Liong, Qiang Xu, Anush Krishnan, Yu Pan, Giancarlo Baldan, and Oscar Beijbom. nuscenes: A multimodal dataset for autonomous driving. *arXiv preprint arXiv:1903.11027*, 2019.
- [7] Holger Caesar, Varun Bankiti, Alex H. Lang, Sourabh Vora, Venice Erin Liong, Qiang Xu, Anush Krishnan, Yu Pan, Giancarlo Baldan, and Oscar Beijbom. nuscenes: A multimodal dataset for autonomous driving, 2020.
- [8] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end object detection with transformers. In Andrea Vedaldi, Horst Bischof, Thomas Brox, and Jan-Michael Frahm, editors, *Computer Vision – ECCV 2020*, pages 213–229, Cham, 2020. Springer International Publishing.
- [9] Sergio Casas, Cole Gulino, Renjie Liao, and Raquel Urtasun. Spaggn: Spatially-aware graph neural networks for relational behavior forecasting from sensor data. In *IEEE Intl. Conf. on Robotics and Automation*. IEEE, 2020.
- [10] Sergio Casas, Wenjie Luo, and Raquel Urtasun. Intentnet: Learning to predict intention from raw sensor data. In *Conf. on Robot Learning*, 2018.
- [11] Sergio Casas, Abbas Sadat, and Raquel Urtasun. Mp3: A unified model to map, perceive, predict and plan. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14403–14412, 2021.
- [12] Ming-Fang Chang, John Lambert, Patsorn Sangkloy, Jagjeet Singh, Slawomir Bak, Andrew Hartnett, De Wang, Peter Carr, Simon Lucey, Deva Ramanan, et al. Argoverse: 3d tracking and forecasting with rich maps. In *CVPR*, 2019.
- [13] Henggang Cui, Thi Nguyen, Fang-Chieh Chou, Tsung-Han Lin, Jeff Schneider, David Bradley, and Nemanja Djuric. Deep kinematic models for kinematically feasible vehicle trajectory predictions. In *IEEE Intl. Conf. on Robotics and Automation*, pages 10563–10569. IEEE, 2020.
- [14] Henggang Cui, Vladan Radosavljevic, Fang-Chieh Chou, Tsung-Han Lin, Thi Nguyen, Tzu-Kuo Huang, Jeff Schneider, and Nemanja Djuric. Multimodal trajectory predictions for autonomous driving using deep convolutional networks. In *IEEE Intl. Conf. on Robotics and Automation*, 2019.
- [15] Zihang Dai, Hanxiao Liu, Quoc V Le, and Mingxing Tan. Coatnet: Marrying convolution and attention for all data sizes. *arXiv preprint arXiv:2106.04803*, 2021.
- [16] A. P. Dempster, N. M. Laird, and D. B. Rubin. Maximum likelihood from incomplete data via the em algorithm. *Journal of the Royal Statistical Society. Series B (Methodological)*, 39(1):1–38, 1977.
- [17] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. In *International Conference on Learning Representations*, 2020.
- [18] Scott Ettinger, Shuyang Cheng, Benjamin Caine, Chenxi Liu, Hang Zhao, Sabeek Pradhan, Yuning Chai, Ben Sapp, Charles Qi, Yin Zhou, Zoey Yang, Aurelien Chouard, Pei Sun, Jiquan Ngiam, Vijay Vasudevan, Alexander McCauley, Jonathon Shlens, and Dragomir Anguelov. Large scale interactive motion forecasting for autonomous driving : The waymo open motion dataset, 2021.
- [19] Jerome H Friedman. *The elements of statistical learning: Data mining, inference, and prediction*. springer open, 2017.
- [20] Jiyang Gao, Chen Sun, Hang Zhao, Yi Shen, Dragomir Anguelov, Congcong Li, and Cordelia Schmid. VectorNet: Encoding hd maps and agent dynamics from vectorized representation. In *CVPR*, 2020.

