FJMP: Factorized Joint Multi-Agent Motion Prediction over Learned Directed Acyclic Interaction Graphs

Luke Rowe, Martin Ethier, Eli-Henry Dykhne, Krzysztof Czarnecki
School of Computer Science, University of Waterloo
{l6rowe, methier, ehdykhne, krzysztof.czarnecki}@uwaterloo.ca

Abstract

Predicting the future motion of road agents is a critical task in an autonomous driving pipeline. In this work, we address the problem of generating a set of scene-level, or joint, future trajectory predictions in multi-agent driving scenarios. To this end, we propose FJMP, a Factorized Joint Motion Prediction framework for multi-agent interactive driving scenarios. FJMP models the future scene interaction dynamics as a sparse directed interaction graph, where edges denote explicit interactions between agents. We then prune the graph into a directed acyclic graph (DAG) and decompose the joint prediction task into a sequence of marginal and conditional predictions according to the partial ordering of the DAG, where joint future trajectories are decoded using a directed acyclic graph neural network (DAGNN). We conduct experiments on the INTER-ACTION and Argoverse 2 datasets and demonstrate that FJMP produces more accurate and scene-consistent joint trajectory predictions than non-factorized approaches, especially on the most interactive and kinematically interesting agents. FJMP ranks 1st on the multi-agent test leaderboard of the INTERACTION dataset.

1. Introduction

Multi-agent motion prediction is an important task in a self-driving pipeline, and it involves forecasting the future positions of multiple agents in complex driving environments. Most existing works in multi-agent motion prediction predict a set of marginal trajectories for each agent [14, 19, 21, 30, 32, 40, 45], and thus fail to explicitly account for agent interactions in the future. This results in trajectory predictions that are not consistent with each other. For example, the most likely marginal prediction for two interacting agents may collide with each other, when in reality a negotiation between agents to avoid collision is far more likely. As scene-consistent future predictions are critical for downstream planning, recent work has shifted toward gen-
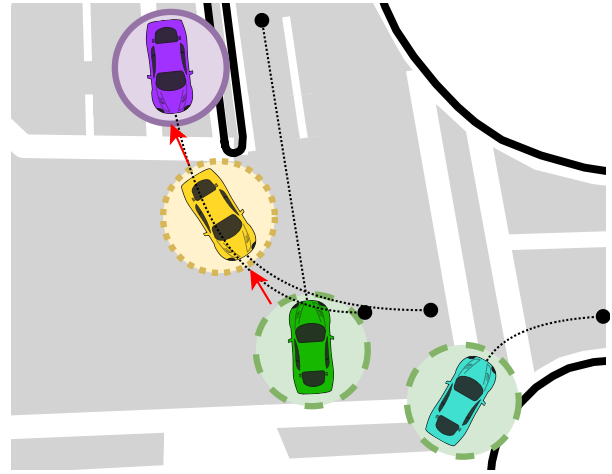

Figure 1. An illustration of the directed acyclic interaction graph, comprised of colored nodes and red arrows. The dotted black lines denote the ground-truth futures over a short time horizon. FJMP first produces marginal predictions for the green (dashed) nodes, followed by a conditional prediction for the yellow (dotted) node and a conditional prediction for the purple (solid) node, with conditioning on the predicted future of the parent nodes in the graph.

erating a set of scene-level, or joint, future trajectory predictions [5, 7, 8, 17, 34, 38], whereby each mode consists of a future trajectory prediction for each agent and the predicted trajectories are consistent with each other.

In this work, we focus on the problem of generating a set of joint future trajectory predictions in multi-agent driving scenarios. Unlike marginal prediction, the joint trajectory prediction space grows exponentially with the number of agents in the scene, which makes this prediction setting particularly challenging. A common approach for this setting is to simultaneously predict the joint futures for all agents in the scene [5, 8, 17, 18, 34]; however, this approach fails to explicitly reason about future interactions in the joint predictions. To address this limitation, recent work has shown that decomposing the joint prediction task of two interacting agents into a marginal prediction for the *influencer* agent

and a conditional prediction for the *reactor* agent, where the reactor's prediction conditions on the predicted future of the influencer, can generate more accurate and scene-consistent joint predictions than methods that generate marginal predictions or simultaneous joint predictions [28, 38]. However, these methods are highly optimized for the joint prediction of only two interacting agents, and they do not efficiently scale to scenes with a large number of interacting agents.

To address these limitations of existing joint motion predictors, in this work we propose **FJMP** – a **F**actorized **J**oint **M**otion **P**rediction framework that efficiently generates joint predictions for driving scenarios with an arbitrarily large number of agents by *factorizing* the joint prediction task into a sequence of marginal and conditional predictions. FJMP models the future scene interaction dynamics as a sparse directed interaction graph, where an edge denotes an explicit interaction between a pair of agents, and the direction of the edge is determined by their influencer-reactor relationship [25, 27, 38], as can be seen in Fig. 1. We propose a mechanism to efficiently prune the interaction graph into a directed acyclic graph (DAG). Joint future trajectory predictions are then decoded as a sequence of marginal and conditional predictions according to the partial ordering of the DAG, whereby marginal predictions are generated for the source node(s) in the DAG and conditional predictions are generated for non-source nodes that condition on the predicted future of their parents in the DAG. To enable this sequential trajectory decoding, we adapt a lightweight directed acyclic graph neural network (DAGNN) [39] architecture for efficiently processing predicted future information through the DAG and decoding the marginal and conditional trajectory predictions. Our main contributions can be summarized as follows:

- We propose FJMP, a novel joint motion prediction framework that generates factorized joint trajectory predictions over sparse directed acyclic interaction graphs. To our knowledge, FJMP is the first framework that enables efficient and scalable factorized joint prediction on scenes with arbitrarily many interacting agents.

- We validate our proposed method on both the multi-agent INTERACTION dataset and the Argoverse 2 dataset and systematically demonstrate that FJMP produces scene-consistent joint predictions for scenes with up to 50 agents that outperform non-factorized approaches, especially on the most interactive and kinematically complex agents. FJMP achieves state-of-the-art performance across several metrics on the challenging multi-agent prediction benchmark of the INTERACTION dataset and ranks 1st on the official leaderboard.

2. Related Work

2.1. Motion Prediction in Driving Scenarios

Given the recent growing interest in autonomous driving, many large-scale motion prediction driving datasets [3, 6, 13, 43, 48] have been publically released, which has enabled rapid progress in the development of data-driven motion prediction methods. Recurrent neural networks (RNNs) are a popular choice for encoding agent trajectories [7, 10, 15–17, 23] and convolutional neural networks (CNNs) are widely used in earlier works to process the birds-eye view (BEV) rasterized encoding of the High-Definition (HD) map [5, 7, 15, 18, 33, 35]. As rasterized HD-map encodings do not explicitly capture the topological structure of the lanes and are constrained by a limited receptive field, recent methods have proposed vectorized [14, 19, 32, 36], lane graph [9, 16, 30], and point cloud [45] representations for the HD-map encoding. Inspired by the success of transformers in both natural language processing [11, 41] and vision [12], several end-to-end transformer-based methods have recently been proposed for motion prediction [18, 20, 21, 32, 34, 36, 49]. However, many of these transformer-based methods are extremely costly in model size and inference speed, which makes them impractical for use in real-world settings. FJMP adopts a LaneGCN-inspired architecture [30] due to its strong performance on competitive benchmarks [6], while retaining a small model size and fast inference speed.

2.2. Interaction Modeling for Motion Prediction

Data-driven methods typically use attention-based mechanisms [17, 18, 30, 32, 34, 49] or graph neural networks (GNNs) [4, 5, 14, 21, 25, 29, 31, 47] to model agent interactions for motion prediction. Recent works have demonstrated the importance of not only modeling agent interactions in the observed agent histories but also reasoning explicitly about the agent interactions that may occur in the future [1, 26–28, 36, 38]. MTR [36] proposes to generate future trajectory hypotheses as an auxiliary task, where the future hypotheses are fed into the interaction module so that it can better reason about future interactions. Multiple works reason about future interactions through predicted pairwise influencer-reactor relationships [27, 28, 38], where the agent who reaches the conflict point first is defined as the influencer, and the reactor otherwise. FJMP uses attention to model interactions in the agent histories and constructs a sparse interaction graph based on pairwise influencer-reactor relationships to model future interactions.

2.3. Joint Motion Prediction

The majority of existing motion prediction systems generate marginal predictions for each agent [2, 9, 10, 14–16, 19–21, 23, 30, 32, 33, 35–37, 40, 42, 45–47, 49]; however, marginal
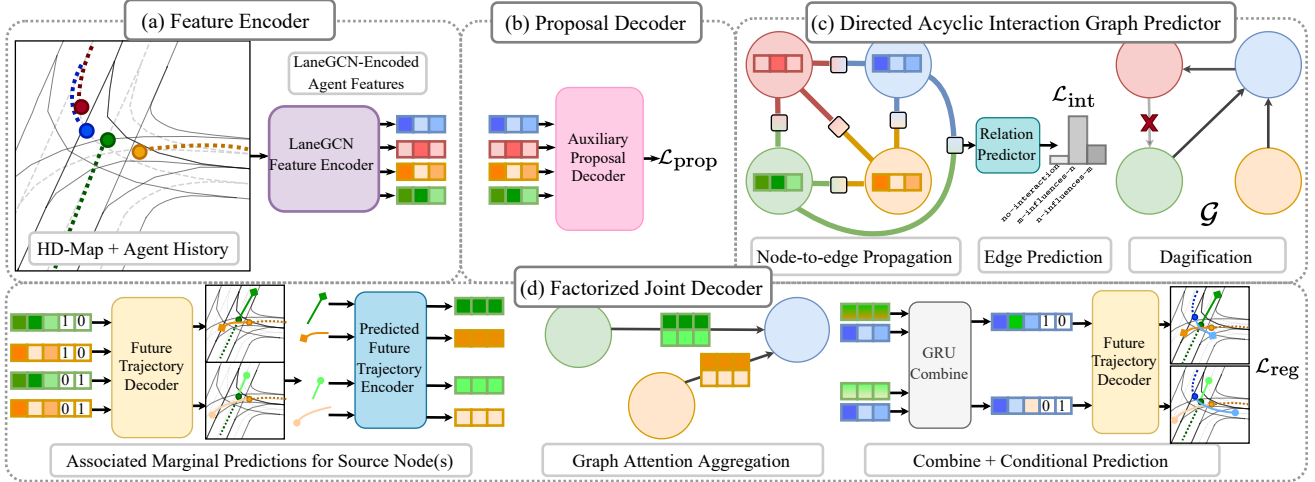
Figure 2. Illustration of the proposed FJMP framework. (a) Agent histories and the HD-Map are first processed by a LaneGCN-inspired feature encoder. (b) During training, the LaneGCN-encoded features are fed into an auxiliary future proposal decoder trained with a regression loss to encourage the LaneGCN features to be *future-aware*. (c) The future-aware LaneGCN-features are processed by a GNN that predicts the pairwise influencer-reactor relationships supervised by a focal loss. A directed interaction graph $\mathcal{G}$ is constructed from the predicted edge probabilities and cycles are removed via an efficient "dagification" procedure. (d). The predicted DAG and future-aware LaneGCN features are fed into a factorized DAGNN-based trajectory decoder (red agent removed for simplicity), which produces K ($K = 2$ shown above) factorized joint futures in parallel and is supervised by a joint regression loss.

predictions lack an association of futures across agents. Recent works have explored generating simultaneous joint predictions [5, 8, 17, 18, 34], but these methods do not explicitly reason about future interactions in the joint predictions. Other works generate joint predictions for two-agent interactive scenarios by selecting K joint futures among all K^2 possible combinations of the marginal predictions [36, 44], which quickly becomes intractable as the number of agents in the scene increases. ScePT [7] proposes to handle the exponentially growing joint prediction space by decomposing joint prediction into the prediction of interactive cliques. However, the density of large cliques imposes a severe computational burden at inference time, which requires ScePT to upper-bound the maximum clique size to 4. To avoid the computational burden associated with dense interaction graphs, FJMP models future interactions as a *sparse* interaction graph consisting only of the strongest interactions, which enables efficient joint decoding over interactive scenarios with many interacting agents.

