

000 001 002 003 004 005 006 007 008 009 010 011 012 013 014 015 016 017 018 019 020 021 022 023 024 025 026 027 028 029 030 031 032 033 034 035 036 037 038 039 040 041 042 043 044 045 046 047 048 049 050 051 052 053 POLYSONA: PARAMETER-EFFICIENT AND MODULAR LATENT BEHAVIOR MODELING FOR TRAFFIC SIMULATION

Anonymous authors

Paper under double-blind review

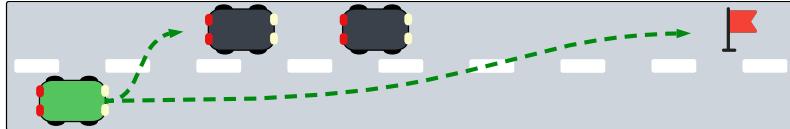
ABSTRACT

In rare but safety-critical driving scenarios, we hypothesize that trajectory outcomes become increasingly multi-modal based on differences between driver style compared to non-critical, common scenarios. However, current approaches for trajectory prediction rarely account for differences in driving style, which may lead to “averaged” driving style in predictions. While average-case behavior may work well in straight driving, easy scenarios, it limits the diversity of outcomes in more complex scenes or in rare events. Extraction of driving style has several benefits, as it enables simulation of counterfactual outcomes in real-world log replays and potentially more accurate predictions through style-consistent predictions. In this paper, we present a parameter-efficient Mixture-of-Experts framework for extraction of latent driving styles in trajectory prediction models. We choose a parameter-efficient approach to reduce forgetting in well-generalized trajectory prediction models, while offering portability of trained driving style modules. We also propose a *Style Consistency Metric* to quantify how often a model’s multi-modal outputs cover the true driving style. In our results, we benchmark different mixture-of-LoRA approaches with our method and show qualitative results that show how the learned experts specialize, and how model saliency changes with our approach. Additional qualitative results can be found on our project website: <https://missanonybloon.github.io/polysona/>

1 INTRODUCTION

As autonomous driving becomes an increasingly accessible technology, handling *mixed autonomy* traffic systems will also become an increasingly important research question. Mixed autonomy traffic systems comprise of both human drivers and autonomous drivers, to varying degrees. This may already be the case in large cities, where autonomous driving has scaled to commercial use (and limited personal use). The general autonomous driving stack approaches driving in sequential modules, each responsible for a specific task: perception, prediction, planning, and control. Prediction, in particular, is responsible for predicting traffic states *that have not occurred yet*. Much of the difficulty in trajectory prediction lies not in the common cases such as lane following and sparser suburban roads, but rather the complex cases with many stressors involved. With more stressors, human driving behavior diverges into a splay of different outcomes; this is expected due to the General Adaptation Syndrome, a spectrum of “fight or flight” responses in humans which has also been shown to influence crowd navigation behavior of humans (Kim et al., 2012).

Currently, this variation in human decision making of other human drivers on the road has not been modeled in trajectory prediction frameworks of autonomous vehicles. Yet, this variable would influence multi-modal outcomes in cases, where an arbitrary decision produces drastically different trajectories, even with a fixed traffic context and route intent. Errors from trajectory prediction can trickle and accumulate down to planning and control modules, resulting in safety implications, especially in risky scenarios. We hypothesize that *driving style* is increasingly impactful during these rare events, where there are typically more stressors pressuring drivers to make decisions in a split second. One motivating example is the case where the driver needs to change lanes to achieve their route goals, but is directly blocked by other drivers in the target lane. If the driver changes lane without modifying their current speed, they would collide with the other drivers. This condition is



054
055
056
057
058 **Figure 1: Motivating Example: Lane change occupied by other vehicles.** The **ego vehicle** has
059 the goal of changing lanes and reaching a particular **waypoint**. At its current velocity, a direct lane
060 change would result in a collision with the other vehicles. Here, the ego vehicle is faced with a natural
061 stressor to make a decision: either decelerate and merge behind the other vehicles, or accelerate and
062 pass the other vehicles first, then merge. Both outcomes are plausible, yet lead to drastically different
063 trajectories by average displacement error metrics, which are computed per-timestep.

064 an example of a stressor on the road. The driver must then choose between two specific outcomes:
065 decelerate and change lanes behind the other drivers, or accelerate and change lanes in front of
066 other drivers. In terms of common trajectory prediction metrics, the difference between the two very
067 plausible outcomes is large; one decision would be penalized, while the other's likelihood would be
068 increased with respect to model parameters. We illustrate this example in Figure 1, where the green
069 vehicle's possible trajectory paths diverge greatly per timestep.

070 Our approach is to model latent driving style as a parameter-efficient Mixture-of-Experts, where
071 we train several expert Low-Rank Adapters (LoRA) (Hu et al., 2022) guided by real world priors
072 on driving style. To guide learned experts towards representations related to driving style, we use
073 a vehicle traffic dynamics adaptation to the Social Forces Model (SFM) for expert routing. This
074 approach is inspired by the use of SFM for other learning-based robotics tasks, such as drone
075 planning (Pang et al., 2021) for collision avoidance among multiple moving agents (i.e. cars) while
076 navigating towards the goal. We motivate the use of parameter-efficient paradigms for its efficiency
077 and baseline performance guarantees. Additionally, adapters make it easy to control for expert
078 behavior, which opens possibilities for counterfactual simulation outcomes.

079 Our main contributions are summarized as follows:

080 1. A parameter-efficient and simple Mixture-of-LoRA approach for latent modeling of driving
081 style in trajectory prediction (Fig. 2).
082 2. A router design guided by social force features from traffic scenarios;
083 3. A style miss rate metric to benchmark simple style consistencies between ground truth
084 trajectories and predictions;
085 4. Qualitative analysis on expert specialization with our MoL approach for latent variable
086 modeling.

088 2 RELATED WORKS

089 2.1 DRIVING STYLE MODELING IN AUTONOMOUS DRIVING

090 Driving style can be useful for several modules in the autonomous driving stack. For policy training,
091 style modeling can be used for personalization of driving policies to maximize rider comfort, espe-
092 cially if the policy is expected to mimic the human's own driving (Karagulle et al., 2024; Schrum et al.,
093 2024). On the other hand, driving style can also be useful for planning, where accurately predicted
094 human driving behavior is essential. Alternative simulated outcomes can also be useful for bench-
095 marking planning modules against counterfactual human behavior. In trajectory forecasting, some
096 work has investigated driving styles for specific maneuvers, such as lane changes (Liu et al., 2021;
097 Hao et al., 2024), or specific context such as highways (Xing et al., 2020) and intersection (Wang et al.,
098 2023) scenarios. Driving style is well-motivated as it can describe variation in plausible outcomes
099 when context and intent are fixed. In our work, context and intent are fixed via the base model; we
100 model latent variable given the sample context and the selected intents from the base model.

101 Several efforts have been made to model driving style from observable trajectory features as a means
102 to extract driving style quantitatively, such as with Graph Neural Networks (GNNs) (Chen et al.,
103 2022; Chandra et al., 2020) or classical clustering approaches based on kinematic properties of
104 trajectories (Hao et al., 2024; Xing et al., 2020). Most recent efforts model driving style with deep
105 learning, using either recurrent networks (Liu et al., 2021; Xing et al., 2020; Choi et al., 2021) or
106 generative models (Kim et al., 2021; Jiao et al., 2022) to extract driving style features implicitly. One
107 common theme of previous work is that driving style is generally considered a latent variable which
can be modeled as a function of kinematic properties such as acceleration and jerk. However, many

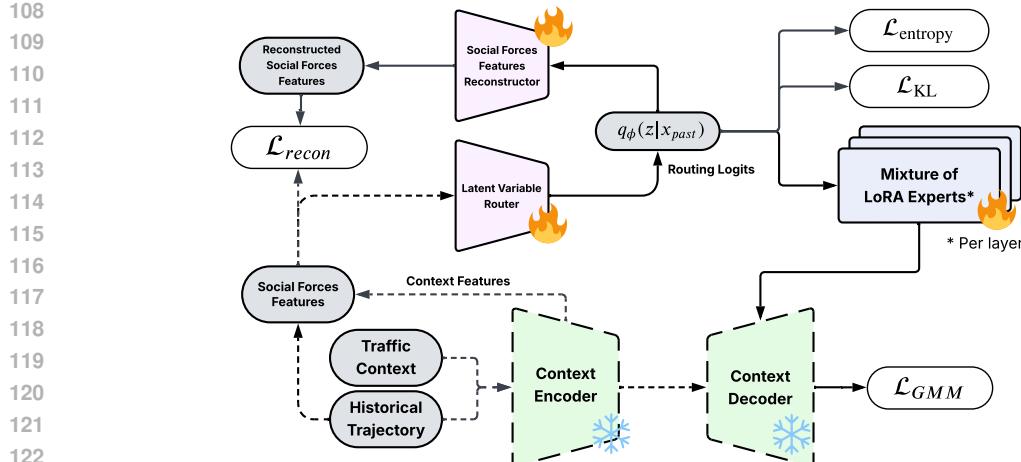


Figure 2: Training Overview. We use a Mixture-of-LoRA (MoL) approach to learn a driving style latent variable for trajectory prediction models. Solid lines indicate gradient flow in backpropagation. In summary, we construct social forces features from historical agent trajectories and train a **routing network** in a Variational Autoencoder-like fashion; in other words, the router predicts logits for discrete latent classes, and a reconstructor network then learns to re-construct social forces features. The predicted routing logits are then used to select experts in the **MoL layers**, which are attached to certain decoder layers of the **base model**. Experts are combined using the weight mixing scheme illustrated in Figure 3. All modules except the **base model** are trained end-to-end.

of these works require an annotated dataset (Liu et al., 2021) or a specialized backbone architecture to extract driving style (Kim et al., 2021; Chandra et al., 2020). The most relevant recent work in controllable behavior modeling with latent variables may be TAE (Jiao et al., 2022), which also uses a VAE-like objective, but models the driving style latent as a continuous instead of categorical variable. There are several limitations of this work. Firstly, the latent variable being continuous makes it difficult to interpret the latent. Secondly, TAE does not achieve competitive results to SOTA trajectory prediction frameworks, which limits the realism of simulated outcomes. We show this comparison in our experimental results to emphasize the performance gap in previous work to our proposed method in realistic prediction.

In our work, we distinguish our goals from previous works in that we want to complement and maintain performance from state-of-the-art trajectory prediction models, whilst also being able to extract generalized categorical representations of driving style.

2.2 PARAMETER-EFFICIENT MIXTURE OF EXPERTS FOR SPECIALIZATION IN DRIVING

Parameter-efficient Mixture-of-Experts (PE-MoE) is a key focus in language modeling and image diffusion, where efficient fine-tuning atop foundation models is essential. Accordingly, prior PE-MoE work has largely centered on language and diffusion tasks. We extend these paradigms to trajectory prediction, aiming to extract a latent variable that captures driving style. PE-MoEs typically use collections of LoRA adapters as experts. Single-LoRA adapters can encode specialized concepts in image diffusion (Gandikota et al., 2024; Gu et al., 2023), while in language tasks, PE-MoEs mitigate catastrophic forgetting (Dou et al., 2023) and match the performance of full fine-tuning at lower cost (Zadouri et al., 2024; Li & Zhou, 2025). Using multiple pre-trained LoRAs also enables reusing task-specific skills. MoE approaches in this context aim to train a robust router to allocate experts per sample. A survey of adapter merging techniques was recently published by Yadav et al. (Yadav et al., 2024). Multi-task learning with LoRAs similarly targets a latent skill representation (task-skill matrix), reusing skills across tasks, akin to PE-MoEs. Notable examples include PolytroponPonti et al. (2023), C-Poly (Wang et al., 2024a), and Hyperformer (Karimi Mahabadi et al., 2021). These are complementary to our work, and our MoE design is compatible with most, with some limitations on the differentiability of global expert routers, which we will discuss in the next section.

