

Interaction-Merged Motion Planning: Effectively Leveraging Diverse Motion Datasets for Robust Planning

Giwon Lee^{1*}, Wooseong Jeong^{1*}, Daehee Park², Jaewoo Jeong¹, and Kuk-Jin Yoon¹

¹Visual Intelligence Lab., KAIST, Korea

²Intelligent Systems and Learning Lab., DGIST, Korea

Abstract

Motion planning is a crucial component of autonomous robot driving. While various trajectory datasets exist, effectively utilizing them for a target domain remains challenging due to differences in agent interactions and environmental characteristics. Conventional approaches, such as domain adaptation or ensemble learning, leverage multiple source datasets but suffer from domain imbalance, catastrophic forgetting, and high computational costs. To address these challenges, we propose **Interaction-Merged Motion Planning (IMMP)**, a novel approach that leverages parameter checkpoints trained on different domains during adaptation to the target domain. IMMP follows a two-step process: pre-merging to capture agent behaviors and interactions, sufficiently extracting diverse information from the source domain, followed by merging to construct an adaptable model that efficiently transfers diverse interactions to the target domain. Our method is evaluated on various planning benchmarks and models, demonstrating superior performance compared to conventional approaches.

1. Introduction

Motion planning [7, 53, 80, 81] is the final step in autonomous robot driving, focused on generating safe and efficient trajectory based on perception and prediction. With the availability of various trajectory datasets [4, 6, 23, 36, 42, 50, 52, 73], data-driven approaches have become more popular due to their ability to adapt to diverse driving scenarios [19–21, 24, 32, 61, 74].

However, existing trajectory datasets exhibit distinct characteristics, including differences in scene composition, agent interaction tendencies, and the difficulties in data acquisition. Indoor datasets are constrained in movement due to static structures, where robots generally move at slower

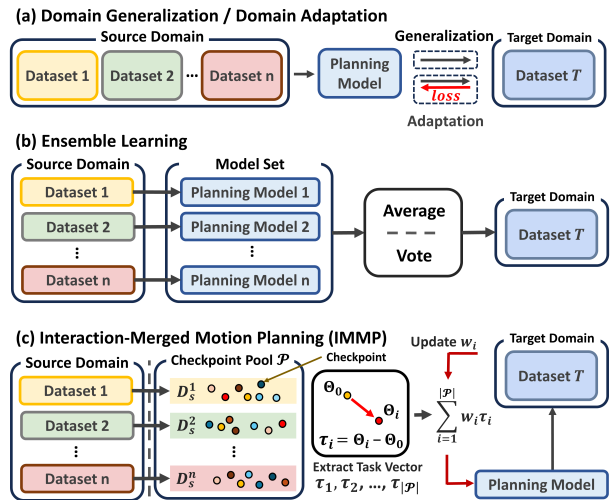


Figure 1. Comparison of different approaches for leveraging the source domains for target domain adaptation. (a) Domain generalization or adaptation directly utilizes relationships within datasets. (b) Ensemble learning combines predictions from multiple planning models, each trained on different datasets. (c) The proposed IMMP collects parameter checkpoints Θ and merges task vectors τ , which capture differences between trained and initial parameters, to improve adaptation to the target domain.

speeds. In contrast, outdoor datasets encompass unconstrained environments and results in a more dynamic maneuver. To account for these diverse environments, Human-Human Interaction (HHI) datasets and Human-Robot Interaction (HRI) datasets are collectively utilized to train a planner model. However, using HHI datasets for planning assumes human locomotion as ego-agent motion, therefore incorporating interaction dynamics that differ from HRI datasets where real robot-human interactions shape agent behaviors. These inherent differences across datasets collectively lead to variations in the nature of agents' maneuvers and their interactions. Holistically accounting for such disparities is crucial towards enhancing the overall adaptability in diverse real-world scenarios.

*Equal contribution to this work.

Our source code is available at: <https://github.com/wooseong97/IMMP>

One straightforward approach is to incorporate multiple source datasets into the training set (Fig. 1 (a)). Training on multiple datasets acquired from disparate environments can leverage richer patterns compared to training on a single dataset [16, 86]. However, naively increasing the number of source datasets fails to incorporate the inherent differences between datasets due to two aspects. First, the most effective source datasets for improving performance on a given target domain are unknown before actually training and fine-tuning a model. Differences in dataset properties—such as agent behaviors, interaction patterns, and environment characteristics—can lead to conflicts, making direct dataset composition non-trivial. These disparities in dataset can cause what we refer to as the *domain imbalance problem*, where certain datasets disproportionately affect the learning process, dominating model updates and affecting different modules or parameters within the planning model. Second, domain imbalance can lead to knowledge interference, similar to catastrophic forgetting in continual learning [34, 55, 57, 58], where newly introduced domains may override previously learned information.

Another approach is to integrate multiple planner models (Fig. 1 (b)). Despite the improvement in generalization on different domains, ensemble approaches require multiple planner models during inference and therefore increases the computational requirement. These challenges suggest that simply aggregating diverse motion datasets or using multiple models are both complex and inefficient approaches towards enhancing motion planning adaptability.

In response, we leverage model merging [11, 45, 68, 69, 72] to integrate knowledge from multiple sources to improve generalization while retaining computational cost. Specifically, model parameters from different checkpoints, each fine-tuned on a specific dataset, are merged. Each model’s parameters encode dataset-specific motion patterns and interaction dynamics, and model merging effectively combines this multi-dataset knowledge to adapt to a new target dataset. However, existing merging methods struggle in motion planning due to two key limitations. First, they fail to construct a sufficient parameter checkpoint pool, as they do not account for diverse evaluation metrics and multi-objective loss functions. Second, they overlook the hierarchical nature of motion planning features, potentially disrupting feature structures or discarding critical information. To address these limitations, we propose a **motion-planning-targeted merging method** that effectively incorporates the unique characteristics of motion datasets.

Our key insight is that transferring agent behaviors and interactions is crucial for motion planning adaptability. Therefore, our method prioritizes preserving and transferring these interactions during the merging process as shown in Fig. 1 (c). Specifically, our approach consists of two main steps. First, during the pre-merging phase, we sep-

arately train motion planning models on diverse datasets. From these models, we construct a pool of metric-specific optimal parameter checkpoints, each maintaining domain-specific interaction patterns. Second, in the merging phase, the checkpoints are merged into a newly initialized model by trainable weights. We partition the planning model parameters into key modules that each encapsulate different components of human-robot interaction: human encoder, robot encoder, interaction encoder, and decoder. In doing so, we divide-and-generalize the heterogeneous nature of human locomotion and interaction. As a result, our proposed Interaction-Merged Motion Planning (IMMP) can mitigate the issues of domain imbalance and catastrophic forgetting with lower costs.

Our main contributions are as follows:

- We propose IMMP, a novel method that merges parameter checkpoints from different domains to enhance adaptation, efficiency, and practicality.
- A two-step process extracts diverse motion features and transfers them by identifying key planning parameters, ensuring seamless integration of source domain information into the target domain.
- Evaluations on various benchmarks show IMMP’s superior adaptability over *ensemble learning*, *domain adaptation*, *domain generalization*, and other *merging-based approaches*, while integrating into existing frameworks.

2. Related Works

2.1. Motion Planning

Motion planning methods mainly utilize grid-based and sampling-based methods. In grid-based methods, the scene is represented as a grid, and planning is performed using search algorithms such as Dijkstra’s or A* algorithm [1, 5, 54, 60, 77]. Sampling-based methods generate random samples and connect them into tree or graph structures to find the optimal plan [3, 30, 31, 35, 62]. Recently, deep learning-based approaches have become increasingly popular due to their adaptability to diverse scenarios. Earlier deep learning-based planning studies adopted a sequential structure, where perception, prediction, and planning were studied independently [9, 13, 47, 64, 70, 74, 83, 84]. Recent trends move towards jointly training forecasters [19–21, 32] or even integrating all components of the autonomous driving model into the learning process [8, 18, 27, 28, 46, 65, 67, 82, 85]. However, such data-driven approaches are inherently vulnerable to dataset biases.