- [21] Thomas Gilles, Stefano Sabatini, Dzmitry Tsishkou, Bogdan Stanciulescu, and Fabien Moutarde. Gohome: Graph-oriented heatmap output for future motion estimation. *arXiv preprint arXiv:2109.01827*, 2021.
- [22] Thomas Gilles, Stefano Sabatini, Dzmitry Tsishkou, Bogdan Stanciulescu, and Fabien Moutarde. HOME: heatmap output for future motion estimation. *CoRR*, abs/2105.10968, 2021.
- [23] Junru Gu, Qiao Sun, and Hang Zhao. Densetnt: Waymo open dataset motion prediction challenge 1st place solution. *CoRR*, abs/2106.14160, 2021.
- [24] Agrim Gupta, Justin Johnson, Li Fei-Fei, Silvio Savarese, and Alexandre Alahi. Social GAN: Socially acceptable trajectories with generative adversarial networks. In *CVPR*, 2018.
- [25] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016.
- [26] Namdar Homayounfar, Wei-Chiu Ma, Shrinidhi Kowshika Lakshmikanth, and Raquel Urtasun. Hierarchical recurrent attention networks for structured online maps. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3417–3426, 2018.
- [27] Joey Hong, Benjamin Sapp, and James Philbin. Rules of the road: Predicting driving behavior with a convolutional model of semantic interactions. In *CVPR*, 2019.
- [28] Siddhesh Khandelwal, William Qi, Jagjeet Singh, Andrew Hartnett, and Deva Ramanan. What-if motion prediction for autonomous driving. *ArXiv*, 2020.
- [29] Juho Lee, Yoonho Lee, Jungtaek Kim, Adam R. Kosiosek, Seungjin Choi, and Yee Whye Teh. Set transformer: A framework for attention-based permutation-invariant neural networks. In Kamalika Chaudhuri and Ruslan Salakhutdinov, editors, *Proceedings of the 36th International Conference on Machine Learning, ICML 2019, 9-15 June 2019, Long Beach, California, USA*, volume 97 of *Proceedings of Machine Learning Research*, pages 3744–3753. PMLR, 2019.
- [30] Namhoon Lee, Wongun Choi, Paul Vernaza, Christopher B Choy, Philip H S Torr, and Manmohan Chandraker. DESIRE: Distant future prediction in dynamic scenes with interacting agents. In *CVPR*, 2017.
- [31] Junwei Liang, Lu Jiang, Kevin Murphy, Ting Yu, and Alexander Hauptmann. The garden of forking paths: Towards multi-future trajectory prediction. In *CVPR*, 2020.
- [32] Ming Liang, Bin Yang, Rui Hu, Yun Chen, Renjie Liao, Song Feng, and Raquel Urtasun. Learning lane graph representations for motion forecasting. *arXiv preprint arXiv:2007.13732*, 2020.
- [33] Francesco Marchetti, Federico Becattini, Lorenzo Seidenari, and Alberto Del Bimbo. Mantra: Memory augmented networks for multiple trajectory prediction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7143–7152, 2020.
- [34] Jean Mercat, Thomas Gilles, Nicole Zoghby, Guillaume Sandou, Dominique Beauvois, and Guillermo Gil. Multi-head attention for joint multi-modal vehicle motion forecasting. In *IEEE Intl. Conf. on Robotics and Automation*, 2020.
- [35] Jiquan Ngiam, Benjamin Caine, Vijay Vasudevan, Zhengdong Zhang, Hao-Tien Lewis Chiang, Jeffrey Ling, Rebecca Roelofs, Alex Bewley, Chenxi Liu, Ashish Venugopal, David Weiss, Benjamin Sapp, Zhifeng Chen, and Jonathon Shlens. Scene transformer: A unified multi-task model for behavior prediction and planning. *CoRR*, abs/2106.08417, 2021.
- [36] Seong Hyeon Park, Gyubok Lee, Manoj Bhat, Jimin Seo, Minseok Kang, Jonathan Francis, Ashwin R Jadhav, Paul Pu Liang, and Louis-Philippe Morency. Diverse and admissible trajectory forecasting through multimodal context understanding. *arXiv preprint arXiv:2003.03212*, 2020.
- [37] Tung Phan-Minh, Elena Corina Grigore, Freddy A Boulton, Oscar Beijbom, and Eric M Wolff. CoverNet: Multimodal behavior prediction using trajectory sets. *arXiv:1911.10298*, 2019.
- [38] Charles R Qi, Hao Su, Kaichun Mo, and Leonidas J Guibas. Pointnet: Deep learning on point sets for 3d classification and segmentation. In *CVPR*, 2017.
- [39] Khaled S. Refaat, Kai Ding, Natalia Ponomareva, and Stéphane Ross. Agent prioritization for autonomous navigation. *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 2060–2067, 2019.
- [40] Nicholas Rhinehart, Kris Kitani, and Paul Vernaza. R2P2: A reparameterized pushforward policy for diverse, precise generative path forecasting. In *ECCV*, 2018.
- [41] Nicholas Rhinehart, Rowan McAllister, Kris Kitani, and Sergey Levine. PRECOG: Prediction conditioned on goals in visual multi-agent settings. In *Intl. Conf. on Computer Vision*, 2019.
- [42] Nicholas Rhinehart, Rowan McAllister, Kris Kitani, and Sergey Levine. Precog: Prediction conditioned on goals in visual multi-agent settings. In *ECCV*, 2019.
- [43] Nicholas Rhinehart, Rowan McAllister, Kris M. Kitani, and Sergey Levine. PRECOG: prediction conditioned on goals in visual multi-agent settings. *CoRR*, abs/1905.01296, 2019.
- [44] Tim Salzmann, Boris Ivanovic, Punarjay Chakravarty, and Marco Pavone. Trajectron++: Multi-agent generative trajectory forecasting with heterogeneous data for control. *arXiv preprint arXiv:2001.03093*, 2020.
- [45] Benjamin Sapp, Yuning Chai, Mayank Bansal, and Dragomir Anguelov. MultiPATH: Multiple probabilistic anchor trajectory hypotheses for behavior prediction. In *Conf. on Robot Learning*, 2019.