Some works apply parameter-efficient fine-tuning approaches to trajectory prediction. Forecast-PEFT (Wang et al., 2024b) uses prompt vectors and LoRA layers to improve trajectory prediction performance in downstream tasks and achieve close results to full fine-tuning, similarly to language

models. On the other hand, Munir et al. explore the possibility of using Large Language Models in trajectory prediction (Munir et al., 2024), thus their application of LoRAs would be in the language domain. Our use case is different; we learn a latent variable across different driving styles in a parameter-efficient manner. In robotics, MoEs have been applied to vehicle trajectory prediction, often focusing on goal or intention inference (Yuan et al., 2024). We instead focus on driving style, with the intention of maintaining driving style continuity between past trajectories and future predictions. MoEs have also improved drone trajectory prediction (Fraser et al., 2023). Alternative latent-variable methods like VAEs (Xu et al., 2022) are more computationally costly and limit latent portability across models.

3 METHODOLOGY

3.1 LATENT DRIVING STYLES WITH EXPERTS

Driving styles are difficult to model because the notion of “driving style” is ill-defined and non-standardized, especially in autonomous driving research. Due to the difficulties of labeling driving style consistency and reliably, many datasets often do not include driving style related information, making this problem challenging—akin to latent variable modeling. In this case, we observe *driving style* as an outcome of latent variable modeling.

In our approach, we make the key assumption that driving style must be strongly correlated to *social forces between agents*. Second-order kinematics is directly related to the actions a driver takes with the steering wheel, throttle, and brake, as input to control a vehicle. That is, the actions patterns that a driver takes on steering, throttle, and brake should be directly correlated with the driving style. The distribution of social forces on each agent over time may be a strong signal related to car-following parameters such as comfortable following distance, preferred velocity, maximum acceleration, and minimum deceleration. Such parameters should not change according to context, and should be constant over time. We illustrate our approach in Figure 2. Our approach first assumes a pre-trained trajectory forecasting model based on encoder-decoder transformer architectures. This model is frozen and serves as a shared representation of general driving behavior. Each MoL layer, depicted in blue, has some variation of the process depicted in Figure 3, which shows the mixing scheme for CAT (Prabhakar et al., 2024), the most recently published method for MoL.

Mixture of LoRA Experts. *Low-rank adapters (LoRA)* (Hu et al., 2022) are a popular method for parameter-efficient finetuning (PEFT), which is popularly used for large language models. LoRA represents finetuning as a low-rank update to pre-trained weight matrices $W_0 \in \mathbb{R}^{d \times k}$:

$$h = W_0x + \Delta Wx + b_0 = W_0x + BAx + b_0 \quad (1)$$

where $B \in \mathbb{R}^{d \times r}$ and $A \in \mathbb{R}^{r \times k}$, and the rank of both matrices $r \ll \min(d, k)$. During training, the pre-trained weights W_0 are frozen and only matrices B and A are updated.

Mixture-of-experts (MoE) is a modular paradigm which consists of N learned experts $\{E_1, \dots, E_N\}$ and a router function which combines outputs from each expert in a weighted voting fashion. The Experts E_n are typically learned jointly and end-to-end, where the objective is for each expert to learn a specialized representation. In the forward pass through a MoE model, the overall output is determined by a weighted “vote” of the individual expert outputs. The voting weights are known as *routing weights*, and a mixture function G determines how the outputs are combined using the routing weights. In our work, G is simply a weighted sum.

For our use case, we would like to learn specialized experts which encode latent driving styles, but also want to avoid the computational burden of full MoEs. Thus, we use Mixture-of-LoRA (MoL),

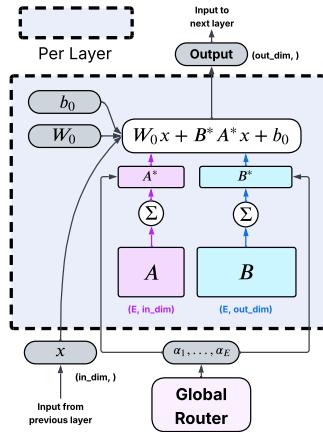


Figure 3: **CAT weight mixing scheme in MoL layers.** In our experiments, we find that the CAT weight-mixing scheme (Prabhakar et al., 2024) performs better on Kalman Hard difficulty scenarios.

216 which trains a set of expert LoRAs, instead of another trajectory prediction model. The mixing
 217 paradigm of the LoRA experts can be interchangeable with any existing paradigm for MoL. Thus, we
 218 benchmark several in our experiments. We note that MoL has not yet been benchmarked in trajectory
 219 prediction so far, to the best of our knowledge. Thus, our experiments are also some of the first
 220 benchmarks on applying parameter-efficient techniques to trajectory models.

221 **Expert Routing based on Social Forces.** We create a global expert router based on social forces and
 222 environment context features from the context encoder of the backbone network. The router is treated
 223 like a variational posterior estimator $q_\phi(z|x_{\text{past}})$ for social forces and context features; we sample
 224 the latent differentiably with Gumbel-Softmax from the predicted latent logits, then reconstruct the
 225 input social forces and context features. We compute social force features from the input trajectory
 226 information based on the Social Forces Model. For a particular intersection between vehicles i and j ,
 227 the repulsion force between them is defined as:

$$F_{ij} = \alpha \cdot e^{(\delta - \|x_j - x_i\|)/\beta} \cdot \hat{\theta}_{ij} \quad (2)$$

231 where α is a hyperparameter scaling repulsion magnitude, β is a hyperparameter scaling the decay
 232 rate as pairwise distances become larger, and δ is a hyperparameter corresponding to preferred, or
 233 comfortable, distance between agents. All three hyperparameters are set to 1 in our experiments.
 234 When constructing social forces features, interactions between a vehicle and itself are masked out,
 235 along with vehicles beyond a radius threshold of 10 meters vehicles not traveling in the same direction.
 236 To combine social forces features with context embeddings of hidden dimension d , we concatenate
 237 both social forces features and context features to a vector of dimension $2d$, then use a single dense
 238 layer to fuse features together, resulting in a final feature dimension of d . This fused representation is
 239 used as input to the global expert router.

240 The expert router architecture is a simple 2-layer multi-layer perceptron: a single dense layer
 241 projecting to a hidden dimension of 64, a normalization layer, ReLU nonlinearity, dropout at $p=0.1$,
 242 and finally a dense layer outputting logits for E classes, or the number of experts. In our experiments,
 243 we use three experts ($E = 3$) by following priors from existing traffic psychology literature, which
 244 suggests that there are three distinct types of drivers (Klauer et al., 2009). In brief experimentation
 245 with higher-complexity global routers, we found that heavier routers overfit quickly and deteriorate
 246 performance.

247 **Loss Objectives.** We use three loss terms to guide the MoL layers towards a driving style representation.
 248 Firstly, we maintain the original trajectory prediction objective, which maximizes the likelihood
 249 of the ground-truth trajectory given a predicted Gaussian Mixture of trajectory outcomes (L_{GMM}).
 250 Secondly, since driving style is an unsupervised variable, we enforce a Variational Autoencoder-like
 251 (VAE-like) objective on the global router, which eventually serves as a “persona classifier”. We have
 252 two terms enforced on the router: a reconstruction term, which reconstructs the social forces features
 253 s extracted in the previous section, and a Kullback-Leibler divergence term, which maximizes the
 254 evidence lower-bound (ELBO), which is a lower bound on the probability of observing the data
 255 generated by the model $p_\theta(x)$. Thus, the loss objective in training \mathcal{L} becomes:

$$\mathcal{L} = \mathcal{L}_{GMM} + \lambda_{\text{recon}} \mathcal{L}_{\text{recon}} + \lambda_{KL} KL(q_\phi(z|x_{\text{past}}) || p(z)) + \lambda_{\text{entropy}} \mathcal{L}_{\text{entropy}} \quad (3)$$

$$\mathcal{L}_{GMM} = \mathbb{E}_{q_\phi}(z|x_{\text{past}}) [-\log p_\theta(x_{\text{future}}|z)] \quad (4)$$

$$\mathcal{L}_{\text{recon}} = \|s_{\text{orig}} - s_{\text{pred}}\|_2 \quad (5)$$

$$\mathcal{L}_{\text{entropy}} = - \sum_z q_\phi(z | x_{\text{past}}) \log q_\phi(z | x_{\text{past}}) \quad (6)$$

262 3.2 MEASURING STYLE CONSISTENCY

264 In trajectory prediction tasks, standard metrics such as *Average Displacement Error (ADE)* and
 265 *Final Displacement Error (FDE)* measure how closely a predicted path follows the ground truth in
 266 Euclidean space. While these metrics capture spatial accuracy, they are agnostic to the underlying
 267 kinematic signatures (e.g. rapid acceleration or hard braking) that distinguish different driving styles.
 268 To address this gap, we introduce a *style consistency metric*, which evaluates whether a model’s
 269 multi-modal outputs “cover” the true driving style. The core of this metric is a Gaussian Mixture
 Model (GMM) that assesses style consistency between the ground truth and its predicted trajectories.

270 **Table 1: Trajectory Prediction Benchmark Performance Comparisons.** We benchmark different
 271 variants of our approach with different parameter-efficient MoL schemes (bottom half) compared
 272 to existing benchmark trajectory prediction models (top half). TAE (Jiao et al., 2022) is the most
 273 recent work in modeling latent driving style to our proposed method, as it does not require any
 274 privileged datasets; as shown, the performance is not competitive to baseline trajectory prediction
 275 models. All non-baseline experiments are averaged over three fixed random seed runs. Green cells
 276 indicate improvements over TAE. Bolded values indicate best-performance in respective columns.
 277 Our method improves upon the baseline model in overall trajectory prediction metrics across all MoL
 278 schemes. *The best all-around results are achieved with the MoV approach (Zadouri et al., 2024),*
 279 which uses IA3 (Liu et al., 2022) for finetuning, instead of LoRA.