Our proposed method is most closely related to M2I [38], which first predicts the influencer-reactor relationship between a pair of interacting agents and then generates a marginal prediction for the influencer agent followed by a conditional prediction for the reactor agent. However, we differ from M2I in three critical ways. First, M2I is designed specifically to perform joint prediction of two interacting agents, as their model design assumes one influencer agent and one reactor agent. In contrast, FJMP naturally scales to an arbitrary number of interacting agents, where

an agent may have multiple influencers and influence multiple reactors. Second, M2I requires a costly inference-time procedure that does not scale to multiple agents whereby N conditional predictions are generated for each marginal prediction, resulting in N^2 joint predictions that are pruned to $K = 6$ based on predicted likelihood. On the contrary, FJMP coherently aligns the joint predictions of a given modality through the DAG and directly produces $K = 6$ factorized joint predictions without any required pruning, which allows the system to seamlessly scale to scenes with an arbitrarily large number of agents. Third, M2I uses separate decoders for the marginal and conditional prediction, whereas FJMP decodes both marginal and conditional predictions using the same decoder, making it more parameter-efficient.

3. FJMP

In this section, we describe our proposed factorized joint motion prediction framework, illustrated in Fig. 2.

3.1. Preliminaries

3.1.1 Proposed Joint Factorization

The goal of multi-agent joint motion prediction is to predict the future T_{fut} timesteps of N dynamic agents in a scene given the past motion of the N agents and the structure of the HD-Map. As there are multiple possible futures for a given past, the joint motion prediction task involves predicting $K > 1$ modalities, whereby each modality consists

of a predicted future for each agent in the scene. We let X and Y denote the past trajectories and future trajectories for all N agents in the scene, respectively, where $X_{\mathcal{S}}$ denotes the past trajectories for all agents in the set $\mathcal{S} \subseteq [N]$, and $Y_{\mathcal{S}}$ is defined similarly. Moreover, we let C be an encoding of the HD-Map context.

We first propose to model the future scene interaction dynamics as a DAG $\mathcal{G} = \{\mathcal{V}, \mathcal{E}\}$, where the vertices $\mathcal{V} = [N]$ correspond to the N dynamic agents in the scene, and a directed edge $e_{mn} \in \mathcal{E}$; $m, n \in [N]$, denotes an explicit interaction between agents m and n whereby m is the influencer and n is the reactor of the interaction. We propose to factorize the joint future trajectory distribution $P(Y|X, C)$ over the DAG $\mathcal{G}$ as follows:

$$P(Y|X, C) = \prod_{n=0}^{N-1} P(Y_{\{n\}}|Y_{\mathrm{pa}_{\mathcal{G}}(n)}, X, C), \qquad (1)$$

where $\mathrm{pa}_{\mathcal{G}}(n)$ denotes the set containing the parents of node n in $\mathcal{G}$. Intuitively, the proposed joint factorization can be interpreted as an inductive bias that encourages accounting for the predicted future of the agent(s) that influence agent n when predicting the future of agent n. We hypothesize that this inductive bias will ease the complexity of learning the joint distribution when compared to methods that produce a joint prediction for all N agents simultaneously.

3.1.2 Input Preprocessing

The past trajectory of a given agent is expressed as a sequence of T_{obs} states, which contains the 2D position, the velocity, and the heading of the agent at each timestep. We denote the past state of agent n at timestep $t, t \in [T_{\mathrm{obs}}]$, by $\mathbf{x}_t^n = [\mathbf{p}_t^n, \mathbf{v}_t^n, \psi_t^n]$, where $\mathbf{p}_t^n \in \mathbb{R}^2$ is the position, $\mathbf{v}_t^n \in \mathbb{R}^2$ is the velocity, and $\psi_t^n \in \mathbb{R}$ is the yaw angle. We are also provided the agent type a^n. As in LaneGCN [30], we convert the sequence of 2D positional coordinates of each agent n to a sequence of coordinate displacements: $\hat{\mathbf{p}}_t^n = \mathbf{p}_t^n - \mathbf{p}_{t-1}^n$ for all t. We encode the HD-Map as a lane graph with M nodes, each denoting the location of the midpoint of a lane centerline segment. Using the lane graph construction proposed in LaneGCN, four adjacency matrices, $\{\mathbf{A}_i\}_{i \in \{\mathrm{pre, suc, left, right}\}}$, $\mathbf{A}_i \in \mathbb{R}^{M \times M}$, are calculated to represent the predecessor, successor, left, and right node connectivities in the lane graph, respectively. Our system takes as input the M lane node positional coordinates, the lane node connectivities $\{\mathbf{A}_i\}_{i \in \{\mathrm{pre, suc, left, right}\}}$, and the preprocessed agent history states $[\hat{\mathbf{p}}_t^n, \mathbf{v}_t^n, \psi_t^n]$ for all $n \in [N], t \in [T_{\mathrm{obs}}]$.

3.2. Feature Encoder

To encode the agent history and HD-map data, we employ a LaneGCN backbone [30] with a few key modifications. For processing the agent histories, we replace

LaneGCN's proposed ActorNet architecture with a gated recurrent unit (GRU) module. For processing the HD-Map, we employ the MapNet architecture, which consists of L graph convolutional operators that enrich the lane node features by propagating them through the lane graph. We then employ the FusionNet architecture introduced in LaneGCN [30] for fusing the map and actor features, but we remove the actor-to-lane (A2L) and lane-to-lane (L2L) modules, keeping only the lane-to-actor (L2A) and actor-to-actor (A2A) modules. We observed a minimal loss in performance when removing the A2L and L2L modules, and we benefited from the reduced parameter count. The output of the LaneGCN feature encoder produces a set of map-aware agent features $H = \{\mathbf{h}_n\}_{n \in [N]}$ for each agent.

3.2.1 Auxiliary Proposal Decoder

While the output of the LaneGCN feature encoder provides informative map-aware agent features, the A2A module only considers agent interactions in the observed past trajectories. However, these features will be used downstream to reason about agent interactions in the future, and thus we desire agent feature representations that are *future-aware* – agent features that are predictive of the future. To this end, we propose to regularize the LaneGCN agent feature representations with an auxiliary pretext task that predicts joint future trajectories on top of the LaneGCN-encoded agent features. We adopt a proposal decoder f_{prop}, which decodes K joint future trajectories from the LaneGCN-encoded features $\{\hat{\mathbf{y}}_{\mathrm{prop}, k}^n\}_{k \in [K]} = f_{\mathrm{prop}}(\mathbf{h}_n)$ and it is supervised by a joint regression loss $\mathcal{L}_{\mathrm{prop}}$. More details of the joint regression loss can be found in Appendix A. We hypothesize that the proposed pretext task will regularize the LaneGCN feature representations, so that it contains future context that will be useful for reasoning about future interactions in the downstream modules. We note that the proposal decoder is discarded at inference time and is only used to regularize features during training.

3.3. Directed Acyclic Interaction Graph Predictor

3.3.1 Interaction Graph Predictor

In order to construct the directed acyclic interaction graph, we first must classify the future interaction label between every pair of agents in the scene. This task can be formulated as a classification task where we classify every edge in a fully-connected undirected interaction graph $\mathcal{G}_U = \{\mathcal{V}, \mathcal{E}_U\}$, where each agent corresponds to a node in $\mathcal{V}$. Similar to [25, 27, 28, 38], given an edge $e_{m,n} \in \mathcal{E}_U$, the classification task assumes three labels: `no-interaction`, `m-influences-n`, and `n-influences-m`, where the ground-truth future interaction label is heuristically determined using their ground-truth future trajectories. Concretely, we employ a collision checker to check for a colli-

sion between agents m and n at all pairs of future timesteps (t_m, t_n) where $|t_m - t_n| \leq \epsilon_I$ for some threshold ϵ_I. Details of the collision checker can be found in Appendix F. We let $\mathcal{C}$ denote the set of timestep pairs where a collision is detected. If $|\mathcal{C}| = 0$, then m and n are not interacting and the edge is labeled `no-interaction`. Otherwise, we identify the first such pair of timesteps $(\hat{t}_m, \hat{t}_n) \in \mathcal{C}$ where a collision is detected:

$$(\hat{t}_m, \hat{t}_n) = \underset{(t_m, t_n) \in \mathcal{C}}{\arg\min} \quad \min\{t_m, t_n\}. \qquad (2)$$

If $|\mathcal{C}| > 0$, then there exists a conflict point between the two agents, and the influencer agent is defined as the agent who reaches the conflict point first. Specifically, if $\hat{t}_m < \hat{t}_n$, then we assign the edge the label `m-influences-n`, and otherwise we assign the edge the label `n-influences-m`.

With the heuristic interaction labels, we train a classifier to predict the interaction type on each edge of $\mathcal{G}_U$. We first initialize the node features of $\mathcal{G}_U$ to the future-aware LaneGCN agent features $\mathbf{h}_n$. We then perform a node-to-edge feature propagation step, where for each edge $e_{m,n}$:

$$\mathbf{h}_{m,n}^e = f_{\text{edge}}\Big(\big[\mathbf{h}_m || \mathbf{h}_n || f_{\text{dist}}(\mathbf{p}_{t_c}^m - \mathbf{p}_{t_c}^n) || \mathbf{a}_{m,n} \big] \Big), \quad (3)$$

where f_{edge} and f_{dist} are 2-layer MLPs, $||$ denotes concatenation along the feature dimension, $t_c := T_{\text{obs}} - 1$ is the present timestep, and $\mathbf{a}_{m,n} = f_{\text{type}}([a_m, a_n])$ is the output of a 2-layer MLP f_{type} applied to the agent types a_m, a_n. We then classify the interaction label using a 2-layer MLP f_{int} with a softmax activation:

$$\hat{r}_{m,n} = \text{softmax}(f_{\text{int}}(\mathbf{h}_{m,n}^e)). \qquad (4)$$

The interaction classifier is trained with a focal loss $\mathcal{L}_{\text{int}} = \mathcal{L}_{\text{focal}}^{\gamma, \alpha}(R, \hat{R})$ with hyperparameters γ and α, where $\hat{R}$ is the predicted interaction label distributions and R is the ground-truth interaction labels. From the predicted interaction label distributions, we can construct a directed interaction graph $\mathcal{G} = \{\mathcal{V}, \mathcal{E}\}$ by selecting the interaction label on each edge with the highest predicted probability. For each pair of agents, we add a directed edge from the predicted influencer to the predicted reactor if an interaction is predicted to exist, and no edge is added otherwise.