- [46] Yichuan Charlie Tang and Ruslan Salakhutdinov. Multiple futures prediction. In *Proceedings of the 33rd International Conference on Neural Information Processing Systems*, Red Hook, NY, USA, 2019. Curran Associates Inc.
- [47] Ekaterina I. Tolstaya, Reza Mahjourian, Carlton Downey, Balakrishnan Varadarajan, Benjamin Sapp, and Dragomir Anguelov. Identifying driver interactions via conditional behavior prediction. In *IEEE International Conference on Robotics and Automation, ICRA 2021, Xi'an, China, May 30 - June 5, 2021*, pages 3473–3479. IEEE, 2021.
- [48] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *NeurIPS*, 2017.
- [49] Maosheng Ye, Tongyi Cao, and Qifeng Chen. Tpcn: Temporal point cloud networks for motion forecasting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11318–11327, 2021.
- [50] Ye Yuan and Kris Kitani. Diverse trajectory forecasting with determinantal point processes. *ICLR*, 2020.
- [51] Manzil Zaheer, Satwik Kottur, Siamak Ravanbakhsh, Barnabas Poczos, Russ R Salakhutdinov, and Alexander J Smola. Deep sets. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017.
- [52] Wenyuan Zeng, Wenjie Luo, Simon Suo, Abbas Sadat, Bin Yang, Sergio Casas, and Raquel Urtasun. End-to-end interpretable neural motion planner. In *CVPR*, 2019.
- [53] Wei Zhan, Liting Sun, Di Wang, Haojie Shi, Aubrey Clausse, Maximilian Naumann, Julius Kümmerle, Hendrik Königshof, Christoph Stiller, Arnaud de La Fortelle, and Masayoshi Tomizuka. INTERACTION Dataset: An INTERnational, Adversarial and Cooperative moTION Dataset in Interactive Driving Scenarios with Semantic Maps. *arXiv:1910.03088*, 2019.
- [54] Hang Zhao, Jiyang Gao, Tian Lan, Chen Sun, Benjamin Sapp, Balakrishnan Varadarajan, Yue Shen, Yi Shen, Yuning Chai, Cordelia Schmid, et al. Tnt: Target-driven trajectory prediction. *arXiv preprint arXiv:2008.08294*, 2020.
- [55] Matthias Zwicker, Hanspeter Pfister, Jeroen Van Baar, and Markus Gross. Surface splatting. In *Proceedings of the 28th annual conference on Computer graphics and interactive techniques*, pages 371–378, 2001.

A Details and Derivation of Aggregation Algorithm

By having an overcomplete trajectory representation that is later aggregated into a fixed small number of trajectories, we attempt to address two kinds of uncertainties in the data:

- *Aleatoric uncertainty*: This is a natural variation in the data. For example an agent can either take a left or right turn or change lanes, etc given the same context information. This level of ambiguity cannot be resolved by increasing the model capacity, but rather the model needs to predict calibrated probabilities for these outcomes. Despite the theoretical possibility of modeling these variations using a small number of output trajectories directly, there are several challenges in learning. Some examples include mode collapse and failure to model these variations due to limited model capacity. Training the model to produce an overcomplete representation forces the model to output a diverse distribution of trajectories and could make it more resistant to mode collapse. Following this up with greedy iterative trajectory aggregation enhances diversity in the final output.
- *Epistemic uncertainty*: This is the variation across model outputs, which typically indicates the model’s failure to capture certain aspects of the scene or input features. Such variations could occur if some models are poorly trained or haven’t seen a particular slice of the data. By doing model ensembling, we attempt to reduce this uncertainty.