Method	# Trainable Params	∇Router	brierFDE↓	minADE↓	minFDE↓	MissRate↓
Autobot (Girgis et al., 2022)	1.5M	-	2.4439	0.8892	1.7817	0.2803
Wayformer (Nayakanti et al., 2022)	15.2M	-	2.5747	0.9348	1.9718	0.3574
MTR (Shi et al., 2022)	65.2M	-	2.1702	0.8645	1.7094	0.3107
MTR+Actions (Zheng et al., 2024)	65.2M	-	2.1605	0.8658	1.7102	0.3208
TAE (Jiao et al., 2022)	249K	-	12.0023	4.8269	12.0023	0.9275
PolySona+MTR+Polytropon (Ponti et al., 2023)	527K	✗	2.1602	0.8623	1.7037	0.3164
PolySona+MTR+C-Poly (Wang et al., 2024a)	969K	✗	2.1613	0.8625	1.7048	0.3169
PolySona+MTR+HyperFormer (Karimi Mahabadi et al., 2021)	527K	✗	2.1601	0.8623	1.7038	0.3159
PolySona+MTR+CAT (Prabhakar et al., 2024)	526K	✓	2.1607	0.8624	1.7041	0.3171
PolySona+MTR+MoV (Zadouri et al., 2024)	195K	✓	2.1588	0.8619	1.7016	0.3158

280 Following the clustering methodology introduced in (Zheng et al., 2025), we argue that a simpler
 281 clustering model can be a more interpretable and robust style evaluator than complex deep latent
 282 approaches. Specifically, we fit a two-component GMM to summary statistics of the ground-truth
 283 trajectories in the validation set. Each trajectory is represented by a feature vector comprising the
 284 maximum absolute acceleration, the variance of acceleration, the variance of speed, and the “gamma”
 285 statistic (variance of jerk divided by mean jerk) as defined in (Murphay et al., 2009), omitting mean
 286 speed to avoid contextual entanglement. We then assign each candidate trajectory to one of the two
 287 mixture components—“normal” or “aggressive.” Finally, we define the *Style Miss Rate (SMR)* as
 288 the fraction of samples for which none of the predicted trajectories share the same style cluster as
 289 the ground truth. A lower SMR indicates better coverage of the driver’s true style, complementing
 290 traditional displacement-based metrics. For clarity, we summarize the SMR formulation below:
 291

$$\text{SMR} = \frac{1}{N} \sum_{n=1}^N \mathbf{1}(\forall i \in \{1, \dots, K\}, s_{n,i} \neq s_n^*) \quad (7)$$

300 where s_n^* is the ground-truth style for sample n , $s_{n,i}$ are its K predicted styles, and $\mathbf{1}(\cdot)$ is the
 301 indicator function. More details and exact formulation can be found in the Appendix.
 302

4 RESULTS

308 **Hardware.** Each experiment is trained on 32 GB memory, two AMD EPYC 7352 24-Core Processors,
 309 and two NVIDIA RTXA5000 GPUs.
 310

311 **Experiment setup.** In our experiments, the base trajectory
 312 prediction model used is Motion Transformer (MTR), a state-
 313 of-the-art, open source, non-autoregressive approach. We train
 314 experiments on the Argoverse 2 dataset (Wilson et al., 2021),
 315 which consists of 250,000 total scenarios. We use 183,333
 316 samples for training and 22,979 samples for validation. Our
 317 experiments were conducted with the UniTraj framework (Feng
 318 et al., 2024) for trajectory prediction. In our experiments, the
 319 task is to predict the next future 6s trajectories, given 2s of agent
 320 history trajectories and road context polylines. Agent trajectories
 321 are sampled at 10Hz, for a total of 20 history frames and
 322 60 future frames. Baseline experiments which do not involve
 323 our proposed method are all trained based on their original
 324 recommended hyperparameter settings. All experiments with
 325 our proposed method are trained on 10 epochs and evaluated on the 10th epoch. All experiments with

326 **Table 2: Style Miss Rate between
 327 baseline model (*) and variants
 328 of our method.**

Method	SMR↓
MTR*	0.6317
TAE	0.7607
Ours+Polytropon	0.2198
Ours+C-Poly	0.2279
Ours+HyperFormer	0.2248
Ours+CAT	0.2293
Ours+MoV	0.2215

324 Table 3: **minADE Comparison by Kalman Difficulty and TDBM Driving Style groups.** We also
 325 compare variants of our method on different MoL schemes to baseline models and the most recent
 326 latent driving style modeling baseline, TAE, across Kalman difficulty categories and TDBM driving
 327 styles (Cheung et al., 2018). In general, we find that *weight mixing provides the greatest improvement*
 328 to samples with hard Kalman difficulty. None of the variants improved minADE on “careful” driving
 329 samples over baseline models; however, “careful” is also the TDBM driving style assigned to the vast
 330 majority of common driving scenarios.

Method	Kalman Difficulty			TDBM Driving Styles			
	Easy	Medium	Hard	Timid	Careful	Reckless	Threatening
Autobot (Girgis et al., 2022)	0.8399	1.2436	1.9054	0.9186	0.9765	0.8906	0.8434
Wayformer (Nayakanti et al., 2022)	0.8764	1.3440	3.0969	0.9676	0.9484	0.9362	0.8861
MTR (Shi et al., 2022)	0.8195	1.1736	3.0148	0.8849	0.7718	0.8670	0.8190
MTR+Actions (Zheng et al., 2024)	0.8212	1.1781	2.5331	0.8846	0.7904	0.8687	0.8194
TAE (Jiao et al., 2022)	4.1203	9.9024	20.3377	4.8066	3.9524	4.4860	4.8664
PolySona+MTR+PolyTropo (Zadouri et al., 2024)	0.8185	1.1690	2.5585	0.8803	0.8665	0.8653	0.8156
PolySona+MTR+C-Poly (Wang et al., 2024a)	0.8187	1.1691	2.5474	0.8803	0.8472	0.8655	0.8166
PolySona+MTR+HyperFormer (Karimi Mahabadi et al., 2021)	0.8186	1.1673	2.5707	0.8803	0.8471	0.8653	0.8156
PolySona+MTR+CAT (Prabhakar et al., 2024)	0.8188	1.1678	2.4978	0.8793	0.8477	0.8655	0.8161
PolySona+MTR+MoV (Zadouri et al., 2024)	0.8180	1.1690	2.5502	0.8803	0.8652	0.8649	0.8150

340 our proposed method are initialized from a MTR+Actions (Zheng et al., 2024) checkpoint trained
 341 on Argoverse 2; thus, there is no need to account for domain shift in fine-tuning. All experiments
 342 are trained to learn three experts (or driving styles) with a prior probability of [0.3, 0.6, 0.1] for each
 343 expert, respectively, which is inspired from driver distributions published by the NHTSA (Klauer
 344 et al., 2009). In the base model, there are six decoder layers, where each decoder layer consists of
 345 agent and map attention blocks, as well as a Gaussian mixture prediction head. We apply Rank-4
 346 MoL to the object attention and the Gaussian Mixture prediction head of the 3rd and 6th decoder
 347 layers (6 decoder layers total). This produces about 30 MoL layers in MTR experiments. More
 348 hyperparameter details for experiment reproduction can be found in the Appendix.

349 **Metrics.** We use four standard metrics from trajectory prediction to measure model performance:

- 351 • **BrierFDE** (↓): error (m) of each prediction mode of Gaussian Mixture, weighted by the
 352 mixture score. Accounts for both trajectory accuracy and model confidence.
- 353 • **minADE** (↓): Minimum displacement error (m) from the ground truth trajectory of the
 354 predicted trajectory modes, averaged over time. Best-case, per-timestep error.
- 355 • **minFDE** (↓): Minimum final displacement error (m) between the closest ending position of
 356 the predicted trajectory and the ground truth trajectory. Best-case, final-timestep error.
- 357 • **Miss Rate** (↓): The rate at which minFDE is greater than 2 meters. When the Miss Rate is
 358 1, no predictions ended within a radius of 2 meters of the final position of the ground truth.

360 4.1 COMPARISONS TO OTHER LATENT BEHAVIOR METHODS AND MoL VARIANTS

361 We compare our method to several baselines in trajectory prediction, including a recent method in
 362 modeling latent driving style, TAE (Jiao et al., 2022). Latent variable modeling is common in modern
 363 trajectory prediction architectures, but may not always pertain specifically to variations in driving
 364 style. For example, regression-based trajectory prediction models such as MTR (Shi et al., 2022) and
 365 Wayformer (Nayakanti et al., 2022) use Gaussian Mixture prediction heads to model latent output
 366 modes based on route intent. Our work focuses on modeling variations in driving style, where we
 367 assume intent is fixed. As mentioned in the related works, other approaches for modeling latent
 368 driving style in a similar fashion either (1) require a privileged dataset involving human subjects
 369 which cannot be released, or (2) do not fix intents when modeling latents. Regardless, we show
 370 comparisons to both SOTA trajectory prediction models (which employ latent intent modeling) and
 371 SOTA driving style latent variable modeling (TAE).

372 In Table 1, we compare the overall performance on Brier FDE, minADE, minFDE, and Miss Rate
 373 across different trajectory prediction architectures (top half), recent work on latent driving style
 374 modeling (TAE (Jiao et al., 2022)), and MoL approaches (bottom half) when applied in our framework.
 375 Each experiment is averaged over three fixed random seed runs, except for the baseline models in
 376 the top half. As demonstrated by significant gaps in performance metrics for trajectory prediction,
 377 our approach, PolySona, shows much better performance on trajectory prediction metrics thanks to
 its compatibility with state-of-the-art trajectory prediction architectures. Aside from this, PolySona

378 is also modular and categorical compared to TAE, enabling more controllable simulation outcomes
 379 with counterfactual latents.
 380

381 C-Poly (Wang et al., 2024a), HyperFormer (Karimi Mahabadi et al., 2021), and Polytropon (Ponti
 382 et al., 2023) are parameter-efficient approaches from multi-task learning (MTL) literature, rather than
 383 from Mixture-of-LoRA. While MTL is typically distinct from MoL work, these MTL methods are
 384 very similar in that they use distinct LoRA matrices per task. However, differently from our use
 385 case, they 1) assume that LoRA adapters are pre-trained and 2) assume task classes are a given. In
 386 other words, expert assignments under these MTL approaches are hard and non-differentiable. This
 387 means that the router cannot be learned end-to-end with the experts. And, since we do not have
 388 access to ground-truth driving style annotations, we also cannot pre-train the LoRA experts. For
 389 MTL experiments, the router only depends on VAE objectives to choose experts, since gradients
 390 from the GMM loss do not propagate back to the router (∇ Routing in Table 1). Table 2 shows SMR
 391 performance across variants of our method; we achieve considerable improvement of *at minimum*
 392 63.7% across all methods compared to the MTR and TAE baselines, with the best SMR score being
 393 with the Polytropon variant (Ponti et al., 2023). We note that improvement on SMR also scales
 394 across methods similarly to other performance metrics in Table 1. In Table 3, we show results for
 395 baseline models and variants of our approach across several Kalman Difficulty Feng et al. (2024)
 396 and TDBM Cheung et al. (2018) subsets of Argoverse 2. Kalman difficulty, originally introduced in
 397 UniTraj Feng et al. (2024), is defined as the magnitude of deviation in meters from ground truth when
 398 a linear Kalman filter is used to predict the future trajectory. In TDBM, both traffic state features and
 399 user study responses are used to determine a linear feature-behavior mapping which maps relative
 400 vehicle features to different driving behavior categories. Overall, all variants of our method improves
 401 both the overall performance metrics and across most scenarios detailed in Table 3. Amongst different
 402 MoL paradigms, we find that MoV (Zadouri et al., 2024) achieves the best all-around performance
 403 boosts. As an additional plus, this also the approach with the lowest number of trainable parameters.
 404

403 4.2 INTERPRETING EXPERT SPECIALIZATIONS

404 We would like to investigate what behavior each expert specializes in from the learned model.
 405 To investigate interpretable differences between experts, we compare the second-order kinematic
 406 properties of predicted trajectories by each expert. We predict Argoverse 2 validation trajectories,
 407 and plot the mean of the absolute value of acceleration (m/s^2), jerk (m/s^3), angular acceleration
 408 (θ/s^2), and angular jerk (θ/s^3). We plot this comparison in Figure 4a.
 409

410 To visualize whether there are non-trivial latent assignments by the MoL layers, we also plot the
 411 t-SNE features computed using router embeddings, which is visualized in Figure 4b. Our router
 412 learns a clear three-class separation among social forces features; we also observe empirically that
 413 reconstruction loss converges well, implying that three latent classes are sufficient to reconstruct the
 414 social forces features of the target agent.
 415

416 4.3 ABLATION STUDY

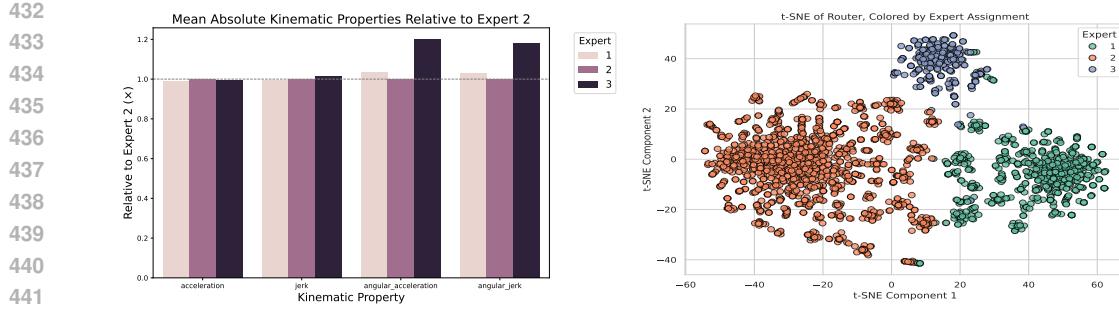
417 We conduct an ablation study on our approach by experimenting with alternative variants with respect
 418 to router expressivity, full fine-tuning, and contribution of each objective term. In Table 4, we record
 419 the percentage difference for each ablation variant of our full approach, which can be referenced in
 420 the last row of Table 1. In general, we find that the *greatest degradation in performance occurs when*
 421 *the reconstruction process or KL divergence is removed* during training. Unsurprisingly, the router
 422 experiences expert collapse when KL loss is removed. Expert collapse also occurs when entropy
 423 loss on the posterior logits is removed, which is aligned with our expectations. Full fine-tuning
 424 ($\sim 65M$ trainable parameters) performs very similarly to parameter-efficient learning; while minADE
 425 improves slightly, we find that all other metrics do degrade slightly.
 426

427 4.4 HOW MANY LATENTS ARE ENOUGH?