Dagification. In order to perform factorized joint prediction over the learned directed interaction graph $\mathcal{G}$, we require $\mathcal{G}$ to be a DAG. We propose to remove cycles from $\mathcal{G}$, or "dagify" $\mathcal{G}$, by iterating through the cycles in $\mathcal{G}$ and removing the edges with the lowest predicted probability. We efficiently enumerate the cycles in $\mathcal{G}$ using Johnson's algorithm [22], which has time complexity $O((|\mathcal{V}| + |\mathcal{E}|)(c+1))$, where c is the number of cycles in $\mathcal{G}$. As the directed interaction graphs are typically sparse ($|\mathcal{V}| \approx |\mathcal{E}|$) with a small number of cycles, for our application Johnson's algorithm runs approximately linear in the number of agents in the scene.

3.4. Factorized Joint Trajectory Decoder

Given the future-aware LaneGCN feature encodings $H = \{\mathbf{h}_n\}_{n \in [N]}$ and the directed acyclic interaction graph $\mathcal{G}$, we perform factorized joint prediction according to the unique partial ordering of $\mathcal{G}$. We parameterize the factorized joint trajectory decoder using an adapted directed acyclic graph neural network (DAGNN) [39]. A DAGNN is a recently proposed architecture that is suited specifically for DAG classification tasks. The originally proposed DAGNN framework performs DAG-level classification tasks on top of the representations of the *leaf* nodes in the DAG, where node features are propagated to the leaf nodes sequentially according to the partial ordering of the DAG. Although we desire to process the agents according to the partial ordering of the interaction graph $\mathcal{G}$, we also aim to use the intermediate updated node features of the DAG to generate conditional future trajectory predictions, and thus we adapt the DAGNN design to fit this criterion. We explain first how to produce a factorized joint prediction using the proposed adapted DAGNN decoder, and then how the proposed decoder is extended to produce multiple joint futures.

The factorized decoder first processes the source node(s) $\mathcal{S}$ in parallel. For each source node $s \in \mathcal{S}$, we first decode a marginal future trajectory prediction:

$$\hat{\mathbf{y}}^s = \text{DECODE}(\mathbf{h}_s), \qquad (5)$$

where DECODE is a residual block followed by a linear layer and $\hat{\mathbf{y}}^s \in \mathbb{R}^{2T_{\text{fut}}}$ is the sequence of predicted future trajectory coordinates. We then encode the predicted future trajectories of each source node $s \in \mathcal{S}$:

$$\mathbf{e}_s = \text{ENCODE}(\hat{\mathbf{y}}^s), \qquad (6)$$

where ENCODE is a 3-layer MLP. For each $s \in \mathcal{S}$, the encoding of the predicted future $\mathbf{e}_s$ is then fed along the outgoing edges of s. Namely, after processing the source nodes $\mathcal{S}$, we update the features of the nodes that are next in the partial ordering of $\mathcal{G}$. For every such node n, we perform the following update:

$$\mathbf{h}_n \leftarrow \text{COMB}\left(\text{AGG}\left(\{\mathbf{e}_m + \mathbf{a}_{mn} | m \in \text{pa}_{\mathcal{G}}(n)\} \right), \mathbf{h}_n \right), \qquad (7)$$

where AGG is a neural network that *aggregates* the node features from n's parents and COMB is a neural network that *combines* this aggregated information with n's features to update the feature representation of n with conditional context about the predicted future of n's parents. $\mathbf{a}_{mn} = f_{\text{type}}^{\text{dec}}([a_m, a_n])$ is the output of a 2-layer MLP $f_{\text{type}}^{\text{dec}}$ applied to the agent types. From here, we let $\mathbf{b}_{mn} := \mathbf{e}_m + \mathbf{a}_{mn}$. Similar to DAGNN [39] which uses additive attention, we

Model	Venue	minADE	minFDE	SMR	CrossCol	CMR
THOMAS [17]	ICLR 2022	0.416	0.968	<u>0.179</u>	0.128	0.252
HDGT [21]	-	<u>0.303</u>	<u>0.958</u>	0.194	0.163	0.236
DenseTNT [19]	ICCV 2021	0.420	1.130	0.224	**0.000**	0.224
AutoBot [18]	ICLR 2022	0.312	1.015	0.193	0.043	0.207
Traj-MAE	-	0.304	0.982	0.185	0.032	<u>0.197</u>
FJMP (Ours)	-	**0.277**	**0.917**	**0.174**	<u>0.008</u>	**0.176**

Table 1. Joint prediction results on the INTERACTION multi-agent **test** set. Methods are sorted by the official ranking metric (CMR). For each metric, the best method is **bolded** and the second-best method is <u>underlined</u>. Lower is better for all metrics.

parameterize AGG using graph attention:

$$\mathbf{m}_n := \mathrm{AGG}\left(\{\mathbf{b}_{mn}|m \in \mathrm{pa}_{\mathcal{G}}(n)\}\right) = \sum_{m \in \mathrm{pa}_{\mathcal{G}}(n)} \alpha_{mn}\mathbf{W}_1\mathbf{b}_{mn},$$

(8)

$$\alpha_{mn} = \frac{\exp(\mathrm{LeakyReLU}(\mathbf{a}^\top[\mathbf{W}_1\mathbf{b}_{mn}||\mathbf{W}_2\mathbf{h}_n]))}{\sum_{k \in \mathrm{pa}_{\mathcal{G}}(n)} \exp(\mathrm{LeakyReLU}(\mathbf{a}^\top[\mathbf{W}_1\mathbf{b}_{kn}||\mathbf{W}_2\mathbf{h}_n]))}.$$

(9)

COMB is parameterized by a GRU recurrent module:

$$\mathbf{h}_n \leftarrow \mathrm{COMB}(\mathbf{m}_n, \mathbf{h}_n) = \mathrm{GRU}(\mathbf{m}_n, \mathbf{h}_n).$$

(10)

As the aggregated message $\mathbf{m}_n$ provides conditional context for updating the representation of $\mathbf{h}_n$, $\mathbf{m}_n$ is treated as the input and $\mathbf{h}_n$ is treated as the hidden state. It is important to note that the roles of the input and hidden state are reversed in the original DAGNN design [39]. The updated representation $\mathbf{h}_n$ for node n is now imbued with conditional context about the predicted future of the parent(s) of n, which can now be fed into DECODE to produce conditional future predictions for agent n. We sequentially continue the process of encoding, aggregating, combining, and decoding according to the DAG's partial order until all nodes in the DAG have a future trajectory prediction. The future trajectory predictions of all nodes are then conglomerated to attain a factorized joint prediction.

Multiple Futures. To extend the DAGNN factorized decoder to produce multiple factorized joint predictions, we simply process K copies of $H = \{\mathbf{h}_n\}_{n \in [N]}$ through the DAG in parallel. To ensure each copy of H generates a different set of futures, we concatenate a one-hot encoding of the modality with $\mathbf{h}_n$ along the feature dimension, for each $n \in [N]$, prior to it being fed into DECODE. This approach is similar to the multiple futures approach proposed in SceneTransformer [34] and we found it to work well in our application. We train the factorized joint predictor to produce diverse multiple futures by training with a winner-takes-all joint regression loss $\mathcal{L}_{\mathrm{reg}}$. More details about the regression loss can be found in Appendix A.

3.5. Training Details

We first train the interaction graph predictor separately using its own feature encoder weights. The interaction graph predictor is trained via gradient descent, where the loss function is defined by:

$$\mathcal{L}_1 = \mathcal{L}_{\mathrm{int}} + \mathcal{L}_{\mathrm{prop}}.$$

(11)

Next, we train the factorized joint decoder using its own feature encoder weights, where the interaction graphs $\mathcal{G}$ are generated with the trained interaction graph predictor. The factorized joint predictor is trained via gradient-descent, where the loss function is defined by:

$$\mathcal{L}_2 = \mathcal{L}_{\mathrm{reg}} + \mathcal{L}_{\mathrm{prop}}.$$

(12)

Similar to M2I [38], we employ teacher forcing of the influencer future trajectories during training, which helps to learn the proper influencer-reactor dynamics.

4. Experiments

Datasets. We evaluate FJMP on the INTERACTION v1.2 multi-agent dataset and the Argoverse 2 dataset, as both have multi-agent evaluation schemes for scenes with many interacting agents and require predicting joint futures for scenes with up to 40 and 56 agents, respectively. However, currently, only INTERACTION has a public benchmark for multi-agent joint prediction. Argoverse 2 contains *scored* and *focal* actors, which are high-quality tracks near the ego vehicle; and *unscored* actors, which are high-quality tracks more than 30 m from the ego vehicle. We evaluate FJMP on (i) only the scored and focal actors; and (ii) all scored, focal, and unscored actors. More details about these datasets can be found in Appendix D.

Evaluation Metrics. We report the following joint prediction metrics: **minFDE** is the final displacement error (FDE) between the ground-truth and closest predicted future trajectory endpoint from the K joint predictions; **minADE** is the average displacement error (ADE) between the ground-truth and closest predicted future trajectory from the K joint predictions; **SMR** is the minimum proportion

Dataset	Actors Evaluated	Model	minFDE	minADE	SCR	SMR	iminFDE	iminADE	iminFDE$_3$	iminADE$_3$	iminFDE$_5$	iminADE$_5$
Interaction	-	Non-Factorized	0.649	0.201	0.005	0.091	0.692	0.213	0.790	0.242	0.859	0.263
		FJMP	0.634	0.196	0.004	0.085	0.676	0.207	0.763	0.234	0.829	0.254
		Δ	0.015	0.005	0.001	0.006	0.016	0.006	0.027	0.008	0.030	0.009
Argoverse 2	Scored	Non-Factorized	1.974	0.838	-	0.354	2.888	1.200	3.300	1.354	3.476	1.419
		FJMP	1.956	0.833	-	0.346	2.812	1.177	3.207	1.322	3.368	1.383
		Δ	0.018	0.005	-	0.008	0.076	0.023	0.093	0.032	0.108	0.036
	All	Non-Factorized	2.013	0.832	-	0.342	3.233	1.300	3.686	1.468	3.868	1.534
		FJMP	1.991	0.823	-	0.338	3.164	1.274	3.607	1.438	3.781	1.500
		Δ	0.022	0.009	-	0.004	0.069	0.026	0.079	0.030	0.087	0.034

Table 2. Non-Factorized Baseline vs. FJMP performance on joint metrics on the INTERACTION and Argoverse 2 validation sets. Lower is better for all metrics. Argoverse 2 lacks agent bounding box information, so SCR is not computed. Δ denotes the difference in performance between FJMP and the Non-Factorized baseline.

of agents whose predicted trajectories "miss" the ground-truth from the K joint predictions, where a miss is defined in Appendix G; and **SCR** is the proportion of modalities where two or more agents collide. The INTERACTION test set additionally reports two joint prediction metrics: **CrossCol** is the same as SCR but does not count ego collisions, and **CMR** is the same as SMR but only considers modalities without non-ego collisions. These six joint prediction metrics do not necessarily capture the performance on the most interactive and challenging cases in the dataset, which is critically important for benchmarking and improving motion prediction systems. To address this limitation, we propose two new interactive metrics: (i) **iminFDE** first identifies the modality k with minimum FDE over all the agents in the scene and then computes the FDE of modality k only over agents that are *interactive*, which we heuristically define as agents with at least one incident edge in the ground-truth sparse interaction graph, where $\epsilon_I = 2.5$ s. (ii) **iminADE** is defined similarly. We found that many of the interactive cases in the datasets contain kinematically simple cases where agents exhibit simple leader-follower behaviour. To evaluate the *challenging* interactive cases, we further remove interactive agents in our evaluation that attain less than d meters in FDE with a constant velocity model. These metrics are denoted **iminFDE$_d$** and **iminADE$_d$**, where we report $d = 3, 5$. Please see Appendix H for more details.