2.2. Leveraging Diverse Motion Datasets

Motion datasets differ in data collection methods, environmental characteristics, dataset size, and diversity. As a result, imitation learning-based models trained for a specific domain often experience performance degradation when de-

ployed in different domains [16, 17, 48, 49]. To improve robustness in target domains, recent studies have investigated domain generalization at both the dataset and architecture levels. UniTraj [16] adopts a dataset-level approach by integrating multiple trajectory datasets into a unified dataset, demonstrating that incorporating diverse data enhances generalization. Ye et al. [78] enhance domain generalization by partitioning a single dataset based on vehicle motion characteristics and representing motion in a reference path-based frame. Wang et al. [66] introduce a module that refines domain-specific velocity and environmental characteristics, while Dong et al. [14] propose a Stepwise Attention Layer for domain alignment at each timestep. Model architecture-level approaches, such as those in [14, 51, 63, 66, 71, 78], are tailored to specific domains or models, limiting their applicability to general models. In contrast, IMMP operates independently of model architecture, making it adaptable across various planning models.

2.3. Model Merging

Model merging [76] integrates parameters from independently trained models to construct a unified model without requiring access to the original training data. This approach can sometimes replace conventional ensemble learning [2, 37] or multi-task learning [25, 33, 40, 44, 56, 79] by leveraging only checkpointed parameters. Model merging techniques can be broadly categorized into *Pre-Merging* and *During Merging*. Pre-Merging methods [29, 41, 45, 59] fine-tune networks before merging to better align their parameters. Fine-tuning within the tangent space further improves weight disentanglement [45], enhancing task arithmetic. During Merging methods [11, 22, 68, 69] focus on how parameters are combined. Task arithmetic [22], which adjusts model accuracy by adding or removing task vectors representing parameter differences, plays a crucial role in this process. Ties-Merging [72] addresses conflicts between models through its trim, elect-sign, and merge process. Despite these advancements, model merging remains largely unexplored in motion planning. Additionally, naive adoption of existing merging approaches has yielded poor results, motivating our proposed merging procedure specifically designed for motion planning.

3. Method

3.1. Problem Definition and Motivation

We define samples from the input space \mathbb{X} as $X \in \mathbb{X}$ and from the output space \mathbb{Y} as $Y \in \mathbb{Y}$. In the motion planning task, the input space is represented as $X = \{\mathcal{X}_{ego}^{-T_{obs}:0}, \mathcal{X}_{surr}^{-T_{obs}:0}, \hat{\mathcal{Y}}_{surr}^{1:T_{fut}}\}$, where $\mathcal{X}_{ego}^{-T_{obs}:0}$ and $\mathcal{X}_{surr}^{-T_{obs}:0}$ are the past trajectory of the ego and surrounding agents, and $\hat{\mathcal{Y}}_{surr}^{1:T_{fut}}$ represents the future predicted trajec-

ries of the surrounding agents. We use $\hat{\cdot}$ to denote prediction. The output space is $Y = \{\mathcal{Y}_{ego}^{1:T_{fut}}\}$, which represents the future plan of the ego agent. T_{obs} and T_{fut} respectively represent the observed timesteps of past trajectories and future timesteps of final plan.

The source domains are represented as $\mathcal{S} = \{D_s^1, \dots, D_s^n\}$, where n is the number of source domains. The target domains, which we aim to adapt to, are denoted as $\mathcal{T} = \{D_t^1, \dots, D_t^m\}$, where m is the number of target domains. Each dataset is expressed as $D_s^i = \{(X_j^{s,i}, Y_j^{s,i})\}_{j=1}^{l_i}$ for the source domain and $D_t^i = \{(X_j^{t,i}, Y_j^{t,i})\}_{j=1}^{l_i}$ for the target domain, where i refers to the i -th dataset, and l_i is the number of samples within it. The objective of this work is to leverage \mathcal{S} to enhance the performance of the planning model on \mathcal{T} .

Due to the varying characteristics of existing motion planning datasets, domain gaps between source and target domains are more pronounced. Conventional methods such as domain adaptation [16] and domain generalization [17] directly utilize source domain samples, incorporating diverse datasets during adaptation. However, these approaches face limitations, as dataset sizes vary significantly, making it difficult to balance their influence during training. Furthermore, this imbalance can also lead to *catastrophic forgetting*, where knowledge from certain domains dominates, causing the model to forget information from other domains. Even with sufficient computational resources, selecting beneficial samples from the source domain \mathcal{S} for adaptation to the target domain \mathcal{T} remains a challenging task. This requires assessing the impact of each sample or domain on performance, which adds complexity to the learning process.

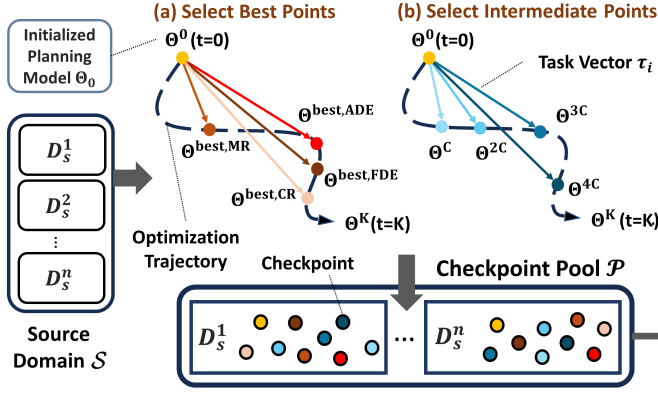
Instead, we propose new adaptation methods that divide the adaptation process into two stages, considering the nature of motion planning: collecting parameter checkpoints (Sec. 3.3) and merging them to transfer information across domains (Sec. 3.4). The overall IMMP framework is introduced in Fig. 2.

3.2. Model Merging and Task Vector

Model merging [76] is used to integrate the parameters of multiple network backbones, each trained on different tasks, and serves as an alternative to traditional multi-task learning [25, 26, 38, 39, 79] by avoiding joint training. This technique is primarily explored in language-image tasks [45, 59, 75] and image classification [22, 72]. Notably, Task Arithmetic [22] is a pivotal work in this domain, introducing the concept of the Task Vector, which enables simple parameter editing by adding or negating targeted information. The task vector is defined as follows:

Definition 1 (Task Vector [22]) *The task vector τ_i for task i is $\tau_i = \Theta_i - \Theta_0$, where Θ_0 and Θ_i are the parameters of the pre-trained and fine-tuned models, respectively.*

Sec 3.3. Interaction Conserving Pre-Merging



Sec 3.4. Interaction Transfer with Merging

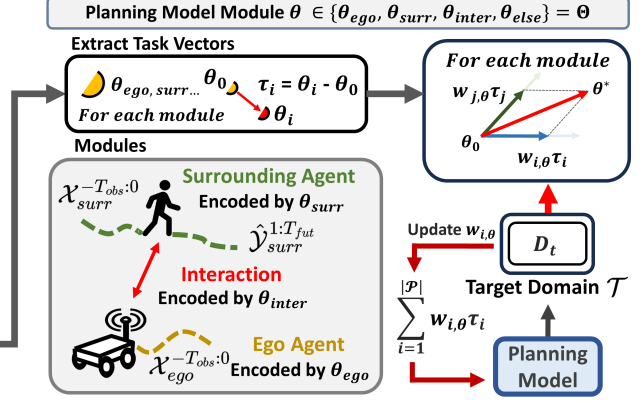


Figure 2. Overview of the proposed IMMP framework. In interaction-conserving pre-merging, parameter checkpoints are selected to preserve the distinct characteristics of different motion planning datasets. In interaction transfer with merging, task vectors are extracted for different modules in the planning model using the collected checkpoint pool \mathcal{P} . These task vectors are then merged by learning their linear weights w to enhance adaptation to the target domain.