428 We also explore how the number of categorical latents modeled
 429 influences performance across trajectory prediction metrics,
 430 grouped by difficult categories. In particular, we train several
 431 variants our model with the CAT weight mixing, varying only
 432 the number of latents modeled. Each variant is trained three

Table 5: minADE by # Latent
 Classes & Kalman Difficulty Level

K	minADE \downarrow			
	Easy	Medium	Hard	Overall
2	0.8108	1.1951	3.5275	0.8589
3	0.8188	1.1678	2.4978	0.8624
4	0.8119	1.1990	3.5116	0.8603
5	0.8120	1.1996	3.5191	0.8605
6	0.8109	1.1986	3.5729	0.8595
8	0.8122	1.2027	3.4217	0.8610



432
433
434
435
436
437
438
439
440
441
442
443
444
445
446
447
448
449
450
451
452
453
454
455
456
457
458
459
460
461
462
463
464
465
466
467
468
469
470
471
472
473
474
475
476
477
478
479
480
481
482
483
484
485
(a) Relative magnitudes of second-order kinematics. (b) t-SNE visualization of router embeddings.

Figure 4: **Visualizations of expert representations in Ours+CAT.** We plot both relative magnitudes of kinematic properties and projected router embeddings by expert. In (a), we show the mean magnitudes of acceleration, jerk, angular acceleration, and angular jerk between experts, relative to Expert 2. We observe that *the learned latents pertain more to angular; lateral kinematics, as opposed to longitudinal kinematics along the lane*. In (b), we pass the same set of scenarios through the trained router and project the embeddings to 2 dimensions with t-SNE, then color each point by the router’s expert prediction. This figure shows that *the learned latents are not collapsing to redundant representations*, which is a common practical concern with latent variable modeling.

Table 4: **Ablation Study.** We evaluate the impact of each component in our MoL training scheme for trajectory prediction using the Ours+CAT variant. All values are averaged over three fixed random seed runs. In this table, we quantify the % improvement in each respective metric. We find that the greatest degradation to performance is from removing the reconstruction process or the KL loss.

Ours	% brierFDE↑	% minADE↑	% minFDE↑	% MissRate↑
+ LinearRouter	-0.153%	0.024%	-0.226%	0.157%
+ Full Finetuning	-0.114%	0.444%	-0.313%	-0.673%
- Social Forces	-0.221%	0.299%	-0.401%	-1.060%
- Context Features	-0.186%	0.096%	-0.363%	-1.199%
- KL Loss	-0.243%	0.136%	-0.439%	-1.669%
- Reconstruction	-0.500%	-0.071%	-0.780%	-1.337%
- Expert Entropy Loss	-0.099%	-0.096%	-0.135%	0.265%

times across three fixed random seeds and averaged, to account for randomness in results. This experiment is motivated by two questions: 1) How many latents are sufficient to model driving style for difficult scenarios? And 2) As number of latents modeled increase, does the performance also increase due to increased expressivity?

Results are shown in Table 5. As we expect based on our prior knowledge, we observe that minADE performance across Medium and Hard Kalman difficulty scenarios is best when modeling $K = 3$, *especially for hard scenarios*. Despite the overall minADE for $K = 3$ being the highest of all values, we find that this is due to representation being skewed towards Easy scenarios. Still, best overall performance results from modeling $K = 2$, where we find the lowest minADE for Easy scenarios.

To answer the two questions above, our experiment results suggested that *three latent classes is sufficient to model driving style for difficult scenarios*, and that *increasing the number of latents does not necessarily increase overall performance of the model*. Drawing from this, we can also consider traditional VAEs to model an infinite number of latents in a continuous vector space—this result may also illustrate the interpolation of performance to the continuous case.

Interestingly, we find that $K = 3$ also corroborates our prior from Traffic Psychology literature, where scientists designing human subjects studies initially divided drivers into two “driving style” groups, but later modified the study to account for three “driving style” groups, as having three groups resulted in better separation between varying levels of risk taking (Dingus et al., 2006).

486 5 CONCLUSION

487 In this paper, we introduced a framework based on Mixture-of-LoRA (MoL) to extract driving style
 488 variables from pre-trained trajectory models. The framework is both parameter-efficient and modular,
 489 making it easy to adapt to existing trajectory prediction works. This framework allows for controllable
 490 routing due to the global router, which learns a latent variable for driving style. Our contributions
 491 also offer a metric focused on driving style consistency and qualitative analysis on model saliency,
 492 which can be useful for interpreting prediction reasoning in future work.

494 6 REPRODUCIBILITY STATEMENT

496 We, the authors, emphasize reproducibility for this work, especially since nearly all previous work
 497 (to the best of our knowledge) do not have reproducible or open-sourced methodologies. To uphold
 498 complete reproducibility of our work, we include an anonymized code link in our supplemental
 499 material, in addition to referencing it on our anonymized project site, which can be found at the
 500 end of the abstract. Additionally, we include all hyperparameter details for the setup in Figure 2 in
 501 Appendix Section 8. Our implementation is based on the unified trajectory prediction framework,
 502 UniTraj (Feng et al., 2024), and can thus be integrated with all supported open source architectures
 503 and datasets. In our experiments, we primarily use Argoverse 2 (Wilson et al., 2021).

504 In the future, we plan to open-source our full implementation, including scripts used to run all fixed
 505 seed experiments and averaged results. By doing so, we hope to support more research in modeling
 506 driving style for autonomous driving.

508 7 USAGE OF LLMs

510 The use of LLMs in our work is strictly limited to two things: implementation of TAE and aiding in
 511 plot styling. While TAE is a relevant comparison for our proposed approach, it is not open sourced;
 512 we were unable to obtain the original source implementation. Thus, we re-implemented TAE to the
 513 best of our ability using details from the original TAE publication, then used LLMs to refine details
 514 on our implementation. All writing was done without the use of LLMs.

516 REFERENCES

518 Rohan Chandra, Aniket Bera, and Dinesh Manocha. Stylepredict: Machine theory of mind for human
 519 driver behavior from trajectories. *CoRR*, abs/2011.04816, 2020. URL <https://arxiv.org/abs/2011.04816>.

522 Chao Chen, Qiang Liu, Xingchen Wang, Chengwu Liao, and Daqing Zhang. semi-traj2graph
 523 identifying fine-grained driving style with gps trajectory data via multi-task learning. *IEEE*
 524 *Transactions on Big Data*, 8(6):1550–1565, 2022. doi: 10.1109/TB DATA.2021.3063048.

525 Ernest Cheung, Aniket Bera, Emily Kubin, Kurt Gray, and Dinesh Manocha. Identifying driver behav-
 526 iors using trajectory features for vehicle navigation. In *2018 IEEE/RSJ International Conference*
 527 *on Intelligent Robots and Systems (IROS)*, pp. 3445–3452. IEEE, 2018.

529 Seungwon Choi, Nahyun Kweon, Chanuk Yang, Dongchan Kim, Hyukju Shon, Jaewoong Choi, and
 530 Kunsoo Huh. Dsa-gan: Driving style attention generative adversarial network for vehicle trajectory
 531 prediction. In *2021 IEEE International Intelligent Transportation Systems Conference (ITSC)*, pp.
 532 1515–1520, 2021. doi: 10.1109/ITSC48978.2021.9564674.

533 Thomas A Dingus, Sheila G Klauer, Vicki Lewis Neale, Andy Petersen, Suzanne E Lee, Jeremy
 534 Sudweeks, Miguel A Perez, Jonathan Hankey, David Ramsey, Santosh Gupta, et al. The 100-car
 535 naturalistic driving study, phase ii-results of the 100-car field experiment. Technical report, United
 536 States. Department of Transportation. National Highway Traffic Safety . . ., 2006.

538 Shihhan Dou, Enyu Zhou, Yan Liu, Songyang Gao, Jun Zhao, Wei Shen, Yuhao Zhou, Zhiheng Xi,
 539 Xiao Wang, Xiaoran Fan, et al. Loramoe: Revolutionizing mixture of experts for maintaining
 world knowledge in language model alignment. *arXiv preprint arXiv:2312.09979*, 4(7), 2023.

540 Lan Feng, Mohammadhossein Bahari, Kaouther Messaoud Ben Amor, Éloi Zablocki, Matthieu Cord,
 541 and Alexandre Alahi. Unitraj: A unified framework for scalable vehicle trajectory prediction.
 542 *arXiv preprint arXiv:2403.15098*, 2024.

543

544 Benjamin Fraser, Adolfo Perrusquía, Dimitrios Panagiotakopoulos, and Weisi Guo. A deep mixture
 545 of experts network for drone trajectory intent classification and prediction using non-cooperative
 546 radar data. In *2023 IEEE Symposium Series on Computational Intelligence (SSCI)*, pp. 1–6, 2023.
 547 doi: 10.1109/SSCI52147.2023.10371877.

548 Rohit Gandikota, Joanna Materzyńska, Tingrui Zhou, Antonio Torralba, and David Bau. Erasing
 549 concepts from diffusion models. In *Proceedings of the 2024 IEEE European Conference on*
 550 *Computer Vision*, 2024. arXiv preprint arXiv:2311.12092.

551

552 Roger Girsch, Florian Golemo, Felipe Codevilla, Martin Weiss, Jim Aldon D’Souza, Samira Ebrahimi
 553 Kahou, Felix Heide, and Christopher Pal. Latent variable sequential set transformers for joint
 554 multi-agent motion prediction. In *International Conference on Learning Representations*, 2022.
 555 URL https://openreview.net/forum?id=Dup_dDqkZC5.

556

557 Yuchao Gu, Xintao Wang, Jay Zhangjie Wu, Yujun Shi, Yunpeng Chen, Zihan Fan, WUYOU
 558 XIAO, Rui Zhao, Shuming Chang, Weijia Wu, Yixiao Ge, Ying Shan, and Mike Zheng Shou.
 559 Mix-of-show: Decentralized low-rank adaptation for multi-concept customization of diffusion
 560 models. In A. Oh, T. Naumann, A. Globerson, K. Saenko, M. Hardt, and S. Levine (eds.),
 561 *Advances in Neural Information Processing Systems*, volume 36, pp. 15890–15902. Curran Asso-
 562 ciates, Inc., 2023. URL https://proceedings.neurips.cc/paper_files/paper/2023/file/3340ee1e4a8bad8d32c35721712b4d0a-Paper-Conference.pdf.