Implementation Details. Our models are trained on 4 NVIDIA Tesla V100 GPUs using the Adam optimizer [24]. The interaction graph predictor and factorized joint decoder are trained with the same hyperparameters. For INTERACTION, we set the batch size to 64 and train for 50 epochs with a learning rate of 1e-3, step-decayed by a factor of 1/5 at epochs 40 and 48. For Argoverse 2, we set the batch size to 128 and train for 36 epochs with a learning rate of 1e-3, step-decayed by a factor of 1/10 at epoch 32. For training the interaction graph predictor, we set $\gamma = 5$ and $\alpha = [1, 4, 4]$. We set $\epsilon_I = 2.5$ s on INTERACTION, and as interactions are comparatively more sparse in Argoverse 2,

we set $\epsilon_I = 6$ s in Argoverse 2. As bounding-box information is not provided with Argoverse 2, the collision checker used to construct interaction labels uses a predefined length and width for each agent type, as listed in Appendix D.2. Our INTERACTION and Argoverse 2 models train in 10 and 15 hours. See Appendix E for more details.

Model	minFDE	minADE	SMR	Prop. Edges	Inf. Time (s)
Non-Factorized	0.649	0.201	0.091	-	**0.0043**
FJMP (Dense)	**0.628**	**0.195**	**0.083**	0.192	0.0531
FJMP	0.631	0.196	0.084	0.048	0.0218

Table 3. Comparison of sparse vs. dense interaction graphs on the INTERACTION validation set. The FJMP model is trained and evaluated using the ground-truth sparse interaction graphs, and FJMP (Dense) is trained and evaluated using dense ground-truth interaction graphs attained via the M2I [38] labeling heuristic. **Prop. Edges** measures the proportion of agent pairs connected in the ground-truth training interaction graphs. **Inf. Time** is the inference time per validation scene on 1 NVIDIA Tesla V100 GPU.

Methods under Comparison. We compare FJMP against the top-performing methods on the INTERACTION multi-agent test set leaderboard [17–19, 21]. FJMP is the only method on the leaderboard that performs factorized joint prediction. To measure the improvement of FJMP over non-factorized approaches, we compare FJMP against a baseline called **Non-Factorized**, which computes K simultaneous joint futures from the feature representations output by the proposed LaneGCN feature encoder.

Results. The joint prediction results for the top-performing methods on the INTERACTION multi-agent test set are shown in Tab. 1, where CMR is the official ranking metric. FJMP performs the best on all evaluated joint metrics with the exception of CrossCol; however, we note that DenseTNT [19] uses a collision-check post-processing scheme to prevent predicted collisions, which explains the 0 cross-collision rate. Crucially, FJMP produces joint predictions that are both more accurate—as demonstrated by its superior performance on minADE, minFDE, and SMR—

and more scene-consistent—as demonstrated by its near-zero collision rate—than non-factorized approaches, which highlights the benefit of the proposed joint factorization.

Model	Prop?	TF?	minFDE	minADE	iminFDE	iminADE
Non-Factorized	✗	✗	1.974	0.838	2.888	1.200
FJMP	✗	✗	1.978	0.843	2.863	1.195
FJMP	✗	✓	1.987	0.844	2.885	1.202
FJMP	✓	✗	<u>1.970</u>	<u>0.837</u>	<u>2.855</u>	<u>1.186</u>
FJMP	✓	✓	**1.956**	**0.833**	**2.812**	**1.177**

Table 4. Ablation study of FJMP on Argoverse 2 validation set on scored agents. **Prop?** denotes whether we include the proposal decoder during training. **TF?** denotes whether we teacher-force the influencer trajectories during training.

Table 2 reports validation results on the INTERACTION and Argoverse 2 datasets, where we compare FJMP against the baseline method without joint factorization. For Argoverse 2, we have two evaluation schemes: (i) we evaluate the joint predictions of the scored and focal agents (Scored), and (ii) we evaluate the joint predictions of the scored, unscored, and focal agents (All) to demonstrate its scalability to scenes with a large number of agents. The results show that the proposed joint factorized predictor consistently provides an improvement in performance over the non-factorized baseline. We expect that FJMP improves the most over the baseline on the interactive cases in the dataset, as the proposed factorization directly enables conditioning the reactor predictions on the predicted futures of their influencers. Importantly, we note that for scenes with no predicted interactions, the factorization becomes a product of marginal predictions and thus FJMP reduces to the non-factorized prediction. As expected, the relative improvement of FJMP over the baseline is larger on the interactive and kinematically interesting cases, as demonstrated by a larger performance improvement on the interactive minFDE/minADE metrics. This indicates that the performance improvement from the joint factorization concentrates on the challenging interactive cases, while still producing accurate joint predictions for the full scene. We refer readers to Appendix L for qualitative examples demonstrating the benefit of factorized prediction.

Table 3 uses the INTERACTION dataset to ablate the design choice of representing the interaction graph sparsely with only the strongest pairwise interactions as edges in the graph. We compare FJMP against a variant of FJMP that uses a different labeling heuristic for the interaction graph, resulting in *denser* interaction graphs. Namely, FJMP (Dense) uses the M2I [38] heuristic: for each pair of agents, an interaction is defined to exist if any pair of future trajectory coordinates in the future trajectory horizon is within a threshold Euclidean distance of each other, where the threshold is taken to be the sum of the lengths of the two agents. The influencer-reactor relationship is determined by who reaches the conflict point first. We found that the M2I heuristic often adds several unnecessary edges, especially in congested scenes—as exemplified in Appendix J. We train and evaluate the FJMP models in Tab. 3 using the ground-truth interaction graphs to precisely compare the different labeling heuristics. The results show that the dense (M2I) interaction graph improves very slightly over the sparse interaction graph; however, we retain most of the improvement over the non-factorized baseline with the sparse interaction graph, which indicates that modeling only the strongest interactions is sufficient to see most of the improvement with joint factorization. Moreover, the sparse interactions contain 75% fewer edges than the dense interaction graph, which accelerates inference by more than 2x.

Table 4 conducts an ablation study on Argoverse 2 where we analyze the effect of using the auxiliary proposal decoder and teacher forcing of the influencer's future trajectories during training. The results indicate that both the auxiliary proposal decoder and teacher forcing is critical for allowing the model to reason appropriately about the influencer-reactor future dynamics. Notably, without the proposal decoder (rows 2 and 3 in Tab. 4), FJMP performs similarly to the non-factorized baseline, which we hypothesize is because the LaneGCN-encoded features do not contain the necessary future context to reason appropriately about the future interactions. Teacher forcing also provides an additional performance benefit by removing the spurious noise in the predicted influencer trajectories, so that the model better learns the proper influencer-reactor dynamics during training.

5. Conclusion

In this paper, we propose FJMP, a factorized joint motion prediction framework for multi-agent interactive driving scenarios. FJMP models the future scene interaction dynamics as a sparse directed acyclic interaction graph, which enables efficient factorized joint prediction. We demonstrate clear performance improvements with our factorized design on the Argoverse 2 and INTERACTION datasets and perform state-of-the-art on the challenging multi-agent INTERACTION benchmark.

Limitations The proposed framework adopts a heuristic labeling scheme to determine the ground-truth interaction graph. We observe a performance-efficiency tradeoff with a significantly denser interaction graph; however, there may exist better heuristics for classifying pairwise future interactions that retain most of the sparsity of the interaction graph without trading off performance. Moreover, long chains of leader-follower behaviour in congested traffic may require costly sequential processing with our method. Finding mechanisms to prune the interaction graph to best trade-off performance and efficiency is a direction we plan to explore in future work.

References

[1] Yutong Ban, Xiao Li, Guy Rosman, Igor Gilitschenski, Ozanan R. Meireles, Sertac Karaman, and Daniela Rus. A deep concept graph network for interaction-aware trajectory prediction. In *Proceedings of the International Conference on Robotics and Automation (ICRA)*, 2022. 2

[2] Prarthana Bhattacharyya, Chengjie Huang, and Krzysztof Czarnecki. Ssl-lanes: Self-supervised learning for motion forecasting in autonomous driving. abs/2206.14116, 2022. 2

[3] Holger Caesar, Varun Bankiti, Alex H. Lang, Sourabh Vora, Venice Erin Liong, Qiang Xu, Anush Krishnan, Yu Pan, Giancarlo Baldan, and Oscar Beijbom. nuscenes: A multi-modal dataset for autonomous driving. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 2

[4] Sergio Casas, Cole Gulino, Renjie Liao, and Raquel Urtasun. Spagnn: Spatially-aware graph neural networks for relational behavior forecasting from sensor data. In *Proceedings of the International Conference on Robotics and Automation (ICRA)*, 2020. 2

[5] Sergio Casas, Cole Gulino, Simon Suo, Katie Luo, Renjie Liao, and Raquel Urtasun. Implicit latent variable model for scene-consistent motion forecasting. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2020. 1, 2, 3

[6] Ming-Fang Chang, John Lambert, Patsorn Sangkloy, Jagjeet Singh, Slawomir Bak, Andrew Hartnett, De Wang, Peter Carr, Simon Lucey, Deva Ramanan, and James Hays. Argoverse: 3d tracking and forecasting with rich maps. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 2

[7] Yuxiao Chen, Boris Ivanovic, and Marco Pavone. Scept: Scene-consistent, policy-based trajectory predictions for planning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 1, 2, 3

[8] Alexander Cui, Sergio Casas, Abbas Sadat, Renjie Liao, and Raquel Urtasun. Lookout: Diverse multi-future prediction and planning for self-driving. In *Proceedings of the IEEE/CVF IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 16087–16096, 2021. 1, 3

[9] Alexander Cui, Sergio Casas, Kelvin Wong, Simon Suo, and Raquel Urtasun. Gorela: Go relative for viewpoint-invariant motion forecasting. abs/2211.02545, 2022. 2

[10] Nachiket Deo, Eric M. Wolff, and Oscar Beijbom. Multimodal trajectory prediction conditioned on lane-graph traversals. In *Proceedings of the Conference on Robot Learning (CoRL)*, 2021. 2

[11] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: pre-training of deep bidirectional transformers for language understanding. In Jill Burstein, Christy Doran, and Thamar Solorio, editors, *Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NACCL-HLT)*, 2019. 2

[12] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2021. 2

[13] Scott Ettinger, Shuyang Cheng, Benjamin Caine, Chenxi Liu, Hang Zhao, Sabeek Pradhan, Yuning Chai, Ben Sapp, Charles R. Qi, Yin Zhou, Zoey Yang, Aurelien Chouard, Pei Sun, Jiquan Ngiam, Vijay Vasudevan, Alexander McCauley, Jonathon Shlens, and Dragomir Anguelov. Large scale interactive motion forecasting for autonomous driving : The waymo open motion dataset. In *Proceedings of the IEEE/CVF IEEE/CVF International Conference on Computer Vision (ICCV)*, 2021. 2

[14] Jiyang Gao, Chen Sun, Hang Zhao, Yi Shen, Dragomir Anguelov, Congcong Li, and Cordelia Schmid. Vectornet: Encoding HD maps and agent dynamics from vectorized representation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 1, 2