The model parameters Θ can then be expressed as a combination of the initial parameters Θ_0 and the task vector set, formulated as follows:

$$\Theta = \Theta_0 + \lambda \sum_{i=1}^{|\mathcal{P}|} w_i \cdot \tau_i \quad (1)$$

where w_i represents the weight of the task vector τ_i , λ is a scaling factor, and $|\mathcal{P}|$ is the number of parameter checkpoints in the pool. The objective of merging is to determine the optimal parameter Θ for the target domain set $\mathcal{T} = \{D_t^1, \dots, D_t^m\}$, formulated as:

$$\Theta^* = \arg \min_{\Theta} \sum_{i=1}^m \sum_{j=1}^{l_i} \mathcal{L}(\Theta, X_j^{t,i}, Y_j^{t,i}) \quad (2)$$

Here, m denotes the number of target domains, and l_i represents the number of data samples in the i -th target dataset. Since Θ_0 and τ_i in Eq. (1) are fixed values, merging methods directly optimize λ and the weight set $\{w_i\}_{i=1}^{|\mathcal{P}|}$ using train loss. Initial model parameters Θ_0 can be a model parameter from scratch or one of any fine-tuned models parameters $\{\Theta_i\}_{i=1}^{|\mathcal{P}|}$. Since the task vector is extracted and linearly weighted during merging, the choice of initial parameters does not drastically affect the results. Various merging paradigms have been proposed [22, 43, 72, 75]. However, their direct adaptation to motion planning is either infeasible or insufficient. We will further elaborate on this in the following section, explaining how we address these challenges in motion planning.

3.3. Interaction-Conserving Pre-Merging

There are two major challenges in obtaining usable model checkpoints before the merging process in motion planning.

First, the motion planning task utilizes multiple metrics, making it challenging to extract the optimal checkpoint. In classification tasks where merging has traditionally been applied, checkpoint selection is based on a single metric such as classification accuracy, simplifying the process. However, the evaluation of planning tasks is based on multiple criteria, including effectiveness (ADE), goal success (FDE, Miss Rate), and safety (Collision Rate; CR) where details on the metrics are discussed in Sec. 4.1. Since motion planning requires evaluation across multiple criteria, most existing planning models [20, 21, 32] employ multiple loss functions, such as trajectory deviation from the ground truth and collision avoidance penalties. For simplicity, we denote the sum of all relevant losses as \mathcal{L}_{total} . As \mathcal{L}_{total} consists of multiple loss terms, individual metrics exhibit trade-offs and fluctuations during training. This complicates the selection of an optimal checkpoint for merging.

Second, prior merging methods rely on pretrained models with large-scale datasets, such as ViT [15] pretrained on ImageNet [12]. However, this setup differs significantly from motion planning, where models are typically smaller, and no large-scale dataset equivalent to ImageNet exists. As a result, previous works in the vision domain were sufficient to extract only a single checkpoint per domain for merging networks. In contrast, motion planning suffers from data scarcity, necessitating additional methods to expand the checkpoint pool for effective merging.

To address these challenges, we introduce an effective interaction-conserving pre-merging approach, as outlined in Algorithm 1. Starting from the parameters Θ^0 initialized with Θ_0 , we iteratively update the model parameters using $\Theta^k = \Theta^{k-1} - \eta \nabla_{\Theta} \mathcal{L}_{total}$ over a total of K iterations with a learning rate η . The superscript of Θ denotes the iteration, while the subscript refers to the checkpoint

Algorithm 1: Interaction-Merged Motion Planning

Data: Source domains $\mathcal{S} = \{D_s^1, \dots, D_s^n\}$,
total iterations K , learning rate η ,
checkpoint interval C , checkpoint pool \mathcal{P} ,
initial model parameters Θ^0

```
1 Initialization:  $\mathcal{P} \leftarrow \{\}$ 
  // Interaction-Conserving Pre-Merging
2 for  $D_s^i$  in  $\mathcal{S}$  do
3   for  $k = 1$  to  $K$  do
4     Update parameters:
       $\Theta^k \leftarrow \Theta^{k-1} - \eta \nabla_{\Theta} \mathcal{L}_{total}$ 
5     if  $k \equiv 0 \pmod{C}$  then
6        $\mathcal{P} \leftarrow \mathcal{P} \cup \{\Theta^k\}$ 
7     for  $m \in \{ADE, FDE, CR, MR\}$  do
8       if  $m^k$  is better than  $m^{best}$  then
9          $\Theta^{best,m} \leftarrow \Theta^k$ 
10         $m^{best} \leftarrow m^k$ 
11 for  $m \in \{ADE, FDE, CR, MR\}$  do
12    $\mathcal{P} \leftarrow \mathcal{P} \cup \{\Theta^{best,m}\}$ 
  // Interaction Transfer with Merging
13 Let  $\mathcal{P} = \{\Theta_1, \dots, \Theta_{|\mathcal{P}|}\}$ 
14 for  $\theta \in \{\theta_{ego}, \theta_{surr}, \theta_{inter}, \theta_{else}\}$  do
15   for  $i = 1$  to  $|\mathcal{P}|$  do
16     Compute task vector:  $\tau_i = \theta_i - \theta$ 
17   Merge parameters:  $\theta^* = \theta_0 + \sum_{i=1}^{|\mathcal{P}|} w_{i,\theta} \cdot \tau_i$ 
18 Update all  $\{w_{i,\theta}\}_{i=1}^{|\mathcal{P}|}$  using the train set and  $\mathcal{L}_{total}$ .
```

pool number. To overcome the first challenge, which is that various evaluation metrics fluctuate during optimization, we extract parameters corresponding to the best performance for each metric. Since different metrics emphasize distinct aspects of agent behavior and interactions, the optimal parameters for different metrics reflect different characteristics. Thus, we collect these best-performing parameters into checkpoint pool \mathcal{P} .

As for the second challenge, we leverage the fact that the feature space in intermediate checkpoints during optimization tends to be more generalized across domain shifts, as they are not overly fitted to the source domains. Based on this, we include these intermediate checkpoints at intervals of C from the optimization trajectory in our checkpoint pool. This process is applied across all source domains.

3.4. Interaction Transfer with Merging

Our goal is to effectively transfer appropriate planning-related information from the source to the target domains. To this end, our idea starts with the observation that each source dataset contains unique ego and surrounding agents

past trajectories and interaction information. Therefore, we separate and respectively merge the modules for each distinct features, enabling a better adaptation on the unique characteristics of the target domain. This approach differs from traditional merging techniques, such as model-level merging and parameter-wise merging. As shown in Tab. 2, extracting key features individually within the source domain is more efficient than traditional merging methods. This demonstrates that, for domain-robust motion planning, it is crucial to properly extract key modules during the merging process. To ensure the general applicability of the IMMP methodology to various planning models, we categorize the modules that are essential in planning networks.

Conventional planning models employ hierarchical structures to encode agent behavior and their interactions. For instance, the trajectory information of the ego agent and surrounding agents is embedded using the modules ϕ_{ego} and ϕ_{surr} , typically implemented with LSTMs.

$$h_{ego} = \phi_{ego}(\mathcal{X}_{ego}^{-T_{obs}:0}; \theta_{ego}) \quad (3)$$

$$h_{surr} = \phi_{surr}(\mathcal{X}_{surr}^{-T_{obs}:0}, \hat{\mathcal{Y}}_{surr}^{1:T_{fut}}; \theta_{surr}) \quad (4)$$

where θ_{ego} and θ_{surr} are the parameters of each module. Interaction information is embedded based on h_{ego} and h_{surr} using the ψ_{inter} module, which is typically implemented with a Transformer.

$$h_{inter} = \psi_{inter}(h_{surr}, h_{ego}; \theta_{inter}) \quad (5)$$

where θ_{inter} is the parameters of ψ_{inter} . Finally, the planning $\hat{\mathcal{Y}}_{ego}^{1:T_{fut}}$ is generated based on the decoder using h_{inter} . The detailed model structure of the planning models used in the experiments is in suppl B.