563

564 Jinbo Hao, Hui Xie, Fan Guo, Yun Chen, and Kang Song. Vehicle trajectory prediction with driving
 565 style identification and intention fusion. In *2024 8th CAA International Conference on Vehicular*
 566 *Control and Intelligence (CVCI)*, pp. 1–6, 2024. doi: 10.1109/CVCI63518.2024.10830035.

567

568 Edward J Hu, yelong shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang,
 569 and Weizhu Chen. LoRA: Low-rank adaptation of large language models. In *International*
 570 *Conference on Learning Representations*, 2022. URL <https://openreview.net/forum?id=nZeVKeFYf9>.

571

572 Ruochen Jiao, Xiangguo Liu, Bowen Zheng, Dave Liang, and Qi Zhu. Tae: A semi-supervised
 573 controllable behavior-aware trajectory generator and predictor. In *2022 IEEE/RSJ International*
 574 *Conference on Intelligent Robots and Systems (IROS)*, pp. 12534–12541. IEEE, 2022.

575

576 Ruya Karagulle, Nikos Aréchiga, Andrew Best, Jonathan DeCastro, and Necmiye Ozay. A safe
 577 preference learning approach for personalization with applications to autonomous vehicles. *IEEE*
 578 *Robotics and Automation Letters*, 9(5):4226–4233, 2024. doi: 10.1109/LRA.2024.3375626.

579

580 Rabeeh Karimi Mahabadi, Sebastian Ruder, Mostafa Dehghani, and James Henderson. Parameter-
 581 efficient multi-task fine-tuning for transformers via shared hypernetworks. In *Annual Meeting of*
 582 *the Association for Computational Linguistics*, 2021.

583

584 Dongchan Kim, Hyukju Shon, Nahyun Kweon, Seungwon Choi, Chanuk Yang, and Kunsoo Huh.
 585 Driving style-based conditional variational autoencoder for prediction of ego vehicle trajectory.
 586 *IEEE Access*, 9:169348–169356, 2021. doi: 10.1109/ACCESS.2021.3138502.

587

588 Sujeong Kim, Stephen J. Guy, Dinesh Manocha, and Ming C. Lin. Interactive simulation of dynamic
 589 crowd behaviors using general adaptation syndrome theory. In *Proceedings of the ACM SIGGRAPH*
 590 *Symposium on Interactive 3D Graphics and Games*, I3D ’12, pp. 55–62, New York, NY, USA, 2012.
 591 Association for Computing Machinery. ISBN 9781450311946. doi: 10.1145/2159616.2159626.
 592 URL <https://doi.org/10.1145/2159616.2159626>.

593

594 Sheila G Klauer, Thomas A Dingus, Vicki L Neale, Jeremy D Sudweeks, and David J Ramsey.
 595 Comparing real-world behaviors of drivers with high versus low rates of crashes and near crashes.
 596 Technical report, 2009.

594 Ziyue Li and Tianyi Zhou. Sparser mixture-of-adapters with cross-layer generalization. In
 595 Luis Chiruzzo, Alan Ritter, and Lu Wang (eds.), *Proceedings of the 2025 Conference of the*
 596 *Nations of the Americas Chapter of the Association for Computational Linguistics: Human*
 597 *Language Technologies (Volume 1: Long Papers)*, pp. 3988–4002, Albuquerque, New Mex-
 598 ico, April 2025. Association for Computational Linguistics. ISBN 979-8-89176-189-6. URL
 599 <https://aclanthology.org/2025.naacl-long.201/>.

600 Haokun Liu, Derek Tam, Muqeeth Mohammed, Jay Mohta, Tenghao Huang, Mohit Bansal, and Colin
 601 Raffel. Few-shot parameter-efficient fine-tuning is better and cheaper than in-context learning.
 602 In Alice H. Oh, Alekh Agarwal, Danielle Belgrave, and Kyunghyun Cho (eds.), *Advances in*
 603 *Neural Information Processing Systems*, 2022. URL <https://openreview.net/forum?id=rBCvMG-JsPd>.

604 Xulei Liu, Yafei Wang, Zhisong Zhou, Kanghyun Nam, Chongfeng Wei, and Chengliang Yin.
 605 Trajectory prediction of preceding target vehicles based on lane crossing and final points generation
 606 model considering driving styles. *IEEE Transactions on Vehicular Technology*, 70(9):8720–8730,
 607 2021. doi: 10.1109/TVT.2021.3098429.

608 Farzeen Munir, Tsvetomila Mihaylova, Shoaib Azam, Tomasz Piotr Kucner, and Ville Kyrki. Explor-
 609 ing large language models for trajectory prediction: A technical perspective. In *Companion of the*
 610 *2024 ACM/IEEE International Conference on Human-Robot Interaction*, HRI '24, pp. 774–778,
 611 New York, NY, USA, 2024. Association for Computing Machinery. ISBN 9798400703232. doi:
 612 10.1145/3610978.3640625. URL <https://doi.org/10.1145/3610978.3640625>.

613 Yi Lu Murphey, Robert Milton, and Leonidas Kiliaris. Driver's style classification using jerk analysis.
 614 In *2009 IEEE Workshop on Computational Intelligence in Vehicles and Vehicular Systems*, pp.
 615 23–28, 2009. doi: 10.1109/CIVVS.2009.4938719.

616 Nigamaa Nayakanti, Rami Al-Rfou, Aurick Zhou, Kratarth Goel, Khaled S. Refaat, and Benjamin
 617 Sapp. Wayformer: Motion forecasting via simple & efficient attention networks, 2022. URL
 618 <https://arxiv.org/abs/2207.05844>.

619 Bo Pang, Tianyang Zhao, Xu Xie, and Ying Nian Wu. Trajectory prediction with latent belief
 620 energy-based model. In *Proceedings of the IEEE/CVF conference on computer vision and pattern*
 621 *recognition*, pp. 11814–11824, 2021.

622 Edoardo Maria Ponti, Alessandro Sordoni, Yoshua Bengio, and Siva Reddy. Combining parameter-
 623 efficient modules for task-level generalisation. In Andreas Vlachos and Isabelle Augenstein (eds.),
 624 *Proceedings of the 17th Conference of the European Chapter of the Association for Computational*
 625 *Linguistics*, pp. 687–702, Dubrovnik, Croatia, May 2023. Association for Computational Lin-
 626 istics. doi: 10.18653/v1/2023.eacl-main.49. URL <https://aclanthology.org/2023.eacl-main.49/>.

627 Akshara Prabhakar, Yuanzhi Li, Karthik Narasimhan, Sham Kakade, Eran Malach, and Samy
 628 Jelassi. Lora soups: Merging loras for practical skill composition tasks, 2024. URL <https://arxiv.org/abs/2410.13025>.

629 Mariah L. Schrum, Emily Sumner, Matthew C. Gombolay, and Andrew Best. Maveric: A data-driven
 630 approach to personalized autonomous driving. *IEEE Transactions on Robotics*, 40:1952–1965,
 631 2024. doi: 10.1109/TRO.2024.3359543.

632 Shaoshuai Shi, Li Jiang, Dengxin Dai, and Bernt Schiele. Motion transformer with global intention
 633 localization and local movement refinement. *Advances in Neural Information Processing Systems*,
 634 2022.

635 Haowen Wang, Tao Sun, Congyun Jin, Yingbo Wang, Yibo Fan, Yunqi Xu, Yuliang Du, and
 636 Cong Fan. Customizable combination of parameter-efficient modules for multi-task learning.
 637 In *The Twelfth International Conference on Learning Representations*, 2024a. URL <https://openreview.net/forum?id=G1Hlubz1fR>.

638 Jifeng Wang, Kaouther Messaoud, Yuejiang Liu, Juergen Gall, and Alexandre Alahi. Forecast-
 639 peft: Parameter-efficient fine-tuning for pre-trained motion forecasting models. *arXiv preprint*
 640 *arXiv:2407.19564*, 2024b.

648 Xiao Wang, Ke Tang, Xingyuan Dai, Jintao Xu, Jinhao Xi, Rui Ai, Yuxiao Wang, Weihao Gu, and
 649 Changyin Sun. Safety-balanced driving-style aware trajectory planning in intersection scenarios
 650 with uncertain environment. *IEEE Transactions on Intelligent Vehicles*, 8(4):2888–2898, 2023.
 651 doi: 10.1109/TIV.2023.3239903.

652 Benjamin Wilson, William Qi, Tanmay Agarwal, John Lambert, Jagjeet Singh, Siddhesh Khandelwal,
 653 Bowen Pan, Ratnesh Kumar, Andrew Hartnett, Jhony Kaesemel Pontes, Deva Ramanan, Peter
 654 Carr, and James Hays. Argoverse 2: Next generation datasets for self-driving perception and
 655 forecasting. In *Proceedings of the Neural Information Processing Systems Track on Datasets and*
 656 *Benchmarks (NeurIPS Datasets and Benchmarks 2021)*, 2021.

657 Yang Xing, Chen Lv, and Dongpu Cao. Personalized vehicle trajectory prediction based on joint
 658 time-series modeling for connected vehicles. *IEEE Transactions on Vehicular Technology*, 69(2):
 659 1341–1352, 2020. doi: 10.1109/TVT.2019.2960110.

660 Pei Xu, Jean-Bernard Hayet, and Ioannis Karamouzas. Socialvae: Human trajectory prediction using
 661 timewise latents. In *European Conference on Computer Vision*, pp. 511–528. Springer, 2022. doi:
 662 10.1007/978-3-031-19772-7_30.

663 Prateek Yadav, Colin Raffel, Mohammed Muqeeth, Lucas Caccia, Haokun Liu, Tianlong Chen, Mohit
 664 Bansal, Leshem Choshen, and Alessandro Sordoni. A survey on model moerging: Recycling and
 665 routing among specialized experts for collaborative learning. *arXiv preprint arXiv:2408.07057*,
 666 2024.

667 Renteng Yuan, Mohamed Abdel-Aty, Qiaojun Xiang, Zijin Wang, and Xin Gu. A temporal multi-
 668 gate mixture-of-experts approach for vehicle trajectory and driving intention prediction. *IEEE*
 669 *Transactions on Intelligent Vehicles*, 9(1):1204–1216, 2024. doi: 10.1109/TIV.2023.3336310.

670 Ted Zadouri, Ahmet Üstün, Arash Ahmadian, Beyza Ermis, Acyr Locatelli, and Sara Hooker.
 671 Pushing mixture of experts to the limit: Extremely parameter efficient moe for instruction tuning.
 672 In *The Twelfth International Conference on Learning Representations*, 2024. URL <https://openreview.net/forum?id=EvDeiLv7qc>.

673 Laura Zheng, Sanghyun Son, Jing Liang, Xijun Wang, Brian Clipp, and Ming C Lin. Deep stochastic
 674 kinematic models for probabilistic motion forecasting in traffic. In *2024 IEEE/RSJ International*
 675 *Conference on Intelligent Robots and Systems (IROS)*, pp. 943–950. IEEE, 2024.

676 Laura Zheng, Hamidreza Yaghoubi Araghi, Tony Wu, Sandeep Thalapanane, Tianyi Zhou, and
 677 Ming C Lin. Quantifying and modeling driving styles in trajectory forecasting. *arXiv preprint*
 678 *arXiv:2503.04994*, 2025.