[15] Thomas Gilles, Stefano Sabatini, Dzmitry Tsishkou, Bogdan Stanciulescu, and Fabien Moutarde. HOME: heatmap output for future motion estimation. In *Proceedings of the IEEE International Intelligent Transportation Systems Conference (ITSC)*, 2021. 2

[16] Thomas Gilles, Stefano Sabatini, Dzmitry Tsishkou, Bogdan Stanciulescu, and Fabien Moutarde. GOHOME: graph-oriented heatmap output for future motion estimation. In *Proceedings of the International Conference on Robotics and Automation (ICRA)*, 2022. 2

[17] Thomas Gilles, Stefano Sabatini, Dzmitry Tsishkou, Bogdan Stanciulescu, and Fabien Moutarde. THOMAS: trajectory heatmap output with learned multi-agent sampling. In *International Conference on Learning Representations (ICLR)*, 2022. 1, 2, 3, 6, 7

[18] Roger Girgis, Florian Golemo, Felipe Codevilla, Martin Weiss, Jim Aldon D'Souza, Samira Ebrahimi Kahou, Felix Heide, and Christopher Pal. Latent variable sequential set transformers for joint multi-agent motion prediction. In *International Conference on Learning Representations (ICLR)*, 2022. 1, 2, 3, 6, 7

[19] Junru Gu, Chen Sun, and Hang Zhao. Densetnt: End-to-end trajectory prediction from dense goal sets. In *Proceedings of the IEEE/CVF International Conference on Computer Vision, (ICCV)*, 2021. 1, 2, 6, 7

[20] Zhiyu Huang, Xiaoyu Mo, and Chen Lv. Multi-modal motion prediction with transformer-based neural network for autonomous driving. In *Proceedings of the International Conference on Robotics and Automation (ICRA)*, 2022. 2

[21] Xiaosong Jia, Penghao Wu, Li Chen, Hongyang Li, Yu Liu, and Junchi Yan. HDGT: heterogeneous driving graph transformer for multi-agent trajectory prediction via scene encoding. abs/2205.09753, 2022. 1, 2, 6, 7

[22] Donald B. Johnson. Finding all the elementary circuits of a directed graph. *SIAM J. Comput.*, 4(1):77–84, 1975. 5

[23] Byeoungdo Kim, SeongHyeon Park, Seokhwan Lee, Elbek Khoshimjonov, Dongsuk Kum, Junsoo Kim, Jeong Soo Kim,

and Jun Won Choi. Lapred: Lane-aware prediction of multi-modal future trajectories of dynamic agents. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. 2

[24] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2015. 7

[25] Sumit Kumar, Yiming Gu, Jerrick Hoang, Galen Clark Haynes, and Micol Marchetti-Bowick. Interaction-based trajectory prediction over a hybrid traffic graph. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2021. 2, 4

[26] Yen-Ling Kuo, Xin Huang, Andrei Barbu, Stephen G. McGill, Boris Katz, John J. Leonard, and Guy Rosman. Trajectory prediction with linguistic representations. In *Proceedings of the International Conference on Robotics and Automation (ICRA)*, 2022. 2

[27] Donsuk Lee, Yiming Gu, Jerrick Hoang, and Micol Marchetti-Bowick. Joint interaction and trajectory prediction for autonomous driving using graph neural networks. abs/1912.07882, 2019. 2, 4

[28] Ding Li, Qichao Zhang, Shuai Lu, Yifeng Pan, and Dongbin Zhao. Conditional goal-oriented trajectory prediction for interacting vehicles with vectorized representation. abs/2210.15449, 2022. 2, 4

[29] Jiachen Li, Fan Yang, Hengbo Ma, Srikanth Malla, Masayoshi Tomizuka, and Chiho Choi. RAIN: reinforced hybrid attention inference network for motion forecasting. In *Proceedings of the IEEE/CVF International Conference on Computer Vision, (ICCV)*, 2021. 2

[30] Ming Liang, Bin Yang, Rui Hu, Yun Chen, Renjie Liao, Song Feng, and Raquel Urtasun. Learning lane graph representations for motion forecasting. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2020. 1, 2, 4

[31] Yongkang Liu, Xuewei Qi, Emrah Akin Sisbot, and Kentaro Oguchi. Multi-agent trajectory prediction with graph attention isomorphism neural network. In *Proceedings of the IEEE Intelligent Vehicles Symposium (IV)*, 2022. 2

[32] Yicheng Liu, Jinghuai Zhang, Liangji Fang, Qinhong Jiang, and Bolei Zhou. Multimodal motion prediction with stacked transformers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. 1, 2

[33] Sriram Narayanan, Ramin Moslemi, Francesco Pittaluga, Buyu Liu, and Manmohan Chandraker. Divide-and-conquer for lane-aware diverse trajectory prediction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. 2

[34] Jiquan Ngiam, Benjamin Caine, Vijay Vasudevan, Zhengdong Zhang, Hao-Tien Lewis Chiang, Jeffrey Ling, Rebecca Roelofs, Alex Bewley, Chenxi Liu, Ashish Venugopal, David Weiss, Benjamin Sapp, Zhifeng Chen, and Jonathon Shlens. Scene transformer: A unified multi-task model for behavior prediction and planning. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2021. 1, 2, 3, 6

[35] Seong Hyeon Park, Gyubok Lee, Jimin Seo, Manoj Bhat, Minseok Kang, Jonathan Francis, Ashwin R. Jadhav, Paul Pu Liang, and Louis-Philippe Morency. Diverse and admissible trajectory forecasting through multimodal context understanding. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2020. 2

[36] Shaoshuai Shi, Li Jiang, Dengxin Dai, and Bernt Schiele. Motion transformer with global intention localization and local movement refinement. abs/2209.13508, 2022. 2, 3

[37] Haoran Song, Di Luan, Wenchao Ding, Michael Yu Wang, and Qifeng Chen. Learning to predict vehicle trajectories with model-based planning. In *Proceedings of the Conference on Robot Learning (CoRL)*, 2021. 2

[38] Qiao Sun, Xin Huang, Junru Gu, Brian C. Williams, and Hang Zhao. M2I: from factored marginal trajectory prediction to interactive prediction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 1, 2, 3, 4, 6, 7, 8

[39] Veronika Thost and Jie Chen. Directed acyclic graph neural networks. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2021. 2, 5, 6

[40] Balakrishnan Varadarajan, Ahmed Hefny, Avikalp Srivastava, Khaled S. Refaat, Nigamaa Nayakanti, Andre Cornman, Kan Chen, Bertrand Douillard, Chi-Pang Lam, Dragomir Anguelov, and Benjamin Sapp. Multipath++: Efficient information fusion and trajectory aggregation for behavior prediction. In *Proceedings of the International Conference on Robotics and Automation (ICRA)*, 2022. 1, 2

[41] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2017. 2

[42] Jingke Wang, Tengju Ye, Ziqing Gu, and Junbo Chen. LTP: lane-based trajectory prediction for autonomous driving. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 2

[43] Benjamin Wilson, William Qi, Tanmay Agarwal, John Lambert, Jagjeet Singh, Siddhesh Khandelwal, Bowen Pan, Ratnesh Kumar, Andrew Hartnett, Jhony Kaesemodel Pontes, Deva Ramanan, Peter Carr, and James Hays. Argoverse 2: Next generation datasets for self-driving perception and forecasting. In *Proceedings of the Neural Information Processing Systems Track on Datasets and Benchmarks 1*, 2021. 2

[44] David Wu and Yunnan Wu. Air^2 for interaction prediction. abs/2111.08184, 2021. 3

[45] Maosheng Ye, Tongyi Cao, and Qifeng Chen. TPCN: temporal point cloud networks for motion forecasting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. 1, 2

[46] Maosheng Ye, Jiamiao Xu, Xunnong Xu, Tongyi Cao, and Qifeng Chen. DCMS: motion forecasting with dual consistency and multi-pseudo-target supervision. abs/2204.05859, 2022. 2

[47] Wenyuan Zeng, Ming Liang, Renjie Liao, and Raquel Urtasun. Lanercnn: Distributed representations for graph-centric motion forecasting. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2021. 2

[48] Wei Zhan, Liting Sun, Di Wang, Haojie Shi, Aubrey Clausse, Maximilian Naumann, Julius Kümmerle, Hendrik

Königshof, Christoph Stiller, Arnaud de La Fortelle, and Masayoshi Tomizuka. INTERACTION dataset: An international, adversarial and cooperative motion dataset in interactive driving scenarios with semantic maps. *CoRR*, abs/1910.03088, 2019. 2

[49] 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 (CVPR)*, 2022. 2

A. Loss Functions

The loss function for training the factorized joint trajectory decoder (Sec. 3.4) is defined by:

$$\mathcal{L}_2 = \mathcal{L}_{\text{reg}} + \mathcal{L}_{\text{prop}}, \tag{13}$$

where $\mathcal{L}_{\text{reg}}(\{\hat{Y}_k\}_{k \in [K]}, Y) := \mathcal{L}_{\ell_1}(\{\hat{Y}_k\}_{k \in [K]}, Y)$ is a scene-level smooth ℓ_1 regression loss applied to the best modality of K joint modalities $\{\hat{Y}_k\}_{k \in [K]}$, where the best modality attains the minimum loss:

$$\mathcal{L}_{\ell_1}(\{\hat{Y}_k\}_{k \in [K]}, Y) = \min_{k \in [K]} \frac{1}{A \cdot T_{\text{fut}}} \sum_{a \in [A]} \sum_{t \in [T_{\text{fut}}]} \text{reg}(\hat{Y}_{t,k}^a - Y_t^a), \tag{14}$$

where Y denotes the ground-truth future trajectory coordinates of all A agents in the scene, $\text{reg}(\mathbf{x}) = \sum_i d(x_i)$, x_i is the i'th element of $\mathbf{x}$, and $d(x)$ is the smooth ℓ_1 loss defined by:

$$d(x) = \begin{cases} 0.5x^2, & \text{if } ||x||_1 \leq 1 \\ ||x||_1 - 0.5, & \text{otherwise.} \end{cases} \tag{15}$$

Similarly, the auxiliary decoder loss $\mathcal{L}_{\text{prop}}$ is a scene-level smooth ℓ_1 loss applied to the best of K joint proposals $\{\hat{Y}_k^{\text{prop}}\}_{k \in [K]}$:

$$\mathcal{L}_{\text{prop}}(\{\hat{Y}_k^{\text{prop}}\}_{k \in [K]}, Y) := \mathcal{L}_{\ell_1}(\{\hat{Y}_k^{\text{prop}}\}_{k \in [K]}, Y). \tag{16}$$

We use the auxiliary proposal loss $\mathcal{L}_{\text{prop}}$ for training both the interaction graph predictor ($\mathcal{L}_1$ in Sec. 3.5) and the factorized joint decoder ($\mathcal{L}_2$ in Sec. 3.5) as both modules require explicit reasoning about interactions in the future trajectories, and thus *future-aware* agent features are beneficial for both modules.