Based on the key modules of the planning model that learn features with significant differences between the source domains, we group the parameters θ_{ego} , θ_{surr} , and θ_{inter} while classifying the remaining decoder and other layers as θ_{else} . Thus, the complete set of model parameters is defined as $\Theta = \{\theta_{ego}, \theta_{surr}, \theta_{inter}, \theta_{else}\}$. We perform merging separately for each parameter group, assigning distinct merging weights $\{w_{i,\theta}\}_{i=1}^{|\mathcal{P}|}$ to each θ in $\{\theta_{ego}, \theta_{surr}, \theta_{inter}, \theta_{else}\}$. This enables interaction-level information merging for effective adaptation in planning models while preserving the feature hierarchy. The overall procedure is outlined in the interaction transfer with merging step in Algorithm 1.

Relying solely on model parameters during adaptation provides two key advantages. First, it eliminates the need to access source domains after extracting the necessary parameter checkpoints, reducing the adaptation cost for the target domain with a pre-constructed checkpoint pool. Second, it effectively mitigates domain imbalance and catastrophic forgetting, which arise from imbalanced dataset compositions that affect previous approaches.

4. Experiments

4.1. Experimental Setup

Datasets & Backbone. In our experiments, we utilize three categories of datasets: ETH-UCY [36, 50] (Human-Human Interaction dataset), CrowdNav [6] (RL algorithm-based Robot dataset), THOR [52] and SIT [4] (Human-Robot Interaction dataset). For ETH-UCY, we leverage the unique characteristics of each sub-dataset while maintaining the overall dataset size by alternately selecting 4 out of 5 scenes to compose source datasets. The excluded scene is denoted as “w/o” in our notation. To extract the key characteristics of each dataset appropriately, we apply a preprocessing step to standardize all data into a uniform format. Specifically, we extract ego-centric x, y coordinates from all datasets and sample the data at 2.5 FPS. Additionally, the observed past trajectories consist of 8 timesteps, while the predicted trajectories of surrounding agents and the planned trajectory of the ego agent are set to 12 timesteps. In our experiments, we vary the target domain across different datasets, using SIT and THOR as target domains. Considering practical applicability, THOR and SIT are selected as target datasets, as they are relatively small so require more information transfer from other datasets. The datasets not included in the target domain are treated as the source domain. We evaluate the performance of our methodology against three planning baseline models: **GameTheoretic** [32], **DIPP** [20], and **DTPP** [21]. Detailed descriptions of each data collection process and the backbone planning models are provided in suppl A and B.

Metrics. The primary evaluation criteria for autonomous robot driving include Effectiveness, Safety, and Goal Success. Average Displacement Error (ADE) is used as the metric for Effectiveness, Collision Rate (CR) for Safety, and Miss Rate (MR) along with Final Displacement Error (FDE) for Goal Success. Additionally, the final inference cost, such as FLOPs or the number of model parameters, is measured as Cost. Since all methods use the same network architecture, Cost is reported as a relative multiplier with respect to the baseline model. More details on the evaluation metrics are provided in suppl C.

Baselines. Our approach is not model-specific and can be applied orthogonally with other methods, so we use non-model-specific methods as baselines and compare them with our IMMP. (1) **Target Only** trains a model exclusively on the target dataset. (2) **Domain Generalization** [17] pre-trains the model on all datasets except the target dataset. (3) **Domain Adaptation** [16] pre-trains the model on all datasets except the target dataset and then finetunes it on the target dataset. (4) **Ensemble** [2, 37] uses multiple models, each on a different source dataset, and combines their trajectories for the final planning. Specifically, we incorporate two ensemble strategies for comparison: winner-takes-all

(WTA) and averaging (AVG). Additionally, we incorporate model merging approaches such as (5) **Averaging** [10, 68], which computes the parameter-wise mean of all individual models, defined as $\theta_{merge} = \sum_{i=1}^n \theta_i / n$. (6) **Task Arithmetic** [22] sums all task vectors and scales them to produce a merged model, formulated as $\theta_{merge} = \theta_0 + \lambda \cdot \sum_{i=1}^n \tau_i$. (7) **Ties Merging** [72] performs the merge in three steps: trimming task vectors with minimal change during training, resolving sign conflicts, and merging only parameters that align with the final signed agreed on.

4.2. Experimental Results

Merging Interaction for Better Initialization. IMMP not only improves generalization for target domains but also provides better initialization for fine-tuning to target domains. Table 1 presents IMMP’s performance after directly merging parameter checkpoints, while IMMP+Finetune shows the results of fine-tuning the merged model parameters used as initialization. Notably, IMMP+Finetune outperforms both Domain-Specific and Domain Adaptation scenarios. This result supports that IMMP offers a strong initialization point for fine-tuning.

Domain Gap and Generalization. Table 1 presents a comparison of IMMP with baseline methods. Domain generalization for motion planning, which incorporates multiple source datasets [16], shows significantly lower performance. Interestingly, training solely on the target domain outperforms domain generalization [17], highlighting the substantial domain disparity between source and target datasets. Additionally, domain adaptation does not always achieve better performance than the target-only approach, further supporting this observation. Ensemble-WTA significantly outperforms Ensemble-AVG, suggesting that only certain source domains effectively contribute to target domain performance. These results indicate potential limitations in previous methods, such as domain imbalance or catastrophic forgetting. In contrast, IMMP leverages parameters trained across different domains, reducing susceptibility to these issues.

Comparison to Merging Approaches. Table 1 shows that the naive adoption of existing model merging techniques leads to poor performance in motion planning. Previous methods fail to consider the feature hierarchy specific to motion planning, where features can be encoded at the agent behavior and interaction levels. As a result, motion planning models exhibit significant disparities after training, making them difficult to merge using conventional methods. Specifically, Averaging and Task Arithmetic can disrupt the hierarchy of features, while Ties Merging tends to discard too much critical information during the trimming process when resolving conflicts. In contrast, IMMP efficiently transfers information from source datasets to the target domain, achieving superior performance.

Table 1. Qualitative comparison of the proposed IMMP with baseline methods. Each baseline is implemented across various planning models. The methods are categorized into conventional approaches for generalization in motion planning, adaptations of other model merging methods in our setting, and the proposed IMMP.