679

680

681

682

683

684

685

686

687

688

689

690

691

692

693

694

695

696

697

698

699

700

701

702 APPENDIX

703

704 A ANONYMIZED CODE

705

706 For an anonymized version of our code, please see: anonymous.4open.science/r/polysona

707

708 B WEIGHT-MIXING IMPLEMENTATION

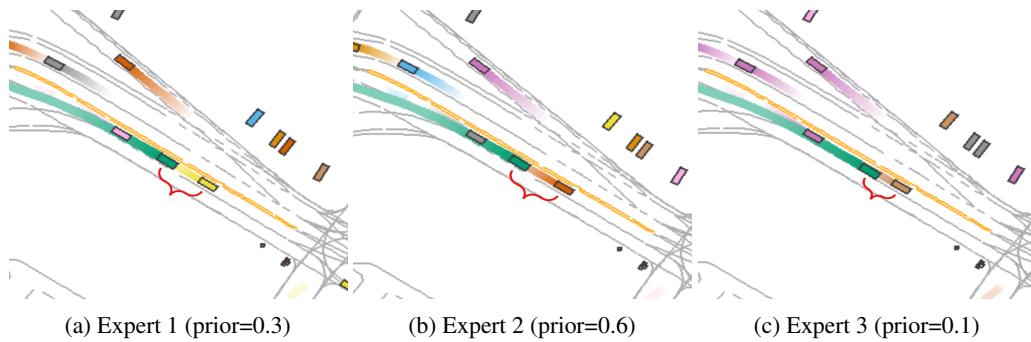
709

710 **PyTorch-Style Forward Pass of Weight-Mixing Mixture-of-LoRA Layer**

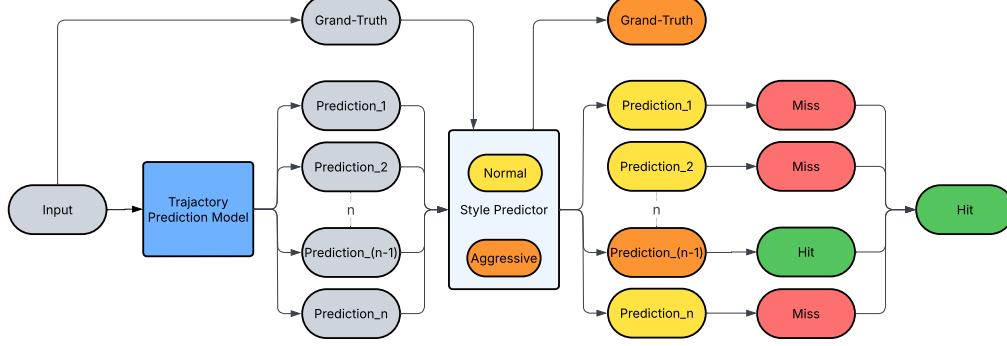
```

711
712     expert_alphas = router(...) # (b, num_experts)
713
714     expert_weight_A = (expert_alphas[..., None, None] * self.
715         expert_weight_A[None]).sum(
716             dim=1
717         ) # (b, r, in_features)
718     expert_weight_B = (expert_alphas[..., None, None] * self.
719         expert_weight_B[None]).sum(
720             dim=1
721         ) # (b, out_features, r)
722
723     output = torch.einsum(
724         "bi,bri->br", x, expert_weight_A
725     ) # (b, in_features) @ (b, r, in_features) -> (b, r)
726     output = torch.einsum(
727         "br,bor->bo", output, expert_weight_B
728     ) # (b, r) @ (b, out_features, r) -> (b, out_features)
729     output = self.dropout(output) # (b, out_features)
730
731     output = F.linear(x, w0, b0) + output # w0x + BAx + b0
732
733
734
735
736
737
738
739
740
741
742
743
744
745
746
747
748
749
750
751
752
753
754
755

```

756 C QUALITATIVE RESULTS ON ARGOVERSE
757

759
760
761
762
763
764
765
766
767
768
769
770 **Figure 5: Qualitative result: Rollouts on a lane change scenario with different experts on**
771 **Ours+CAT.** We show how different experts behave under the same lane change scenario, where
772 the goal is to lane change between two agents. While rollouts from Expert 1 and 2 are very similar
773 (maintaining at least one vehicle length from the leading vehicle), Expert 3, which has the lowest
774 prior probability, predicts a trajectory with much smaller headway distance. Animations of these
775 scenarios are better visualized on our project website, linked in the abstract.

776 D STYLE CONSISTENCY METRIC VISUALIZATION
777

779
780
781
782
783
784
785
786
787
788
789
790 **Figure 6: Style Consistency Metric.** Given a driving scenario (“Input”), the black-box trajectory
791 predictor generates n candidate futures ($\text{Prediction}_1, \dots, \text{Prediction}_n$). A learned style predictor
792 then assigns each candidate—and the true future (“Ground-Truth”—to one of two clusters (e.g.
793 “Normal” vs. “Aggressive”). If at least one of the n predicted trajectories shares the same style
794 label as the ground-truth, the sample is marked a *Hit*; otherwise it is a *Miss*. This hit/miss outcome
795 directly measures whether the model’s multi-modal outputs *cover* the driver’s actual style, beyond
796 conventional displacement errors.

797 E STYLE CONSISTENCY METRIC
798

801 To explicitly measure a model’s ability to cover the correct driving style, as shown in Figure 6, we
802 propose the *Style Miss Rate* a style consistency metric based on kinematic clustering following the
803 clustering methodology introduced in (Zheng et al., 2025), and a hit/miss criterion:

804 **Extract kinematic statics:** For each trajectory τ , compute a feature vector

$$\phi(\tau) = [\max_t |a(t)|, \text{Var}(a), \text{Var}(v), \gamma]^T \in \mathbb{R}^d, \quad (8)$$

805 where a is acceleration (with $\max_t |a(t)|$ denoting the peak absolute acceleration over the trajectory,
806 i.e. the highest instantaneous acceleration magnitude), v is speed, and

810

$$\gamma = \frac{\text{Var}(j(t))}{\mathbb{E}[j(t)]}$$

813 is the jerk-variance ratio as defined in Murphrey et al. (2009).

814 **Learn style clusters:** Fit a Gaussian Mixture Model (GMM) with $k = 2$ on the set of all ground-truth
815 kinematic embeddings in the evaluation set $\{\phi(\tau_n^*)\}_{n=1}^N$, yielding a cluster assignment function

817

$$818 \quad C(\phi) \in \{"\text{normal}", "\text{aggressive}"\}. \quad (9)$$

819

820 Normal and aggressive are assigned based on the mean speed of each cluster. A cluster with higher
821 mean speed will be assigned as aggressive.

822 **Assign styles to predictions:** For each sample n , let $\{\hat{\tau}_{n,i}\}_{i=1}^6$ be the six predicted trajectories.
823 Define

$$824 \quad s_n^* = C(\phi(\tau_n^*)), \quad s_{n,i} = C(\phi(\hat{\tau}_{n,i})). \quad (10)$$

825

826 **Define hit/miss:** A *hit* occurs if at least one predicted style matches the ground-truth style:

$$827 \quad \text{Hit}_n = \{\exists i : s_{n,i} = s_n^*\}, \quad \text{Miss}_n = 1 - \text{Hit}_n. \quad (11)$$

828

829 **Style Miss Rate:** The overall metric is

830

$$831 \quad \text{SMR} = \frac{1}{N} \sum_{n=1}^N \text{Miss}_n. \quad (12)$$

832

833 By construction, the SMR goes beyond pure spatial accuracy: it measures whether the model
834 “covers” the driver’s true style among its multi-modal outputs. A style-agnostic predictor may achieve
835 low ADE/FDE by clustering its modes around average behavior, but will incur a high miss rate
836 on aggressive samples. In contrast, a style-aware model—conditioned on inferred driving-style
837 embeddings—should include at least one candidate trajectory whose kinematics align with the true
838 style, yielding a lower SMR.

839

840

841

842

843

844

845

846

847

848

849

850

851

852

853

854

855

856

857

858

859

860

861

862

863

864 F STYLE MISS RATE EVALUATION ON ABLATION VARIANTS
865

866 To further assess how each component of our mixture-of-experts framework contributes to style
867 coverage, we compute the Style Miss Rate (SMR) on the same ablation variants presented in Table 4.
868 That is, for each model variant—removing reconstruction, KL loss, entropy regularization, etc.—we
869 evaluate how often none of its multi-modal predictions match the true driving style cluster. The
870 resulting SMR values are reported in Table 6. This analysis shows that the ablations which most
871 degrade traditional error metrics (e.g. reconstruction and KL removal) also incur the largest increases
872 in SMR, indicating a direct link between component contributions and the model’s ability to cover
873 the driver’s style.

874
875 Table 6: **Style Miss Rate (SMR) for each ablation variant.**

Ours	SMR \downarrow
+LinearRouter	0.2246
+Full Finetuning	0.2033
-Social Forces	0.2229
-Context Features	0.2236
-KL Loss	0.2267
-Reconstruction	0.2263
-Expert Entropy Loss	0.2260

918 **G HOW DOES THE MODEL REASON? TAKING A LOOK AT THE SALIENCY
919 MAPS.**

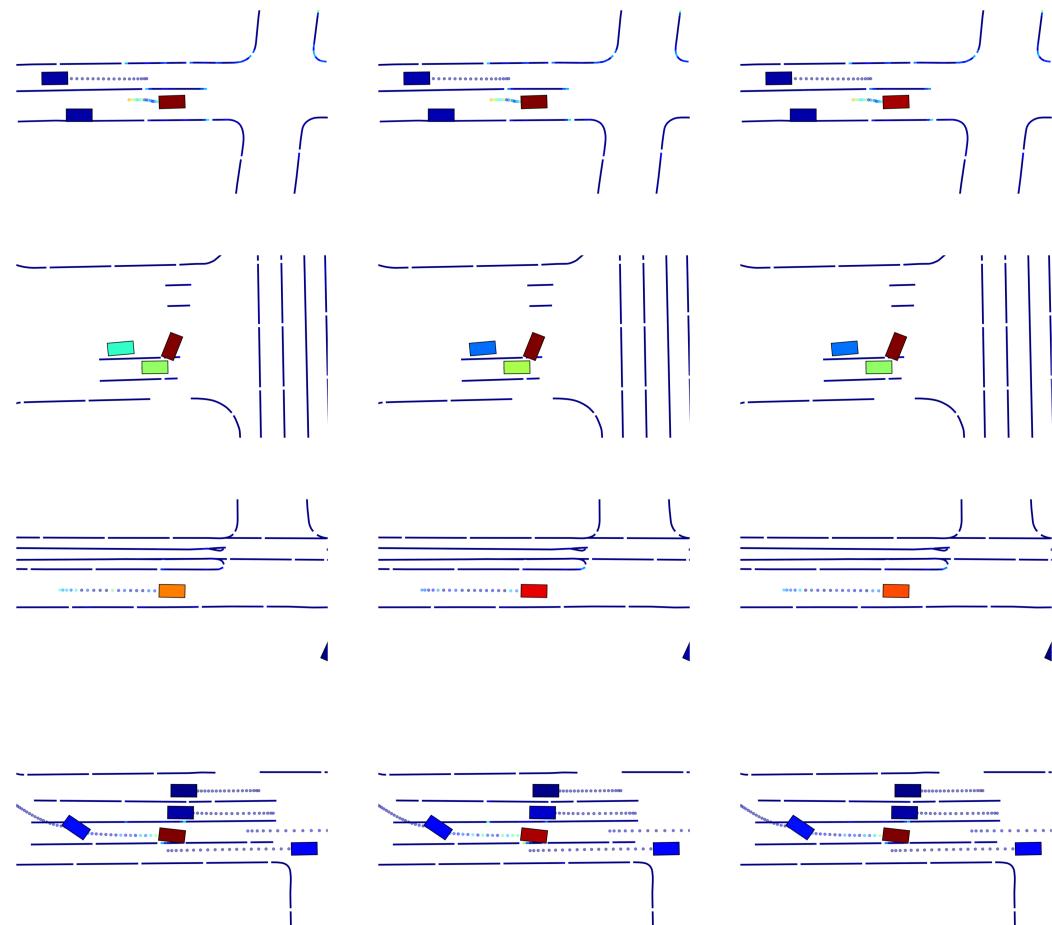
921 To gain insight into which input features the model attends when forecasting agent trajectories, we
922 compute saliency maps by measuring the sensitivity of the most likely predicted trajectory with
923 respect to each input feature. Formally, let $X = \{A_{in}, M_{in}\}$ denote the concatenation of historical
924 object trajectories A_{in} and map polylines M_{in} . If \hat{p}_τ is the probability of trajectory τ , then let
925 $\hat{\tau}^* = \arg \max_\tau \hat{p}_\tau$ be the most likely predicted trajectory after a forward pass from a model. Then,
926 we compute the gradient