B. FJMP System Diagram

B.1. Training Time

Figure 3 illustrates a high-level schematic of the FJMP architecture training stages at training time. We note that Feature Encoder 1 and Feature Encoder 2 consist of the same architecture as described in Sec. 3.2, but use separate weights.
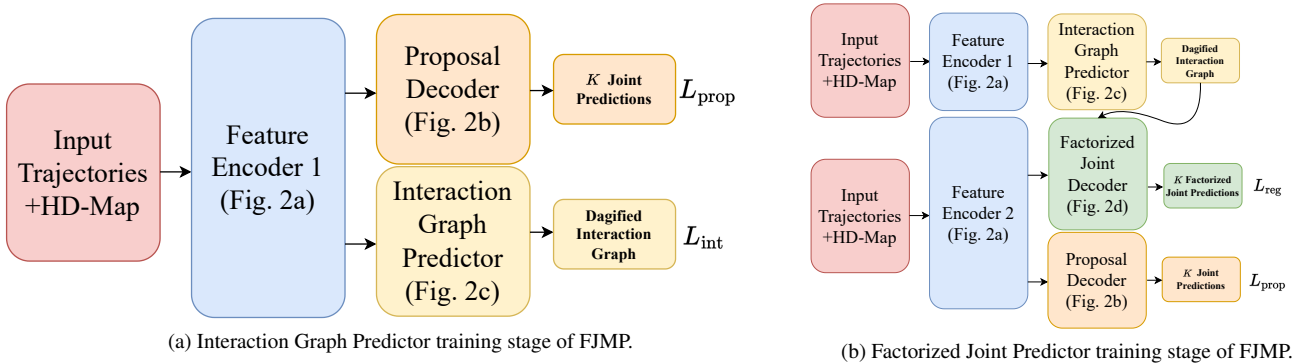

(a) Interaction Graph Predictor training stage of FJMP.

(b) Factorized Joint Predictor training stage of FJMP.

Figure 3. High-level schematic of the training stages of FJMP.

B.2. Inference Time

Figure 4 illustrates a high-level schematic of the FJMP architecture and data flow at inference time. We note that at inference time the proposal decoders are removed.
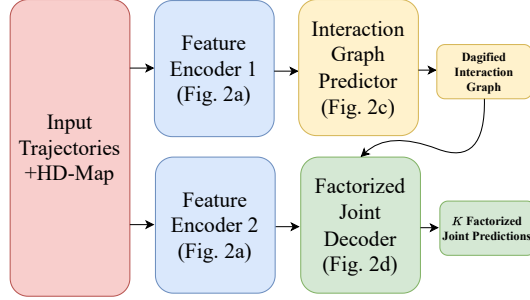
12

Figure 4. High-level schematic of the FJMP architecture at inference time.

C. Non-Factorized Baseline

We explain the non-factorized baseline described in Sec. 4 in more detail. The non-factorized baseline uses the same feature encoder architecture as FJMP, but the factorized joint decoder is replaced with a DECODE module consisting of a residual block and linear layer for simultaneously decoding K joint future trajectory coordinates, where diverse futures are obtained by appending a one-hot encoding of the modality index to the agent feature representation before feeding it into DECODE, as is done in FJMP. The DECODE module is the same architecture as the DECODE module used in FJMP. The non-factorized baseline is trained with the scene-level winner-takes-all smooth ℓ_1 loss $\mathcal{L}_{\text{reg}}$ that is described in Appendix A. The non-factorized baseline is trained with the same training hyperparameters as FJMP.

D. Datasets

D.1. INTERACTION

INTERACTION requires predicting 3 seconds into the future given 1 second of past observations sampled at 10 Hz. INTERACTION contains 47,584 training scenes, 11,794 validation scenes, and 2,644 test scenes. A scene consists of a 4 s sequence of observations (3 s past, 1 s future) for each agent. INTERACTION contains pedestrians, bicyclists, and vehicles as context agents but only requires predicting vehicles in their multi-agent challenge.

D.2. Argoverse 2

Argoverse 2 requires predicting 6 seconds into the future given 5 seconds of past observations sampled at 10 Hz. Argoverse 2 contains 199,908 training scenes and 24,988 validation scenes. A scene consists of an 11 s sequence of observations (5 s past, 6 s future) for each agent. Argoverse 2 requires predicting 5 agent types: vehicle, pedestrian, bicyclist, motorcyclist, and bus. As bounding box length/width information is not provided in the Argoverse 2 dataset, we use the following predefined length/width in meters for each agent type to construct the interaction labels (length/width): vehicle (4.0/2.0), pedestrian (0.7/0.7), bicyclist (2.0/0.7), motorcyclist (2.0/0.7), bus (12.5/2.5).

E. Training Details

E.1. INTERACTION

The hidden dimension of FJMP is 128 except for the GRU history encoder, which has a hidden dimension of 256. The output of the GRU encoder is mapped from dimension 256 to 128 with a linear layer. We set $K = 6$. During training, we center and rotate the scene on a random agent, as an input normalization step. During validation and test time, we center and rotate the scene on the agent closest to the centroid of the agents' current positions. We use 2 MapNet layers, 2 L2A layers, and 2 A2A layers, where the L2A and A2A distance thresholds are set to 20 m and 100 m, respectively. We use all agents in the scene for context that contains a ground-truth position at the present timestep. As centerline information is not provided in INTERACTION, for each lanelet we interpolate P evenly-spaced centerline points, where $P = \min\{10, \max\{L, R\}\}$ and L, R are the number of points on the lanelet's left and right boundaries, respectively; that is, we restrict long lanelets to have a maximum of 10 evenly-spaced centerline points. At validation time, we consider for evaluation all vehicles that contain a ground-truth position at both the present and final timesteps.

Actors Evaluated	Model	SMR$_{\text{Argoverse2}}$
Scored	Non-Factorized	0.267
	FJMP	0.260
	Δ	0.007
All	Non-Factorized	0.262
	FJMP	0.258
	Δ	0.004

Table 5. Non-Factorized Baseline vs. FJMP performance on Argoverse 2 SMR metric on the Argoverse 2 validation set. Δ denotes the difference in performance between FJMP and the Non-Factorized baseline.

E.2. Argoverse 2

The details in Appendix E.1 apply to Argoverse 2 with the following exceptions. At validation time, we center on the ego vehicle. We increase the number of MapNet layers to 4 in Argoverse 2 to handle the larger amount of unique roadway. The L2A threshold is set to 10 m as the centerline points are comparatively more dense in Argoverse 2 than in INTERACTION. We use all scored, unscored, and focal agents in the scene for context that contains a ground-truth position at the present timestep. In the *Scored* validation setting (see Tab. 2), we consider for evaluation all scored and focal agents with a ground-truth position at both the present and final timesteps. In the *All* validation setting (see Tab. 2), we consider for evaluation all scored, unscored, and focal agents with a ground-truth position at both the present and final timesteps.

F. Collision Checker

To construct the interaction labels as described in Sec. 3.3, a collision checker is used to identify collisions between all pairs of timesteps in the future trajectories. We use the collision checker provided with the INTERACTION dataset. At each timestep, the collision checker defines each agent by a list of circles, and two agents are defined as colliding if the Euclidean distance between any two circles' origins of the given two agents is lower than the following threshold:

$$\epsilon_C := \frac{w_i + w_j}{\sqrt{3.8}}, \tag{17}$$

where w_i, w_j are the widths of agents i, j.

G. Miss Rate

For both Argoverse 2 and INTERACTION, we use the definition of a miss used in the INTERACTION dataset: a prediction is considered a "miss" if the longitudinal or latitudinal distance between the prediction and ground-truth endpoint is larger than their corresponding thresholds, where the latitudinal threshold is $\epsilon_{\text{lat}} := 1\,\text{m}$ and the longitudinal threshold is:

$$\epsilon_{\text{long}} := \begin{cases} 1, & \text{if } v \leq 1.4\,\text{m/s} \\ 1 + \frac{v-1.4}{11-1.4}, & \text{if } 1.4\,\text{m/s} \leq v \leq 11\,\text{m/s} \\ 2, & \text{otherwise}, \end{cases} \tag{18}$$

where v is the ground-truth velocity at the final timestep. We note that Argoverse 2 officially defines a miss as a prediction whose endpoint is more than 2 m from the ground-truth endpoint; however, we report all miss rate numbers in Tab. 2 using the miss rate definition in INTERACTION as it is a more robust measure of miss rate that takes into account the agent's velocity. For completeness, we report miss rate numbers for Argoverse 2 using the Argoverse 2 definition of a miss in Tab. 5.

H. Constant Velocity Model

In Sec. 4, we identify the kinematically complex interactive agents in the datasets by filtering for agents that attain at least d m in FDE with a constant velocity model. An interactive agent is defined as an agent with at least one incident edge in the ground-truth interaction graph, where $\epsilon_I = 2.5\,\text{s}$, as is explained in Sec. 4. In this section, we describe the constant velocity model in more detail. The constant velocity model computes the average velocity over the observed timesteps and unrolls a
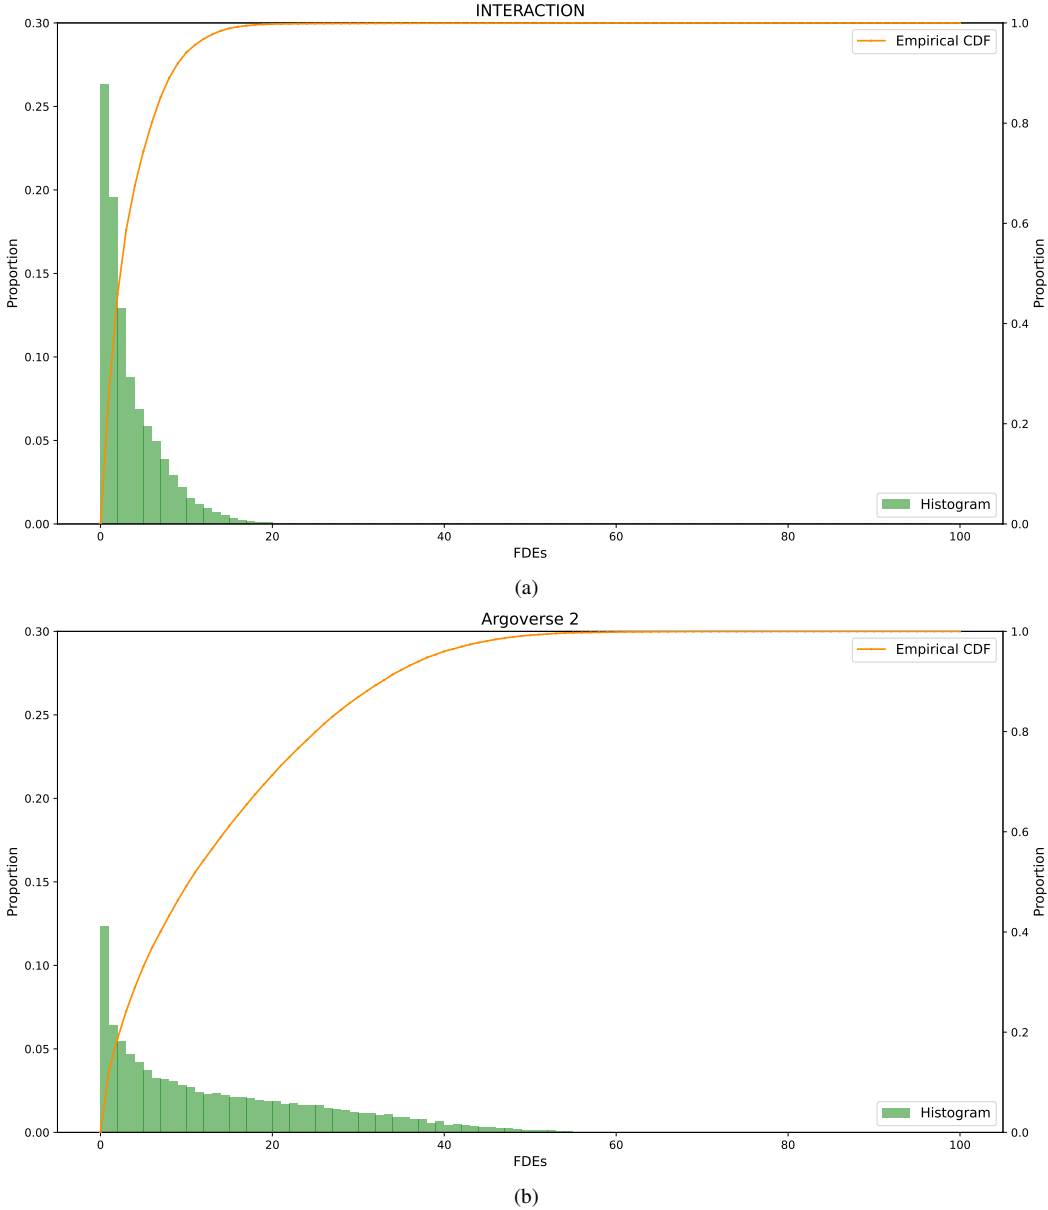
Figure 5. Histogram of FDEs on interacting agents in (a) the INTERACTION dataset, and (b) the Argoverse 2 dataset. The left y-axis corresponds to the histogram and the right y-axis corresponds to the empirical cumulative distribution function (CDF).