Target Dataset		SIT [4]				THOR [52]				Cost
Model	Method	Effectiveness ADE ↓	Safety Col. Rate ↓	Goal Success FDE ↓ Miss Rate ↓		Effectiveness ADE ↓	Safety Col. Rate ↓	Goal Success FDE ↓ Miss Rate ↓		
GameTheoretic [32]	Domain Generalization [17]	0.8338	9.87E-04	1.8594	0.9355	0.3804	1.21E-03	0.8705	0.6957	×1
	Domain Adaptation [16]	0.4388	1.26E-03	1.0611	0.7201	0.1133	2.90E-04	0.2516	0.1268	×1
	Target Only	0.4343	3.41E-04	0.9014	0.6272	0.1003	2.64E-04	0.2153	0.0929	×1
	Ensemble-WTA [2]	0.3695	5.75E-05	0.8283	0.6185	0.2112	4.03E-04	0.4082	0.2996	×7
	Ensemble-AVG [2]	0.5415	1.89E-04	1.1617	0.8159	0.3142	4.86E-04	0.6181	0.6027	×7
	Averaging [10, 68]	0.6726	2.34E-04	1.4340	0.9611	0.2742	5.17E-04	0.5624	0.5711	×1
	Task Arithmetic [22]	0.4132	1.37E-04	0.8936	0.7364	0.2679	3.87E-04	0.5651	0.4812	×1
	Ties Merging [72]	1.1876	5.53E-04	2.2440	0.9872	0.5253	2.49E-03	0.8212	0.6201	×1
	IMMP	0.3380	5.12E-05	0.7626	0.6446	0.1165	3.48E-04	0.2562	0.1330	×1
IMMP + Finetune	0.3157	4.28E-05	0.7300	0.5934	0.0975	2.56E-04	0.2108	0.0912	×1	
DTPP [21]	Domain Generalization [17]	1.0116	1.18E-03	2.1255	0.9350	0.3446	1.42E-03	0.6873	0.5948	×1
	Domain Adaptation [16]	0.4640	2.45E-04	1.0116	0.7573	0.1460	2.50E-04	0.2291	0.0772	×1
	Target Only	0.4832	5.73E-04	0.9839	0.6852	0.1528	2.57E-04	0.2394	0.0910	×1
	Ensemble-WTA [2]	0.6125	2.01E-04	1.2067	0.8217	0.2241	5.15E-04	0.4591	0.4127	×7
	Ensemble-AVG [2]	0.7460	3.36E-04	1.5439	0.9175	0.2145	3.41E-04	0.4095	0.3302	×7
	Averaging [10, 68]	0.6303	2.79E-04	1.2699	0.8641	0.3075	6.75E-04	0.5807	0.4901	×1
	Task Arithmetic [22]	0.6062	3.73E-04	1.2138	0.8542	0.2885	4.18E-04	0.5057	0.4693	×1
	Ties Merging [72]	1.0741	2.25E-03	2.0873	0.9384	0.5141	2.15E-03	0.6764	0.5514	×1
	IMMP	0.4388	2.58E-04	0.9067	0.6220	0.1958	2.46E-04	0.3210	0.1558	×1
IMMP + Finetune	0.3793	4.22E-04	0.8336	0.6336	0.1470	2.17E-04	0.2421	0.0905	×1	
DIPP [20]	Domain Generalization [17]	1.3268	1.09E-03	2.7111	0.9756	0.3398	1.27E-03	0.6983	0.6163	×1
	Domain Adaptation [16]	0.4697	3.93E-04	1.0331	0.7886	0.2660	9.68E-04	0.5272	0.4773	×1
	Target Only	0.5671	6.68E-04	0.9801	0.7253	0.1771	2.17E-04	0.2713	0.1156	×1
	Ensemble-WTA [2]	0.7784	4.36E-04	1.5916	0.8275	0.2604	7.00E-04	0.4874	0.4435	×7
	Ensemble-AVG [2]	0.9737	5.79E-04	1.9442	0.9750	0.2562	5.46E-04	0.5134	0.4754	×7
	Averaging [10, 68]	1.4383	1.74E-03	2.7763	0.9820	0.4329	1.47E-03	0.7534	0.6710	×1
	Task Arithmetic [22]	1.2614	1.66E-03	2.4944	0.9564	0.4196	1.39E-03	0.7120	0.6915	×1
	Ties Merging [72]	1.6028	1.85E-03	3.0906	1.0000	0.5141	2.15E-03	0.6764	0.5515	×1
	IMMP	0.5112	1.60E-04	0.9358	0.7944	0.2335	5.87E-04	0.3219	0.1842	×1
IMMP + Finetune	0.4096	5.27E-04	0.8915	0.6789	0.1204	1.94E-04	0.2236	0.0769	×1	

Table 2. Impact of merging granularity in IMMP for SIT Datasets with GameTheoretic. The granularity levels include Model-level, Parameter-level, and Interaction-level.

Granularity	Effectiveness	Safety	Goal Success	
	ADE ↓	Col. Rate ↓	FDE ↓	Miss Rate ↓
Model	0.3687	7.15E-05	0.8365	0.8002
Parameter	0.3798	9.16E-05	0.7754	0.7433
Interaction	0.3380	5.12E-05	0.7626	0.6446

Qualitative Results. In Fig. 3, we qualitatively analyze how task vectors extracted from each dataset are utilized. Figure. 3 (a) to (g) illustrate the inference results in the target domain SIT using GameTheoretic models trained on individual source domains, while Fig. 3 (h) presents the inference result of the planner after interaction transfer merging in the target domain. The weights above the inference results represent the actual merging weights applied by the proposed IMMP. Notably, there is a meaningful relationship between the merging weights and the performance of individual domains. As shown in Fig. 3 (a) and (g), source datasets with poor inference performance in the target do-

main are assigned lower weights, whereas well-performing domain, such as Fig. 3 (f), is assigned higher weight and utilized more extensively. This demonstrates that our approach effectively prioritizes source domain datasets that are most beneficial for the target domain.

4.3. Ablation Study

Merging Granularity. Table 2 evaluates the impact of merging granularity in SIT datasets using GameTheoretic. We consider three levels: ‘Model-level,’ where identical merging weights apply to all parameters; ‘Parameter-level,’ where merging is parameter-wise; and ‘Interaction-level,’ which merges components related to past ego-trajectory, surrounding agents, and their interactions. Results show that Interaction-level merging achieves the best performance across all metrics. Model-level merging fails to adapt interactions and agent behaviors from source to target domains, as full-model transfer lacks flexibility for domain-specific interactions. Parameter-level merging allows finer adaptation but disrupts the feature hierarchy of trajectory

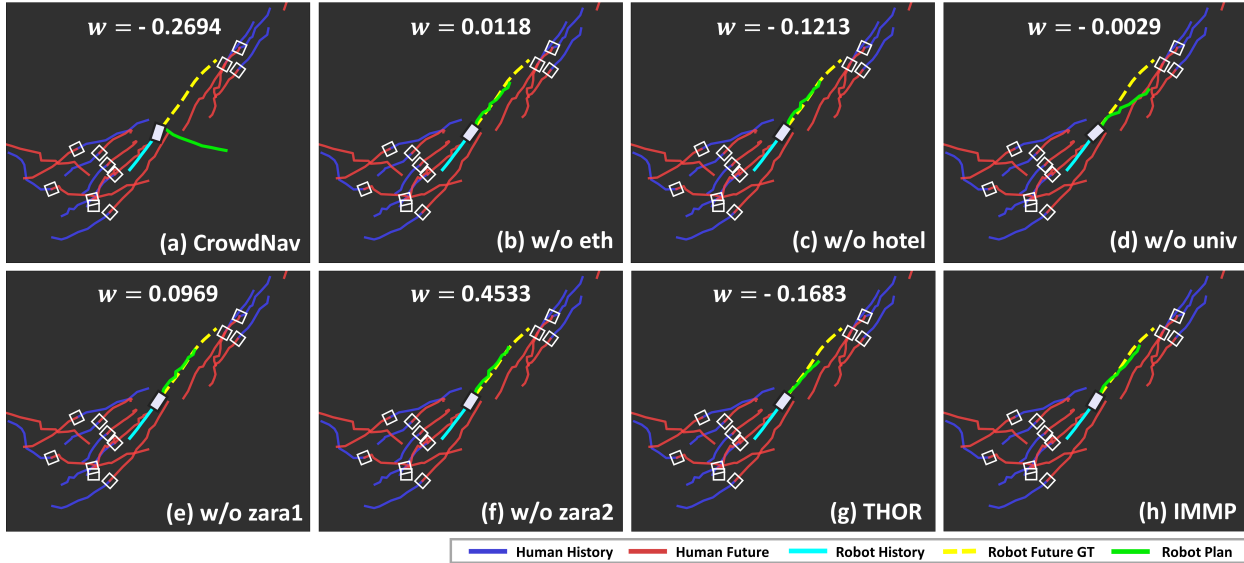


Figure 3. Qualitative results on the SIT dataset with GameTheoretic planning model. (a) to (g) represent the inference results of SIT using models trained on individual dataset from the source domains. The values above, denoted by w , indicate the average of contributions of task vectors per epoch and per module for each source domain in the IMMP planner.

Table 3. Ablation study on checkpoint types and pool size for merging. IMMP planning performance with best checkpoints (across all metrics) and intermediate-epoch for planning (P) / forecaster (F) models; fine-tuning results included.

Selection Method		Finetune	Effectiveness	Safety	Goal Success	
All Metric	Epoch		ADE ↓	Col. Rate ↓	FDE ↓	Miss Rate ↓
✓			0.3646	8.12E-05	0.8063	0.6516
✓	✓ (P)		0.3543	6.44E-05	0.7730	0.6196
✓	✓ (F,P)		0.3380	5.12E-05	0.7626	0.6446
		✓	0.3203	5.40E-05	0.7365	0.5743
✓		✓	0.3186	5.09E-05	0.7358	0.5981
✓	✓ (P)	✓	0.3176	4.31E-05	0.7399	0.5906
✓	✓ (F,P)	✓	0.3157	4.28E-05	0.7300	0.5934

encoding and interaction modeling, introducing instability. While it offers slight improvements over Model-level merging, it does not outperform Interaction-level merging in most cases. These findings underscore the need for proper merging granularity to retain interaction structures and improve planning model adaptation in the target domain.