For ease of exposition, we assume each trajectory to be composed of a single time point; the same computations are applied to each time step in a future sequence. The output Ψ is a Gaussian mixture model (GMM) distribution with M' modes on the future position:

$$\mathbf{p}(x, y; \Psi) = \sum_{j=1}^{M'} q_j \mathcal{N}(x, y; \mu_j, \Sigma_j) \quad (20)$$

We formulate the aggregation as obtaining an M -mode GMM $\bar{\Psi}$ which minimizes the KL-divergence $D_{\text{KL}}(\Psi || \bar{\Psi})$. This is equivalent to maximizing the expected log likelihood $\mathbf{p}(x, y; \bar{\Psi})$ of a sample point (x, y) drawn from the overcomplete distribution Ψ :

$$f(\bar{\Psi}) = \int_{\mathbf{x}} p(\mathbf{x}; \Psi) \log p(\mathbf{x}; \bar{\Psi}) d\mathbf{x} \quad (21)$$

Assuming the overcomplete distribution approximates the real distribution, this is roughly equivalent to fitting the compact distribution to the real data, but with the added benefits described above. Directly maximizing (21) is intractable.

Hence we attempt to employ an Expectation-Maximization-like algorithm to obtain a local maximum. The difference in the objective function between an old and new value may be written as

$$f(\bar{\Psi}') - f(\bar{\Psi}) = \int_{\mathbf{x}} p(\mathbf{x}; \bar{\Psi}) \log \left[\frac{p(\mathbf{x}; \bar{\Psi}')}{p(\mathbf{x}; \bar{\Psi})} \right] d\mathbf{x} \quad (22)$$

Denoting the hidden variable h to be a mixture in the compact representation, we may write:

$$\log \left[\frac{p(\mathbf{x}; \bar{\Psi}')}{p(\mathbf{x}; \bar{\Psi})} \right] \quad (23)$$

$$= \sum_h p(h|\mathbf{x}; \bar{\Psi}) \log \left[\frac{p(\mathbf{x}; \bar{\Psi}')}{p(\mathbf{x}; \bar{\Psi})} \right] \quad (24)$$

$$= \sum_h p(h|\mathbf{x}; \bar{\Psi}) \log \left[\frac{p(h, \mathbf{x}; \bar{\Psi}') p(h|\mathbf{x}; \bar{\Psi})}{p(h|\mathbf{x}; \bar{\Psi}') p(h, \mathbf{x}; \bar{\Psi})} \right] \quad (25)$$

$$= \sum_h p(h|\mathbf{x}; \bar{\Psi}) \log \left[\frac{p(h, \mathbf{x}; \bar{\Psi}')}{p(h, \mathbf{x}; \bar{\Psi})} \right] + \sum_h p(h|\mathbf{x}; \bar{\Psi}) \log \left[\frac{p(h|\mathbf{x}; \bar{\Psi})}{p(h|\mathbf{x}; \bar{\Psi}')} \right] \quad (26)$$

$$= \sum_h p(h|\mathbf{x}; \bar{\Psi}) \log \left[\frac{p(h, \mathbf{x}; \bar{\Psi}')}{p(h, \mathbf{x}; \bar{\Psi})} \right] + D_{\text{KL}}(p(h|\mathbf{x}; \bar{\Psi}) \| p(h|\mathbf{x}; \bar{\Psi}')) \quad (27)$$

$$\geq \sum_h p(h|\mathbf{x}; \bar{\Psi}) \log \left[\frac{p(h, \mathbf{x}; \bar{\Psi}')}{p(h, \mathbf{x}; \bar{\Psi})} \right] \quad (28)$$

Thus

$$f(\bar{\Psi}') - f(\bar{\Psi}) \geq \int_{\mathbf{x}} p(\mathbf{x}; \bar{\Psi}) \sum_h p(h|\mathbf{x}; \bar{\Psi}) \log \left[\frac{p(h, \mathbf{x}; \bar{\Psi}')}{p(h, \mathbf{x}; \bar{\Psi})} \right] d\mathbf{x} \quad (29)$$