future trajectory using the calculated constant velocity. Namely, the average velocity is calculated as:

$$\mathbf{v}_{\text{avg}} = \frac{1}{T_{\text{obs}}} \sum_{t \in [T_{\text{obs}}]} \mathbf{v}_t, \tag{19}$$

where $\mathbf{v}_t$ is the ground-truth velocity at timestep t. Using the constant velocity model, we calculate the agent-level FDE of all interactive agents in the INTERACTION and Argoverse 2 validation sets, respectively, where the FDE distributions are plotted in Fig. 5. We observe that a large proportion of the interactive agents have low FDE with a constant velocity model, especially in the INTERACTION dataset. By filtering out these kinematically simple agents, as is done in Sec. 4, we can assess the model's joint prediction performance on agents that are both interactive and kinematically complex. In Tab. 6, we report the number of interactive agents in the INTERACTION and Argoverse 2 validation sets that attain at least d m in FDE, for $d = 0, 3, 5$. We note that $d = 0$ corresponds to the number of interactive agents in the respective validation sets.

15

Dataset	d	Count
INTERACTION	0	50821
(112994)	3	20977
	5	13011
Argoverse 2	0	39055
(248719)	3	29591
	5	26121

Table 6. Number of *interactive* agents in the INTERACTION and Argoverse 2 datasets that attain at least d m in FDE with a constant velocity model. In parentheses, we include the total number of evaluated agents (interactive + non-interactive) in the respective validation sets.
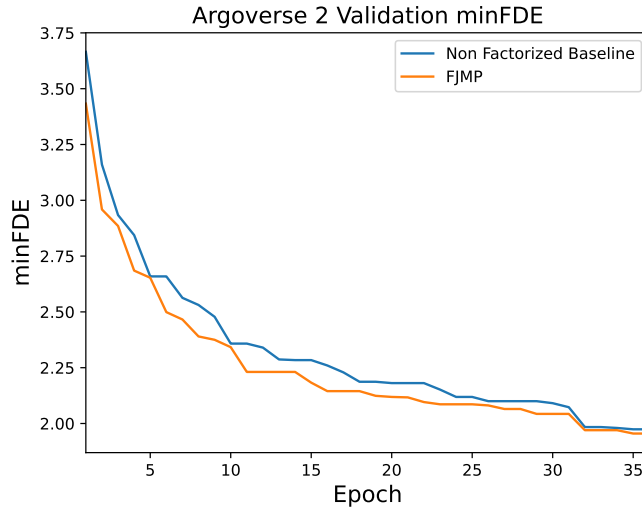

Figure 6. For every epoch e, we plot the best validation minFDE up to epoch e of the Non-Factorized Baseline model and the FJMP model on the Argoverse 2 validation set.

I. Inductive Bias of Joint Factorization

The proposed joint factorization explained in Sec. 3.1.1 can be interpreted as an inductive bias that encourages reasoning about the predicted future of the agents that influence agent n when predicting the future of agent n. In Sec. 3.1.1, we hypothesize that this inductive bias eases the complexity of learning the joint distribution. To support this hypothesis, in Fig. 6, we plot the best validation minFDE of both the non-factorized baseline model and FJMP after every epoch of training. Fig. 6 demonstrates that FJMP consistently attains a lower validation minFDE throughout training, with larger relative gains over the baseline earlier in training.

J. FJMP vs. M2I Interaction Graphs

Figure 7 illustrates the ground-truth interaction graph of a congested scene according to the FJMP and M2I heuristics, respectively. We observe that the M2I heuristic adds several superfluous edges, which would lead to unnecessary additional computation for the factorized decoder.

K. Interaction Graph Predictor Performance

Table 7 reports the proportion of `no-interaction`, `m-influences-n`, and `n-influences-m` edges in the IN-TERACTION and Argoverse 2 training sets. Due to the severe class imbalance, we employ a focal loss when training the interaction graph predictor, as explained in Sec. 3.3.1. The edge type accuracies of the proposed interaction graph predictor

(a) Interaction graph generated with FJMP labeling heuristic.

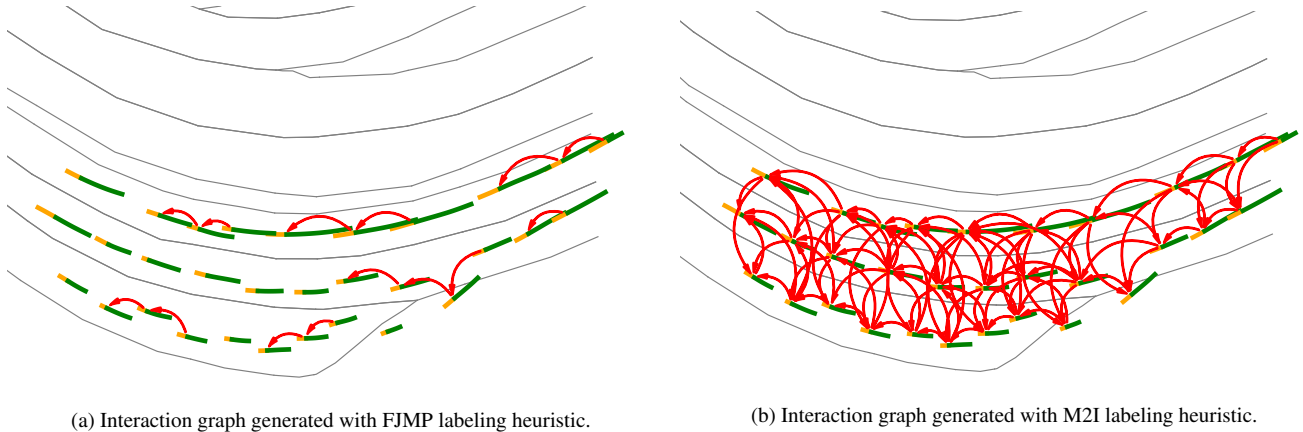(b) Interaction graph generated with M2I labeling heuristic.

Figure 7. Comparison of FJMP and M2I labeling heuristics on a congested scene from the INTERACTION dataset. The ground-truth pasts are indicated in yellow and the ground-truth futures are indicated in green. Lane boundaries are depicted as grey lines. Each red arrow points from an influencer agent to its corresponding reactor agent. We note that two agents at the bottom-right of the scene are on the shoulder of the lane.

Dataset	Edge Type	Edge Type Proportion
INTERACTION	`no-interaction`	0.952
	`m-influences-n`	0.040
	`n-influences-m`	0.008
Argoverse 2	`no-interaction`	0.971
	`m-influences-n`	0.015
	`n-influences-m`	0.013

Table 7. Edge type proportions in the INTERACTION and Argoverse 2 training set interaction graphs with the FJMP labeling heuristic.

Dataset	Edge Type	Edge Type Accuracy
INTERACTION	`no-interaction`	0.994
	`m-influences-n`	0.932
	`n-influences-m`	0.937
Argoverse 2	`no-interaction`	0.991
	`m-influences-n`	0.863
	`n-influences-m`	0.858

Table 8. Accuracy of each edge type on the INTERACTION and Argoverse 2 validation sets with the FJMP interaction graph predictor.

on the INTERACTION and Argoverse 2 validation sets are reported in Tab. 8.

K.1. Ground-truth Interaction Graph Performance

Table 9 compares the performance of FJMP with two modified versions of FJMP: (1) we replace the predicted interaction graphs at inference time with the ground-truth interaction graphs; and (2) we replace the predicted interaction graphs during training and inference time with the ground-truth interaction graphs. The results in Tab. 9 indicate that the choice of interaction graph has a considerable effect on the performance of the factorized joint predictor, as indicated by a 10 cm improvement in iminFDE with the ground-truth interaction graph at inference time over the predicted interaction graph. Moreover, when the model is trained and evaluated with the ground-truth interaction graphs, we see a substantial increase in performance over FJMP with the learned interaction graphs. This indicates that further refinement of the interaction graph predictor may yield additional performance improvements with our FJMP design, which we leave to future work.

Model	Train IG	Inference IG	minFDE	minADE	iminFDE	iminADE
FJMP	Learned	Learned	1.991	0.823	3.164	1.274
FJMP	Learned	Ground-truth	1.975	0.819	3.064	1.229
FJMP	Ground-truth	Ground-truth	**1.920**	**0.802**	**2.952**	**1.215**

Table 9. FJMP with ground-truth vs learned interaction graphs at training and inference time on the Argoverse 2 validation set, All setting. For each metric, the best model is **bolded**. **Train IG** indicates the interaction graphs that are used during training, where **Learned** denotes the predicted interaction graphs from the interaction graph predictor and **Ground-truth** denotes the interaction graphs obtained from the labeling heuristic. The **Inference IG** column is interpreted similarly.

L. Qualitative Results

L.1. Argoverse 2

In this section, we show qualitative results on scenes in the Argoverse 2 validation set where we show side-by-side comparisons between FJMP and the Non-Factorized Baseline. In Fig. 8, Fig. 9, and Fig. 10, for each row, the left panel shows the non-factorized baseline predictions, the middle panel shows FJMP predictions, and the right panel shows the predicted DAG. We visualize only the best scene-level modality to avoid clutter. In Fig. 8, we show examples where FJMP reasons properly about future interactions between vehicles and crossing pedestrians, leading to more scene-consistent predictions than the non-factorized baseline. In Fig. 9, we show examples where FJMP reasons properly in scenes with interactive left-turns. In contrast, the non-factorized baseline incorrectly predicts conservative behaviour where the left-turning vehicle avoids the passing vehicle's trajectory. In Fig. 10, we show qualitative examples where FJMP correctly identifies chains of leader-follower interactions, which in turn leads to more accurate leader-follower predictions than the non-factorized baseline. In Fig. 12, we illustrate two failure cases of the FJMP model. In the first case, an erroneous interaction graph prediction biases the model to predict incorrect leader-follower behaviour. In the second case, an erroneous influencer future prediction negatively biases the downstream reactor prediction.