$$\nabla_X \hat{\tau}^* = \frac{\partial \hat{\tau}^*}{\partial X}$$

927 via back-propagation, and form the saliency map
928

$$S(X) = \log(\|\nabla_X \hat{\tau}^*\|_2 + 1).$$

929 We use logarithmic scaling above to better display nuances in smaller gradient magnitudes. For
930 visualizing this saliency map, we render the map polylines and the agents' historical trajectories
931 colored using $S(X)$ on the "jet" color scheme (dark blue to green to dark red). Warmer colors
932 highlight map segments or agent trajectories that the model deems more important for predicting
933 future trajectories. Likewise, lighter colors highlight areas of less importance. Each agent vehicle
934 is also colored on the same scale based on the maximum saliency value of its historical trajectory. We
935 generate this visualization for MTR+Actions, Ours+CAT, and Ours+MoV:
936



970 **Figure 7: Visualization of scenario feature saliency (Continued on the next page).** Saliency
971 maps are visualized for 10 randomly selected scenarios from Argoverse on MTR+Actions (Left),
972 Ours+CAT (Middle), and Ours+MoV (Right).

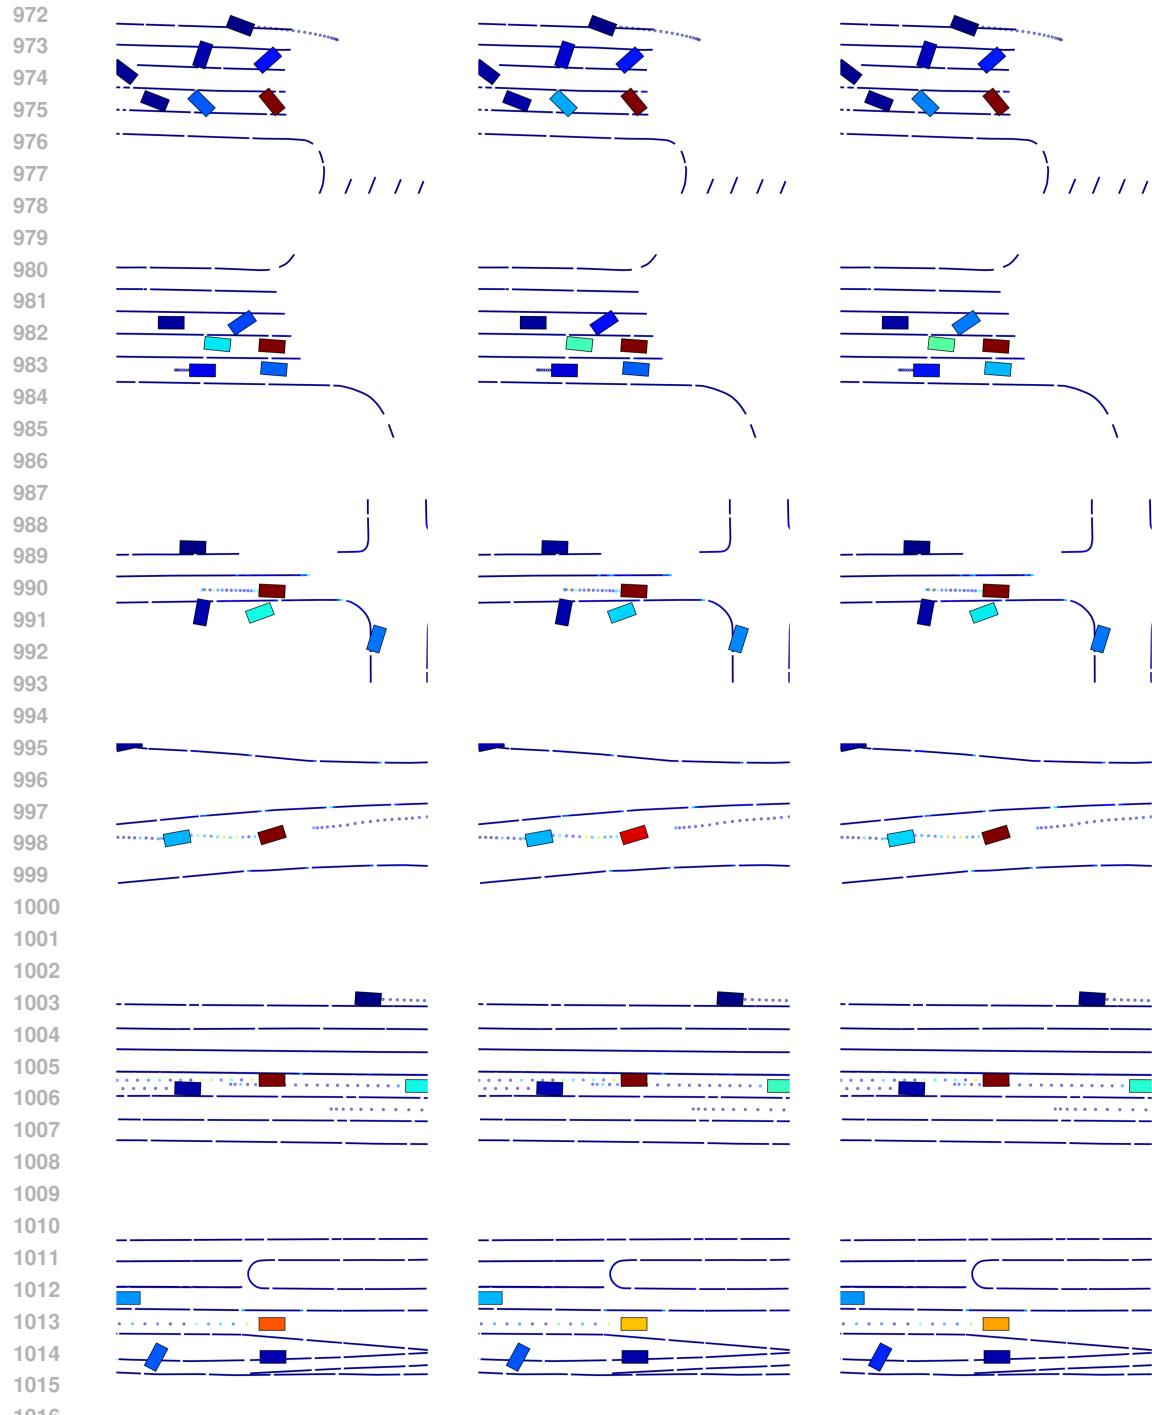
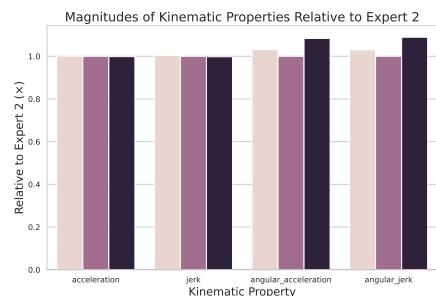
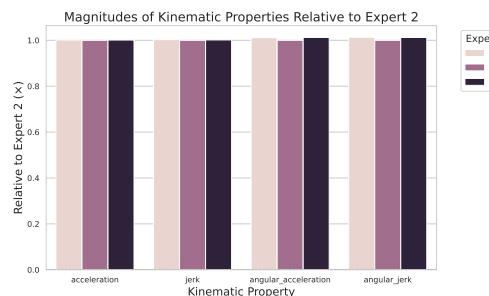


Figure 7: **Visualization of scenario feature saliency.** Saliency maps are visualized for 10 randomly selected scenarios from Argoverse on MTR+Actions (Left), Ours+CAT (Middle), and Ours+MoV (Right). The top 2 scenarios feature mostly stationary agents. In each scenario and model, the ego vehicle has the most saliency (colored red), followed by nearby agents and map polylines.

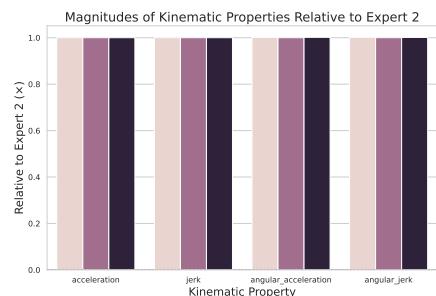
1026 H ADDITIONAL VISUALIZATIONS FOR RELATIVE KINEMATICS BY EXPERT



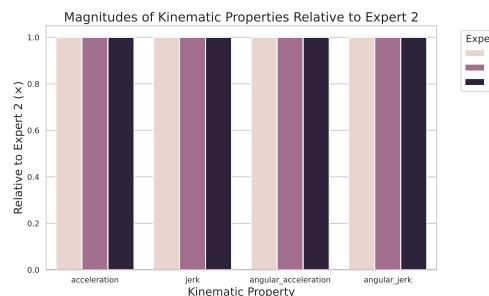
(a) C-Poly



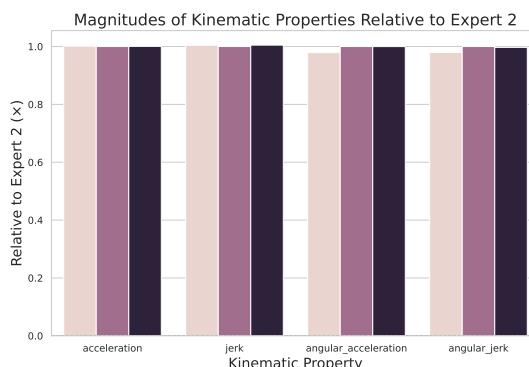
(b) HyperFormer



(c) Polytropon



(d) MoV



(e) CAT

Figure 8: **Kinematic magnitude comparisons for all variants of our approach.** Experiments run with seed 0 are plotted.