The right hand side is called the Q function in the EM algorithm. Maximizing the Q function with respect to $\bar{\Psi}'$ ensures that the likelihood increases at least as much when we update the parameters to $\bar{\Psi}'$. Noting that $p(h, \mathbf{x}; \bar{\Psi}') = p(h; \bar{\Psi}') p(\mathbf{x}; \bar{\Psi}'|h) = \bar{q}'_h \mathcal{N}(\mathbf{x} - \bar{\mu}'_h, \bar{\Sigma}'_h)$ and factoring out the terms independent of $\bar{\Psi}'$, we find the update that maximizes the lower bound to be

$$\bar{\Psi} \leftarrow \operatorname{argmax}_{\bar{\Psi}'} \int_{\mathbf{x}} p(\mathbf{x}; \bar{\Psi}) \sum_h p(h|\mathbf{x}; \bar{\Psi}) \log [\bar{q}'_h \mathcal{N}(\mathbf{x} - \bar{\mu}'_h, \bar{\Sigma}'_h)] d\mathbf{x} \quad (30)$$

$$\bar{\Psi} \leftarrow \operatorname{argmax}_{\bar{\Psi}'} \sum_{i=1}^{M'} q_i \int_{\mathbf{x}} p(\mathbf{x}; \mu_i, \Sigma_i) \sum_h p(h|\mathbf{x}; \bar{\Psi}) \log [\bar{q}'_h \mathcal{N}(\mathbf{x} - \bar{\mu}'_h, \bar{\Sigma}'_h)] d\mathbf{x} \quad (31)$$

The second equation follows from the fact that the overcomplete distribution Ψ is a mixture of M' Gaussians. The updates can be solved as follows.

$$\begin{aligned} \bar{q}'_h &\leftarrow \sum_{i=1}^{M'} q_i \int_{\mathbf{x}} p(\mathbf{x}; \mu_i, \Sigma_i) p(h|\mathbf{x}; \bar{\Psi}) d\mathbf{x} \\ \bar{\mu}'_h &\leftarrow \frac{1}{\bar{q}'_h} \sum_{i=1}^{M'} q_i \int_{\mathbf{x}} \mathbf{x} p(\mathbf{x}; \mu_i, \Sigma_i) p(h|\mathbf{x}; \bar{\Psi}) d\mathbf{x} \\ \bar{\Sigma}'_h &\leftarrow \frac{1}{\bar{q}'_h} \sum_{i=1}^{M'} q_i \int_{\mathbf{x}} \mathbf{x} p(\mathbf{x}; \mu_i, \Sigma_i) (\mathbf{x} - \bar{\mu}'_h)(\mathbf{x} - \bar{\mu}'_h)^T p(h|\mathbf{x}; \bar{\Psi}) d\mathbf{x} \end{aligned}$$

where $p(h|\mathbf{x}; \bar{\Psi})$ is the posterior probability that a given sample \mathbf{x} is sampled from the h^{th} component of the mixture model specified by $\bar{\Psi}$ (here we use the previous estimate for $\bar{\Psi}$). This can be computed as:

$$p(h|\mathbf{x}; \bar{\Psi}) = \frac{\bar{q}_h \mathcal{N}(\mathbf{x}; \bar{\mu}_h, \bar{\Sigma}_h)}{\sum_{k=1}^M \bar{q}_k \mathcal{N}(\mathbf{x}; \bar{\mu}_k, \bar{\Sigma}_k)} \quad (32)$$

Notice the resemblance with standard GMM, except where $p(\mathbf{x}; \mu_i, \Sigma_i)$ is a dirac delta function in the standard setting (since the input data in standard GMM is a set of points instead of a distribution). Unlike standard GMM, these

expectations (integrations) in the above EM updates are hard to compute in closed form. Instead we employ the approximation for any function $g(\mathbf{x})$

$$\begin{aligned} & \mathbb{E}_{\mathbf{x} \sim \mathcal{N}(x; \mu, \Sigma)} [g(\mathbf{x}) p(h|\mathbf{x}; \bar{\Psi})] \\ & \approx \mathbb{E}_{\mathbf{x} \sim \mathcal{N}(x; \mu, \Sigma)} [g(\mathbf{x}) p(h|\mu; \bar{\Psi})] \\ & = p(h|\mu; \bar{\Psi}) \mathbb{E}_{\mathbf{x} \sim \mathcal{N}(x; \mu, \Sigma)} [g(\mathbf{x})]. \end{aligned}$$