Impact of Checkpoint Types and Pool Size. Table 3 evaluates how the composition of the checkpoint pool \mathcal{P} affects planning performance using GameTheoretic. ‘All Metric’ selects checkpoints based on the best performance across all evaluation metrics, while its absence defaults to ADE-based selection. ‘Epoch’ includes intermediate checkpoints along the optimization trajectory. Given the separate forecaster and planner networks in GameTheoretic, we independently select their checkpoints. For instance, (P) includes only intermediate checkpoints of the planning network, while (F,P) incorporates both. Fine-tuning on the target domain and increasing the checkpoint pool size improve performance after both merging and fine-tuning. Notably, selecting check-

points using multiple metrics significantly enhances their corresponding evaluation scores, as seen in the superior performance of ‘All Metric’ over ADE-only selection. Intermediate checkpoints further benefit adaptation, preventing overfitting to the source domain. However, these results do not necessarily imply that the pre-merging phase is essential, as IMMP still achieves promising performance compared to previous approaches when using only checkpoints selected based on ‘ADE’.

5. Conclusion

In this paper, we present the Interaction-Merged Motion Planning (IMMP), which effectively leverages diverse motion planning datasets. By utilizing parameter checkpoints trained on different domains, IMMP addresses domain imbalance and catastrophic forgetting that hinder conventional adaptation approaches. The proposed two-step process—interaction-conserving pre-merging and interaction-level merging—preserves agent behavior and interactions, resulting in a more adaptable planning model. Experimental results demonstrate that IMMP outperforms traditional approaches, achieving superior performance.

Acknowledgment

This work was supported by the National Research Foundation of Korea(NRF) grant funded by the Korea government(MSIT) (NRF2022R1A2B5B03002636), and by the Institute of Information & communications Technology Planning & Evaluation (IITP) grant funded by the Korea government(MSIT) (No. RS-2024-00457882, AI Research Hub Project).

References

- [1] Hub Ali, Dawei Gong, Meng Wang, and Xiaolin Dai. Path planning of mobile robot with improved ant colony algorithm and mdp to produce smooth trajectory in grid-based environment. *Frontiers in neurorobotics*, 14:44, 2020. 2
- [2] Devansh Arpit, Huan Wang, Yingbo Zhou, and Caiming Xiong. Ensemble of averages: Improving model selection and boosting performance in domain generalization. *Advances in Neural Information Processing Systems*, 35:8265–8277, 2022. 3, 6, 7
- [3] Oktay Arslan and Panagiotis Tsiotras. Use of relaxation methods in sampling-based algorithms for optimal motion planning. In *2013 IEEE International Conference on Robotics and Automation*, pages 2421–2428. IEEE, 2013. 2
- [4] Jong Wook Bae, Jungho Kim, Junyong Yun, Changwon Kang, Jeongseon Choi, Chanhyeok Kim, Junho Lee, Jungwook Choi, and Jun Won Choi. Sit dataset: socially interactive pedestrian trajectory dataset for social navigation robots. *Advances in Neural Information Processing Systems*, 36:24552–24563, 2023. 1, 6, 7
- [5] Wei Cheah, Hassan Hakim Khalili, Simon Watson, Peter Green, and Barry Lennox. Grid-based motion planning using advanced motions for hexapod robots. In *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 3573–3578. IEEE, 2018. 2
- [6] Changan Chen, Yuejiang Liu, Sven Kreiss, and Alexandre Alahi. Crowd-robot interaction: Crowd-aware robot navigation with attention-based deep reinforcement learning. In *2019 international conference on robotics and automation (ICRA)*, pages 6015–6022. IEEE, 2019. 1, 6
- [7] Li Chen, Penghao Wu, Kashyap Chitta, Bernhard Jaeger, Andreas Geiger, and Hongyang Li. End-to-end autonomous driving: Challenges and frontiers. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2024. 1
- [8] Zhili Chen, Maosheng Ye, Shuangjie Xu, Tongyi Cao, and Qifeng Chen. Ppad: Iterative interactions of prediction and planning for end-to-end autonomous driving. In *European Conference on Computer Vision*, pages 239–256. Springer, 2024. 2
- [9] Jie Cheng, Xiaodong Mei, and Ming Liu. Forecast-mae: Self-supervised pre-training for motion forecasting with masked autoencoders. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 8679–8689, 2023. 2
- [10] Leshem Choshen, Elad Venezian, Noam Slonim, and Yoav Katz. Fusing finetuned models for better pretraining. *arXiv preprint arXiv:2204.03044*, 2022. 6, 7
- [11] Caglar Demir, Arnab Sharma, and Axel-Cyrille Ngonga Ngomo. Adaptive stochastic weight averaging. *arXiv preprint arXiv:2406.19092*, 2024. 2, 3
- [12] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009. 4
- [13] Yonghao Dong, Le Wang, Sanping Zhou, and Gang Hua. Sparse instance conditioned multimodal trajectory prediction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9763–9772, 2023. 2
- [14] Yonghao Dong, Le Wang, Sanping Zhou, Gang Hua, and Changyin Sun. Recurrent aligned network for generalized pedestrian trajectory prediction. *arXiv preprint arXiv:2403.05810*, 2024. 3
- [15] 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. *arXiv preprint arXiv:2010.11929*, 2020. 4
- [16] Lan Feng, Mohammadhossein Bahari, Kaouther Mes-saoud Ben Amor, Éloi Zablocki, Matthieu Cord, and Alexandre Alahi. Unitrax: A unified framework for scalable vehicle trajectory prediction. In *European Conference on Computer Vision*, pages 106–123. Springer, 2024. 2, 3, 6, 7
- [17] Thomas Gilles, Stefano Sabatini, Dzmitry Tsishkou, Bogdan Stanculescu, and Fabien Moutarde. Uncertainty estimation for cross-dataset performance in trajectory prediction. *arXiv preprint arXiv:2205.07310*, 2022. 3, 6, 7
- [18] Yihan Hu, Jiazhi Yang, Li Chen, Keyu Li, Chonghao Sima, Xizhou Zhu, Siqi Chai, Senyao Du, Tianwei Lin, Wenhai Wang, et al. Planning-oriented autonomous driving. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 17853–17862, 2023. 2
- [19] Zhiyu Huang, Haochen Liu, and Chen Lv. Gameformer: Game-theoretic modeling and learning of transformer-based interactive prediction and planning for autonomous driving. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3903–3913, 2023. 1, 2
- [20] Zhiyu Huang, Haochen Liu, Jingda Wu, and Chen Lv. Differentiable integrated motion prediction and planning with learnable cost function for autonomous driving. *IEEE transactions on neural networks and learning systems*, 2023. 4, 6, 7
- [21] Zhiyu Huang, Peter Karkus, Boris Ivanovic, Yuxiao Chen, Marco Pavone, and Chen Lv. Dtp: Differentiable joint conditional prediction and cost evaluation for tree policy planning in autonomous driving. In *2024 IEEE International Conference on Robotics and Automation (ICRA)*, pages 6806–6812. IEEE, 2024. 1, 2, 4, 6, 7
- [22] Gabriel Ilharco, Marco Tulio Ribeiro, Mitchell Wortsman, Suchin Gururangan, Ludwig Schmidt, Hannaneh Hajishirzi, and Ali Farhadi. Editing models with task arithmetic. *arXiv preprint arXiv:2212.04089*, 2022. 3, 4, 6, 7
- [23] Jaewoo Jeong, Daehee Park, and Kuk-Jin Yoon. Multi-agent long-term 3d human pose forecasting via interaction-aware trajectory conditioning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1617–1628, 2024. 1
- [24] Jaewoo Jeong, Seohee Lee, Daehee Park, Giwon Lee, and Kuk-Jin Yoon. Multi-modal knowledge distillation-based human trajectory forecasting. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 24222–24233, 2025. 1
- [25] Woosong Jeong and Kuk-Jin Yoon. Quantifying task priority for multi-task optimization. In *Proceedings of the*

- IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 363–372, 2024. 3
- [26] Woosong Jeong and Kuk-Jin Yoon. Selective task group updates for multi-task optimization. *arXiv preprint arXiv:2502.11986*, 2025. 3
- [27] Xiaosong Jia, Penghao Wu, Li Chen, Jiangwei Xie, Conghui He, Junchi Yan, and Hongyang Li. Think twice before driving: Towards scalable decoders for end-to-end autonomous driving. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 21983–21994, 2023. 2
- [28] Bo Jiang, Shaoyu Chen, Qing Xu, Bencheng Liao, Jiajie Chen, Helong Zhou, Qian Zhang, Wenyu Liu, Chang Huang, and Xinggong Wang. Vad: Vectorized scene representation for efficient autonomous driving. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 8340–8350, 2023. 2
- [29] Ruochen Jin, Bojian Hou, Jiancong Xiao, Weijie Su, and Li Shen. Fine-tuning linear layers only is a simple yet effective way for task arithmetic. *arXiv preprint arXiv:2407.07089*, 2024. 3
- [30] Sertac Karaman and Emilio Frazzoli. Sampling-based algorithms for optimal motion planning. *The international journal of robotics research*, 30(7):846–894, 2011. 2
- [31] Lydia E Kavraki, Petr Svestka, J-C Latombe, and Mark H Overmars. Probabilistic roadmaps for path planning in high-dimensional configuration spaces. *IEEE transactions on Robotics and Automation*, 12(4):566–580, 1996. 2
- [32] Kushal Kedia, Prithwish Dan, and Sanjiban Choudhury. A game-theoretic framework for joint forecasting and planning. In *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 6773–6778. IEEE, 2023. 1, 2, 4, 6, 7
- [33] Alex Kendall, Yarin Gal, and Roberto Cipolla. Multi-task learning using uncertainty to weigh losses for scene geometry and semantics. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7482–7491, 2018. 3
- [34] James Kirkpatrick, Razvan Pascanu, Neil Rabinowitz, Joel Veness, Guillaume Desjardins, Andrei A Rusu, Kieran Milan, John Quan, Tiago Ramalho, Agnieszka Grabska-Barwinska, et al. Overcoming catastrophic forgetting in neural networks. *Proceedings of the national academy of sciences*, 114(13):3521–3526, 2017. 2
- [35] James J Kuffner and Steven M LaValle. Rrt-connect: An efficient approach to single-query path planning. In *Proceedings 2000 ICRA. Millennium Conference. IEEE International Conference on Robotics and Automation. Symposia Proceedings (Cat. No. 00CH37065)*, pages 995–1001. IEEE, 2000. 2
- [36] Alon Lerner, Yiorgos Chrysanthou, and Dani Lischinski. Crowds by example. In *Computer graphics forum*, pages 655–664. Wiley Online Library, 2007. 1, 6
- [37] Zirui Li, Yunlong Lin, Cheng Gong, Xinwei Wang, Qi Liu, Jianwei Gong, and Chao Lu. An ensemble learning framework for vehicle trajectory prediction in interactive scenarios. In *2022 IEEE Intelligent Vehicles Symposium (IV)*, pages 51–57. IEEE, 2022. 3, 6
- [38] Bo Liu, Xingchao Liu, Xiaojie Jin, Peter Stone, and Qiang Liu. Conflict-averse gradient descent for multi-task learning. *Advances in Neural Information Processing Systems*, 34:18878–18890, 2021. 3
- [39] Bo Liu, Yihao Feng, Peter Stone, and Qiang Liu. Famo: Fast adaptive multitask optimization. *Advances in Neural Information Processing Systems*, 36:57226–57243, 2023. 3
- [40] Bo Liu, Yihao Feng, Peter Stone, and Qiang Liu. Famo: Fast adaptive multitask optimization. *Advances in Neural Information Processing Systems*, 36, 2024. 3
- [41] Tian Yu Liu, Aditya Golatkar, and Stefano Soatto. Tangent transformers for composition, privacy and removal. *arXiv preprint arXiv:2307.08122*, 2023. 3
- [42] Roberto Martin-Martin, Mihir Patel, Hamid Rezatofighi, Abhijeet Sheno, JunYoung Gwak, Eric Frankel, Amir Sadeghian, and Silvio Savarese. Jrdb: A dataset and benchmark of egocentric robot visual perception of humans in built environments. *IEEE transactions on pattern analysis and machine intelligence*, 45(6):6748–6765, 2021. 1
- [43] Michael S Matena and Colin A Raffel. Merging models with fisher-weighted averaging. *Advances in Neural Information Processing Systems*, 35:17703–17716, 2022. 4
- [44] Aviv Navon, Aviv Shamsian, Idan Achituve, Haggai Maron, Kenji Kawaguchi, Gal Chechik, and Ethan Fetaya. Multi-task learning as a bargaining game. *arXiv preprint arXiv:2202.01017*, 2022. 3
- [45] Guillermo Ortiz-Jimenez, Alessandro Favero, and Pascal Frossard. Task arithmetic in the tangent space: Improved editing of pre-trained models. *Advances in Neural Information Processing Systems*, 36, 2024. 2, 3
- [46] Chenbin Pan, Burhaneddin Yaman, Tommaso Nesti, Abhirup Mallik, Alessandro G Allievi, Senem Velipasalar, and Liu Ren. Vlp: Vision language planning for autonomous driving. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14760–14769, 2024. 2
- [47] Daehee Park, Hobin Ryu, Yunseo Yang, Jegyeong Cho, Jiwon Kim, and Kuk-Jin Yoon. Leveraging future relationship reasoning for vehicle trajectory prediction. In *The Eleventh International Conference on Learning Representations*, 2023. 2
- [48] Daehee Park, Jaewoo Jeong, and Kuk-Jin Yoon. Improving transferability for cross-domain trajectory prediction via neural stochastic differential equation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 10145–10154, 2024. 3
- [49] Daehee Park, Jaeseok Jeong, Sung-Hoon Yoon, Jaewoo Jeong, and Kuk-Jin Yoon. T4p: Test-time training of trajectory prediction via masked autoencoder and actor-specific token memory. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15065–15076, 2024. 3
- [50] Stefano Pellegrini, Andreas Ess, Konrad Schindler, and Luc Van Gool. You’ll never walk alone: Modeling social behavior for multi-target tracking. In *2009 IEEE 12th international conference on computer vision*, pages 261–268. IEEE, 2009. 1, 6