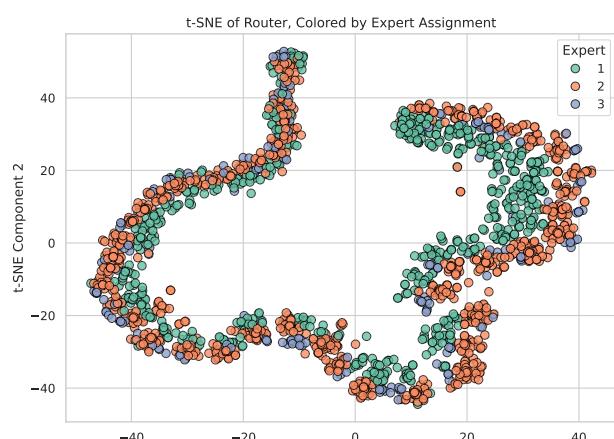
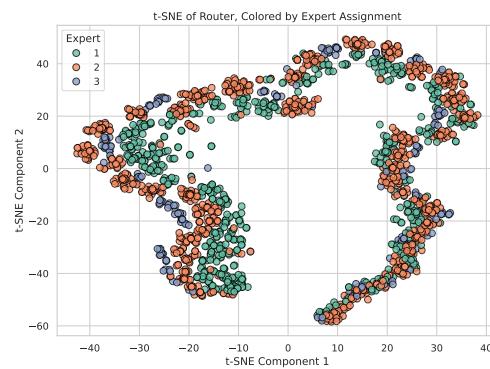
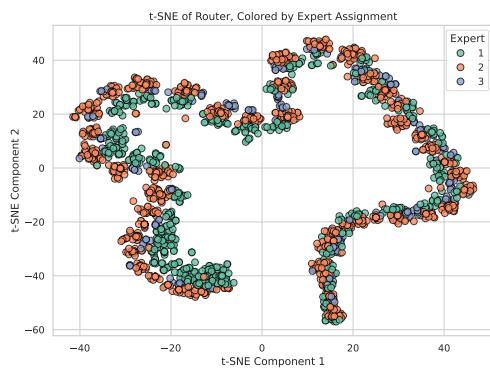
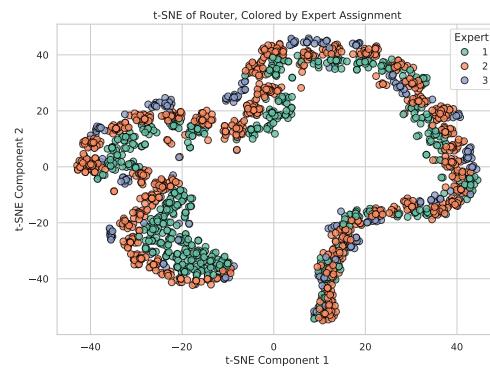
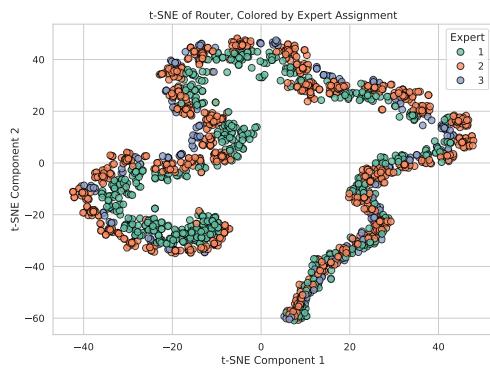
1080 I ADDITIONAL T-SNE PLOTS FOR ROUTER EMBEDDINGS BY EXPERT
1081
1082

Figure 9: **t-SNE visualization of router embeddings for all variants of our approach.** Experiments run with seed 0 are plotted.

1134 **J BASELINE MTR MODEL HYPERPARAMETERS**
1135
1136
11371138 Table 7: Hyperparameters used for training the MTR baseline model.
1139

Hyperparameter	Value
Context Encoder	
NAME	MTREncoder
NUM_OF_ATTN_NEIGHBORS	7
NUM_INPUT_ATTR_AGENT	39
NUM_INPUT_ATTR_MAP	29
NUM_CHANNEL_IN_MLP_AGENT	256
NUM_CHANNEL_IN_MLP_MAP	64
NUM_LAYER_IN_MLP_AGENT	3
NUM_LAYER_IN_MLP_MAP	5
NUM_LAYER_IN_PRE_MLP_MAP	3
D_MODEL	256
NUM_ATTN_LAYERS	6
NUM_ATTN_HEAD	8
DROPOUT_OF_ATTN	0.1
USE_LOCAL_ATTN	True
Motion Decoder	
NAME	MTRDecoder
NUM_MOTION_MODES	6
D_MODEL	512
NUM_DECODER_LAYERS	6
NUM_ATTN_HEAD	8
MAP_D_MODEL	256
DROPOUT_OF_ATTN	0.1
NUM_BASE_MAP_POLYLINES	256
NUM_WAYPOINT_MAP_POLYLINES	128
LOSS_WEIGHTS.cls	1.0
LOSS_WEIGHTS.reg	1.0
LOSS_WEIGHTS.vel	0.5
NMS_DIST_THRESH	2.5
Training	
max_epochs	40
learning_rate	0.0001
learning_rate_sched	[22, 24, 26, 28]
optimizer	AdamW
scheduler	lambdaLR
grad_clip_norm	1000.0
weight_decay	0.01
lr_decay	0.5
lr_clip	0.000001
WEIGHT_DECAY	0.01
train_batch_size	64
eval_batch_size	64
Data	
max_num_agents	64
map_range	100
max_num_roads	768
max_points_per_lane	20
manually_split_lane	True
point_sampled_interval	1
num_points_each_polyline	20
vector_break_dist_thresh	1.0
predict_actions	True

1188 **K POLYSONA MODEL HYPERPARAMETERS**
11891190
1191 Table 8: Hyperparameters used for training the PolySona models. Rows highlighted in yellow
1192 indicate differences from the baseline MTR configuration.

1193	1194	Hyperparameter	Value
1195	Context Encoder		
1196	NAME	MTREncoder	
1197	NUM_OF_ATTN_NEIGHBORS	7	
1198	NUM_INPUT_ATTR_AGENT	39	
1199	NUM_INPUT_ATTR_MAP	29	
1200	NUM_CHANNEL_IN_MLP_AGENT	256	
1201	NUM_CHANNEL_IN_MLP_MAP	64	
1202	NUM_LAYER_IN_MLP_AGENT	3	
1203	NUM_LAYER_IN_MLP_MAP	5	
1204	NUM_LAYER_IN_PRE_MLP_MAP	3	
1205	D_MODEL	256	
1206	NUM_ATTN_LAYERS	6	
1207	NUM_ATTN_HEAD	8	
1208	DROPOUT_OF_ATTN	0.1	
1209	USE_LOCAL_ATTN	True	
Motion Decoder			
1210	NAME	PolySonaDecoder	
1211	NUM_MOTION_MODES	6	
1212	INTENTION_POINTS_FILE	cluster_64_center_dict_6s.pkl	
1213	D_MODEL	512	
1214	NUM_DECODER_LAYERS	6	
1215	NUM_ATTN_HEAD	8	
1216	MAP_D_MODEL	256	
1217	DROPOUT_OF_ATTN	0.1	
1218	NUM_BASE_MAP_POLYLINES	256	
1219	NUM_WAYPOINT_MAP_POLYLINES	128	
1220	LOSS_WEIGHTS.cls	1.0	
1221	LOSS_WEIGHTS.reg	1.0	
1222	LOSS_WEIGHTS.vel	0.5	
1223	NMS_DIST_THRESH	1.0	
Training			
1224	max_epochs	10	
1225	learning_rate	0.001	
1226	learning_rate_sched	[22, 24, 26, 28]	
1227	optimizer	AdamW	
1228	scheduler	polynomialLR (power=2)	
1229	grad_clip_norm	1000.0	
1230	weight_decay	0.00	
1231	lr_decay	0.5	
1232	lr_clip	0.000001	
1233	train_batch_size	256	
1234	eval_batch_size	256	
1235	predict_actions	True	
1236	lora_rank	4	
1237	freeze_encoder	True	
1238	freeze_decoder	True	
1239	attention_only	False	
1240	num_personas	3	
1241	prior	[0.3, 0.6, 0.1]	
1242	λ_{recon}	50	
1243	λ_{KL}	50	
1244	λ_{entropy}	25	
1245	seed	0 / 1 / 2	
Data			
1246	max_num_agents	64	
1247	map_range	100	
1248	max_num_roads	768	
1249	max_points_per_lane	20	
1250	manually_split_lane	True	
1251	point_sampled_interval	1	
1252	num_points_each_polyline	20	
1253	vector_break_dist_thresh	1.0	

1242 L IMPACT OF RANK ON PERFORMANCE
12431244 Table 9: Comparison of Ours+CAT Across Different Ranks.
1245

1246

Rank	brierFDE \downarrow	minADE \downarrow	minFDE \downarrow	MissRate \downarrow
2	2.1593	0.8573	1.7059	0.3151
4	2.1607	0.8624	1.7041	0.3171
8	2.1578	0.8578	1.7042	0.3120
16	2.1668	0.8610	1.7102	0.3151

1251
1252 Table 10: Comparison of Ours+CAT Across Different Ranks, Grouped by Kalman Difficulty
1253 and TDBM Driving Styles.
1254

1255

Rank	Kalman Difficulty			TDBM Driving Styles			
	Easy	Medium	Hard	Timid	Careful	Reckless	Threatening
2	0.8120	1.1675	3.9875	0.8903	0.8865	0.8577	0.8172
4	0.8188	1.1678	2.4978	0.8793	0.8477	0.8655	0.8161
8	0.8130	1.1641	3.9882	0.8913	0.9833	0.8581	0.8183
16	0.8149	1.1758	4.2696	0.8944	0.8858	0.8613	0.8225

1261
1262
1263
1264
1265
1266
1267
1268
1269
1270
1271
1272
1273
1274
1275
1276
1277
1278
1279
1280
1281
1282
1283
1284
1285
1286
1287
1288
1289
1290
1291
1292
1293
1294
1295

1296 M STANDARD DEVIATION TABLE
12971298 Table 11: **Standard Deviation of Trajectory Prediction Benchmark Performance Comparisons.**
1299

Method	brierFDE \downarrow	minADE \downarrow	minFDE \downarrow	MissRate \downarrow
Ours+Polytron Ponti et al. (2023)	0.0004	0.0004	0.0004	0.0009
Ours+C-Poly Wang et al. (2024a)	0.0026	0.0011	0.0025	0.0003
Ours+HyperFormer Karimi Mahabadi et al. (2021)	0.0017	0.0004	0.0017	0.0010
Ours+CAT Prabhakar et al. (2024)	0.0019	0.0012	0.0018	0.0010
Ours+MoV Zadouri et al. (2024)	0.0010	0.0007	0.0010	0.0004

1306
1307 Table 12: **Standard Deviation of minADE Comparison by Kalman Difficulty and TDBM Driving
1308 Style groups.**
1309

Method	Kalman Difficulty			TDBM Driving Styles			
	Easy	Medium	Hard	Timid	Careful	Reckless	Threatening
Ours+PolyTropo Zadouri et al. (2024)	0.0003	0.0032	0.0040	0.0005	0.0038	0.0004	0.0005
Ours+C-Poly Wang et al. (2024a)	0.0013	0.0003	0.0146	0.0009	0.0286	0.0012	0.0011
Ours+HyperFormer Karimi Mahabadi et al. (2021)	0.0006	0.0013	0.0060	0.0006	0.0330	0.0004	0.0007
Ours+CAT Prabhakar et al. (2024)	0.0012	0.0051	0.0682	0.0014	0.0340	0.0012	0.0012
Ours+MoV Zadouri et al. (2024)	0.0007	0.0006	0.0041	0.0008	0.0026	0.0007	0.0006

1350

Table 13: Hyperparameter sweep of # of classes used in PolySona training.

1351

1352

1353

1354

1355

1356

1357

1358

1359

1360

1361

1362

1363

1364

1365

1366

1367

1368

1369

1370

1371

1372

1373

1374

1375

1376

1377

1378

1379

1380

1381

1382

1383

1384

1385

1386

1387

1388

1389

1390

1391

1392

1393

1394

1395

1396

1397

1398

1399

1400

1401

1402

1403

N HYPERPARAMETER SWEET: NUMBER OF LATENT CLASSES MODELED