In other words, we assume that the posterior probability of any output cluster only depends on the mean of the overcomplete cluster centroid inside the expectation. This approximation is reasonable since most samples drawn from the distribution would be concentrated around the mean. Furthermore as we increase the number of cluster centroids in the overcomplete representation, the variance within each overcomplete cluster centroid becomes smaller yielding more focus around the mean. The set of updates can now be solved in closed form as follows:

$$\begin{aligned} \bar{q}_h & \leftarrow \sum_{i=1}^{M'} q_i p(h|\mu_i; \bar{\Psi}) \\ \bar{\mu}_h & \leftarrow \frac{1}{\bar{q}_h} \sum_{i=1}^{M'} q_i p(h|\mu_i; \bar{\Psi}) \mathbf{x} \\ \bar{\Sigma}_h & \leftarrow \frac{1}{\bar{q}_h} \sum_{i=1}^{M'} q_i p(h|\mu_i; \bar{\Psi}) [\Sigma_i + (\mu_i - \bar{\mu}_h)(\mu_i - \bar{\mu}_h)^T] \end{aligned}$$

Since EM is a local optimization method, careful initialization of GMM parameters is important. Our initialization criterion of GMM centroids is to maximize the probability that future point lies within τ distance from at least one centroid:

$$\bar{\mu}_{1:M} = \operatorname{argmax}_{\bar{\mu}_{1:M}} \sum_{i=1}^{M'} q_i \max_{\bar{\mu} \in \bar{\mu}_{1:M}} \mathbb{I}(\|\mu_i - \bar{\mu}\|_2 \leq \tau) \quad (33)$$

Unfortunately, directly optimizing (33) is NP-hard. So instead, we select an M -sized subset of $\mu_{1:M'}$ in a greedy fashion to maximize (33)⁶.

⁶Note that this subset selection problem is *submodular*, which means that a greedy method is guaranteed to achieve at least $1 - 1/e$ of the optimal subset value.

B Multipath baseline system diagram

Figure 10 visualizes the splat rendering rasterization of input points, and the backbone and output architecture. See Section 5.3 for details.

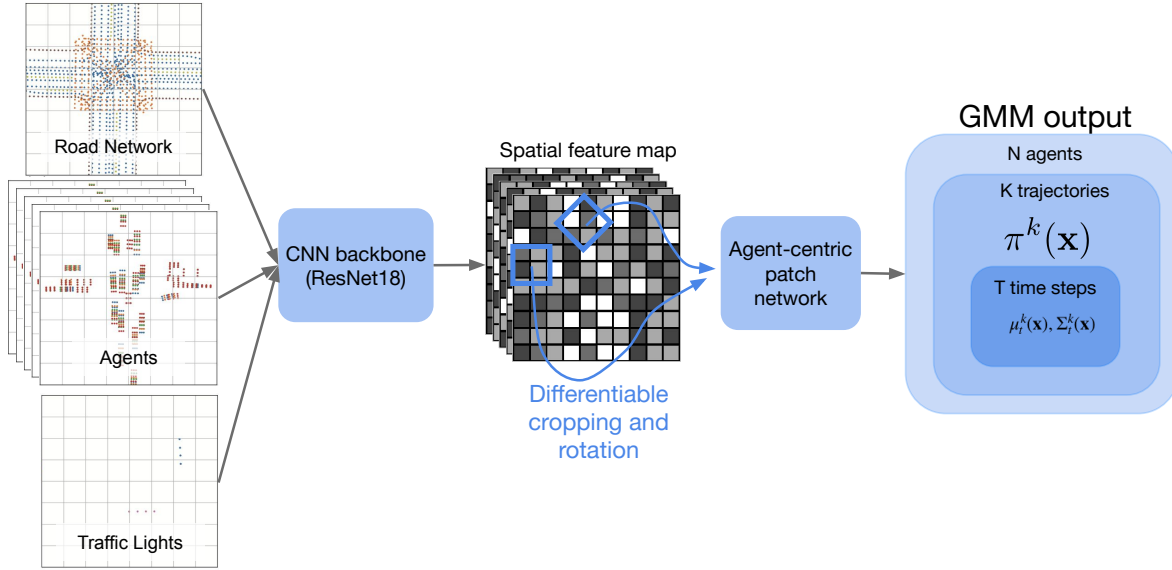


Figure 10: Original MultiPath architecture, with details described in Section 5.3.