- [51] Tangwen Qian, Yile Chen, Gao Cong, Yongjun Xu, and Fei Wang. Adaptraj: A multi-source domain generalization framework for multi-agent trajectory prediction. In *2024 IEEE 40th International Conference on Data Engineering (ICDE)*, pages 5048–5060. IEEE, 2024. 3
- [52] Andrey Rudenko, Tomasz P Kucner, Chittaranjan S Swaminathan, Ravi T Chadalavada, Kai O Arras, and Achim J Lilienthal. Thör: Human-robot navigation data collection and accurate motion trajectories dataset. *IEEE Robotics and Automation Letters*, 5(2):676–682, 2020. 1, 6, 7
- [53] Abbas Sadat, Sergio Casas, Mengye Ren, Xinyu Wu, Pranaab Dhawan, and Raquel Urtasun. Perceive, predict, and plan: Safe motion planning through interpretable semantic representations. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXIII 16*, pages 414–430. Springer, 2020. 1
- [54] Jacques Saraydaryan, Fabrice Jamel, and Olivier Simonin. Navigation in human flows: planning with adaptive motion grid. In *IROS Workshop CrowdNav*, 2018. 2
- [55] Jonathan Schwarz, Wojciech Czarnecki, Jelena Luketina, Agnieszka Grabska-Barwinska, Yee Whye Teh, Razvan Pascanu, and Raia Hadsell. Progress & compress: A scalable framework for continual learning. In *International conference on machine learning*, pages 4528–4537. PMLR, 2018. 2
- [56] Dmitry Senushkin, Nikolay Patakin, Arseny Kuznetsov, and Anton Konushin. Independent component alignment for multi-task learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 20083–20093, 2023. 3
- [57] Hanul Shin, Jung Kwon Lee, Jaehong Kim, and Jiwon Kim. Continual learning with deep generative replay. *Advances in neural information processing systems*, 30, 2017. 2
- [58] Konstantin Shmelkov, Cordelia Schmid, and Karteek Alahari. Incremental learning of object detectors without catastrophic forgetting. In *Proceedings of the IEEE international conference on computer vision*, pages 3400–3409, 2017. 2
- [59] Anke Tang, Li Shen, Yong Luo, Yibing Zhan, Han Hu, Bo Du, Yixin Chen, and Dacheng Tao. Parameter efficient multi-task model fusion with partial linearization. *arXiv preprint arXiv:2310.04742*, 2023. 3
- [60] Georg Tanzmeister, Martin Friedl, Dirk Wollherr, and Martin Buss. Efficient evaluation of collisions and costs on grid maps for autonomous vehicle motion planning. *IEEE Transactions on Intelligent Transportation Systems*, 15(5):2249–2260, 2014. 2
- [61] Hanqing Wang, Wei Liang, Luc Van Gool, and Wenguan Wang. Dreamwalker: Mental planning for continuous vision-language navigation. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 10873–10883, 2023. 1
- [62] Jiankun Wang, Wenzheng Chi, Chenming Li, Chaoqun Wang, and Max Q-H Meng. Neural rrt*: Learning-based optimal path planning. *IEEE Transactions on Automation Science and Engineering*, 17(4):1748–1758, 2020. 2
- [63] Jifeng Wang, Kaouther Messaoud, Yuejiang Liu, Juergen Gall, and Alexandre Alahi. Forecast-peft: Parameter-efficient fine-tuning for pre-trained motion forecasting models. *arXiv preprint arXiv:2407.19564*, 2024. 3
- [64] Mingkun Wang, Xinge Zhu, Changqian Yu, Wei Li, Yuexin Ma, Ruochun Jin, Xiaoguang Ren, Dongchun Ren, Mingxu Wang, and Wenjing Yang. Ganet: Goal area network for motion forecasting. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1609–1615. IEEE, 2023. 2
- [65] Yuqi Wang, Jiawei He, Lue Fan, Hongxin Li, Yuntao Chen, and Zhaoxiang Zhang. Driving into the future: Multiview visual forecasting and planning with world model for autonomous driving. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14749–14759, 2024. 2
- [66] Zhibo Wang, Jiayu Guo, Haiqiang Zhang, Ru Wan, Junping Zhang, and Jian Pu. Bridging the gap: Improving domain generalization in trajectory prediction. *IEEE Transactions on Intelligent Vehicles*, 2023. 3
- [67] Xinshuo Weng, Boris Ivanovic, Yan Wang, Yue Wang, and Marco Pavone. Para-drive: Parallelized architecture for real-time autonomous driving. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15449–15458, 2024. 2
- [68] Mitchell Wortsman, Gabriel Ilharco, Samir Ya Gadre, Rebecca Roelofs, Raphael Gontijo-Lopes, Ari S Morcos, Hongseok Namkoong, Ali Farhadi, Yair Carmon, Simon Kornblith, et al. Model soups: averaging weights of multiple fine-tuned models improves accuracy without increasing inference time. In *International conference on machine learning*, pages 23965–23998. PMLR, 2022. 2, 3, 6, 7
- [69] Mitchell Wortsman, Gabriel Ilharco, Jong Wook Kim, Mike Li, Simon Kornblith, Rebecca Roelofs, Raphael Gontijo Lopes, Hannaneh Hajishirzi, Ali Farhadi, Hongseok Namkoong, et al. Robust fine-tuning of zero-shot models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 7959–7971, 2022. 2, 3
- [70] Yi Xu and Yun Fu. Adapting to length shift: Flexilength network for trajectory prediction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15226–15237, 2024. 2
- [71] Yi Xu, Lichen Wang, Yizhou Wang, and Yun Fu. Adaptive trajectory prediction via transferable gnn. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 6520–6531, 2022. 3
- [72] Prateek Yadav, Derek Tam, Leshem Choshen, Colin A Raffel, and Mohit Bansal. Ties-merging: Resolving interference when merging models. *Advances in Neural Information Processing Systems*, 36:7093–7115, 2023. 2, 3, 4, 6, 7
- [73] Zhi Yan, Tom Duckett, and Nicola Bellotto. Online learning for human classification in 3d lidar-based tracking. In *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 864–871. IEEE, 2017. 1
- [74] Brian Yang, Huangyuan Su, Nikolaos Gkanatsios, Tsung-Wei Ke, Ayush Jain, Jeff Schneider, and Katerina Fragkiadaki. Diffusion-es: Gradient-free planning with diffusion for autonomous and instruction-guided driving. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15342–15353, 2024. 1, 2

- [75] Enneng Yang, Zhenyi Wang, Li Shen, Shiwei Liu, Guibing Guo, Xingwei Wang, and Dacheng Tao. Adamerging: Adaptive model merging for multi-task learning. *arXiv preprint arXiv:2310.02575*, 2023. 3, 4
- [76] Enneng Yang, Li Shen, Guibing Guo, Xingwei Wang, Xiaochun Cao, Jie Zhang, and Dacheng Tao. Model merging in llms, mllms, and beyond: Methods, theories, applications and opportunities. *arXiv preprint arXiv:2408.07666*, 2024. 3
- [77] Zhengcai Yang, Yunzhong Hu, and Youbing Zhang. Path-planning strategy for lane changing based on adaptive-grid risk-fields of autonomous vehicles. *World Electric Vehicle Journal*, 13(10):175, 2022. 2
- [78] Luyao Ye, Zikang Zhou, and Jianping Wang. Improving the generalizability of trajectory prediction models with frenet-based domain normalization. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*, pages 11562–11568. IEEE, 2023. 3
- [79] Tianhe Yu, Saurabh Kumar, Abhishek Gupta, Sergey Levine, Karol Hausman, and Chelsea Finn. Gradient surgery for multi-task learning. *Advances in neural information processing systems*, 33:5824–5836, 2020. 3
- [80] Ekim Yurtsever, Jacob Lambert, Alexander Carballo, and Kazuya Takeda. A survey of autonomous driving: Common practices and emerging technologies. *IEEE access*, 8:58443–58469, 2020. 1
- [81] Wenyuan Zeng, Shenlong Wang, Renjie Liao, Yun Chen, Bin Yang, and Raquel Urtasun. Dsdnet: Deep structured self-driving network. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXI 16*, pages 156–172. Springer, 2020. 1
- [82] Wenzhao Zheng, Ruiqi Song, Xianda Guo, Chenming Zhang, and Long Chen. Genad: Generative end-to-end autonomous driving. In *European Conference on Computer Vision*, pages 87–104. Springer, 2024. 2
- [83] Zikang Zhou, Luyao Ye, Jianping Wang, Kui Wu, and Kejie Lu. Hivt: Hierarchical vector transformer for multi-agent motion prediction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8823–8833, 2022. 2
- [84] Zikang Zhou, Jianping Wang, Yung-Hui Li, and Yu-Kai Huang. Query-centric trajectory prediction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 17863–17873, 2023. 2
- [85] Zixiang Zhou, Yu Wan, and Baoyuan Wang. Avatargpt: All-in-one framework for motion understanding planning generation and beyond. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1357–1366, 2024. 2
- [86] Yuanshao Zhu, James Jianqiao Yu, Xiangyu Zhao, Xuetao Wei, and Yuxuan Liang. Unitraj: Universal human trajectory modeling from billion-scale worldwide traces. *arXiv preprint arXiv:2411.03859*, 2024. 2