L.2. INTERACTION

Figure 11 shows qualitative results of FJMP on various scenes in the INTERACTION dataset, with all $K = 6$ scene-level modalities visualized. We emphasize FJMP's ability to produce accurate and scene-consistent predictions for scenes with a large number of interacting agents.
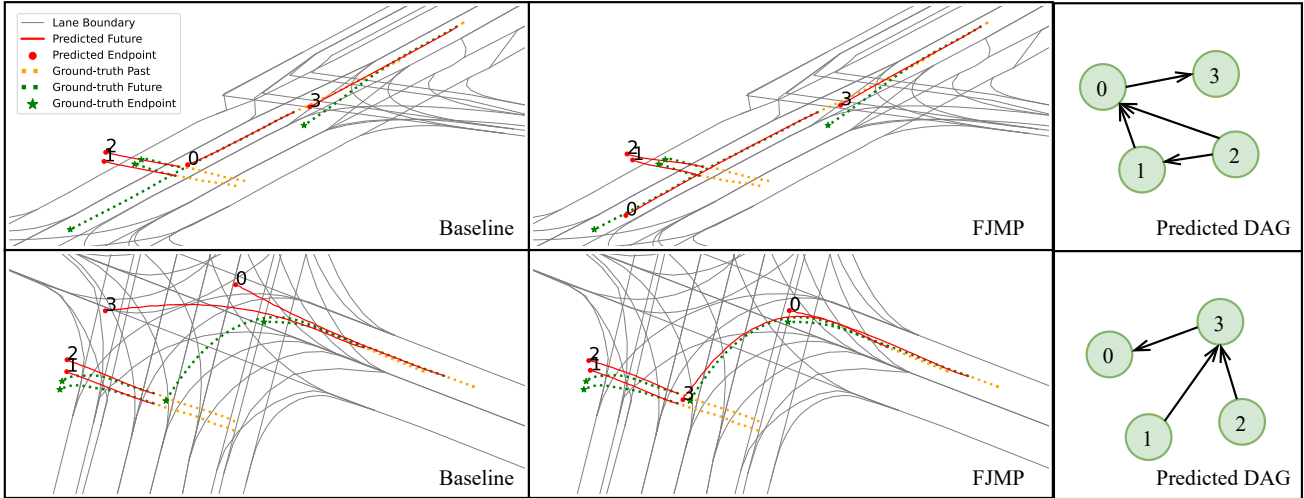
Figure 8. Qualitative examples of vehicles interacting with crossing pedestrians in the Argoverse 2 validation set. All predicted DAGs match the ground-truth DAG. In both scenes, FJMP correctly identifies the crossing pedestrians (1 and 2 in both scenes) as influencers to the interacting vehicle (0 in the top row and 3 in the bottom row), which leads to more scene-consistent future predictions. Additionally, in the bottom row, more scene-consistent leader-follower behaviour between 0 and 3 is captured by FJMP.
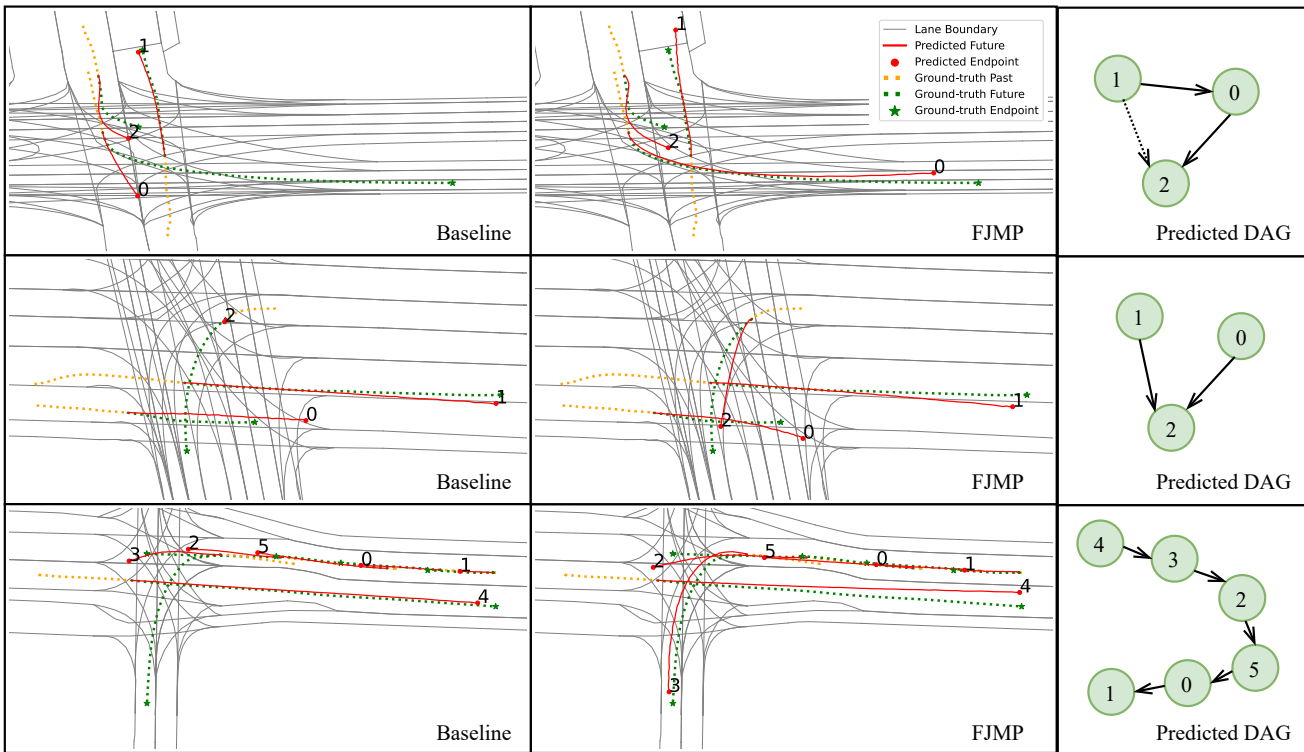

Figure 9. Qualitative examples of left-turn interactive scenes in the Argoverse 2 validation set. Solid arrows denote predicted edges in the DAG that match the ground-truth DAG and dashed arrows denote edges in the ground-truth DAG that are not in the predicted DAG. In all scenes, FJMP correctly identifies the passing vehicle as the influencer and the left-turning vehicle as the reactor. The Non-Factorized baseline consistently predicts overly conservative behaviour that avoids the influencer trajectory. In contrast, FJMP consistently captures the proper left-turn behaviour.
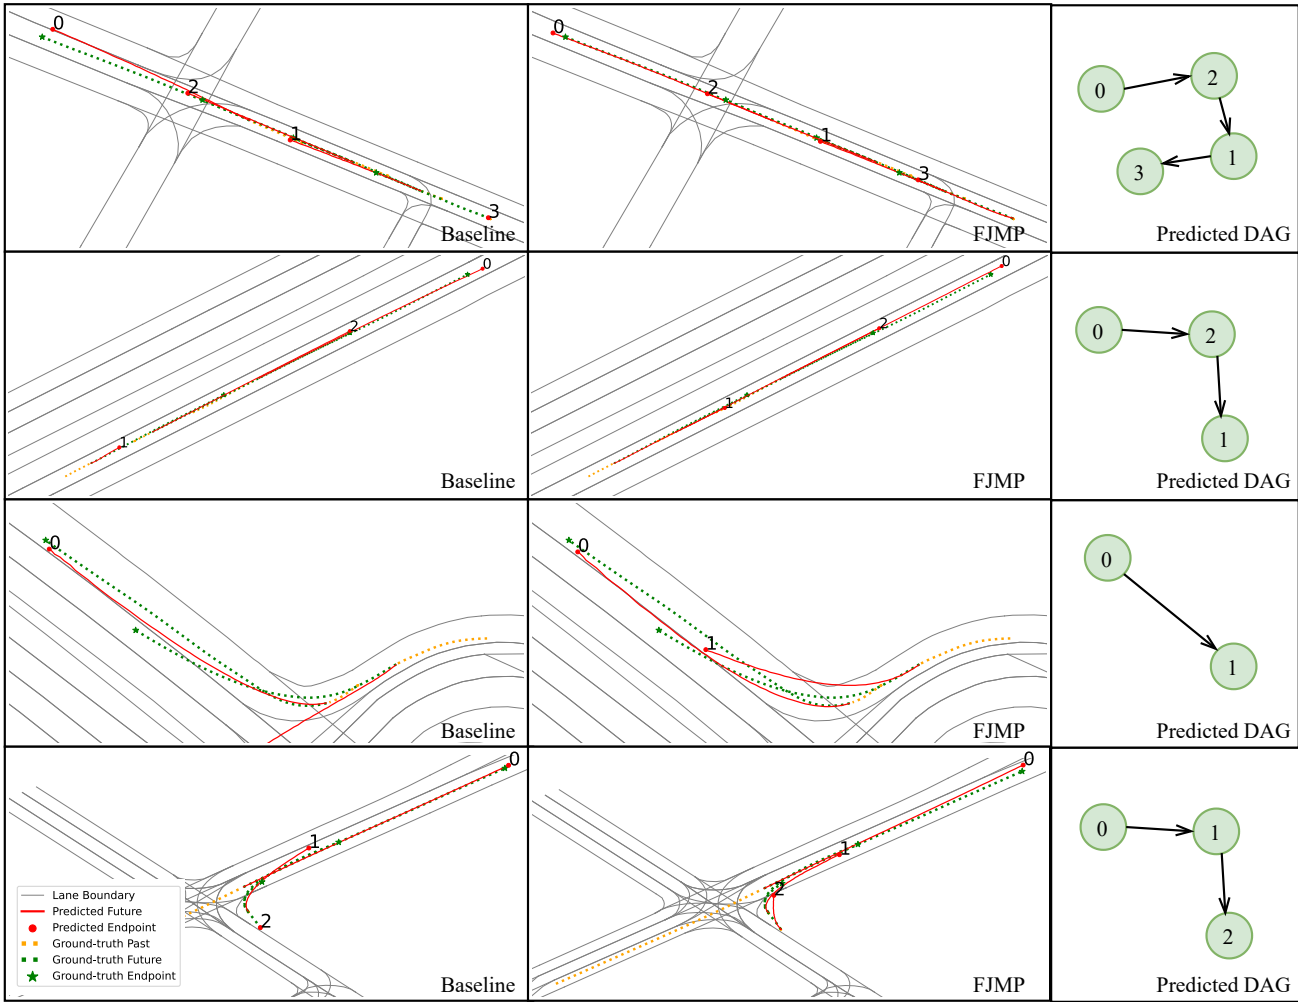
19

Figure 10. Qualitative examples of leader-follower interactive scenes in the Argoverse 2 validation set. All predicted DAGs match the ground-truth DAG. In all of the above scenes, FJMP correctly predicts chains of influencer-reactor relationships. The non-factorized baseline often produces conservative behaviour for the trailing vehicle (see the first, second, and fourth rows). In contrast, FJMP predicts more proper leader-follower behaviour for the trailing vehicles (leaf nodes in the DAG). In the third row, the right-turn mode of the vehicle 1 is missed by the non-factorized baseline, whereas FJMP correctly identifies the right-turn mode due to correctly identifying the leader-follower interaction between vehicles 0 and 1.
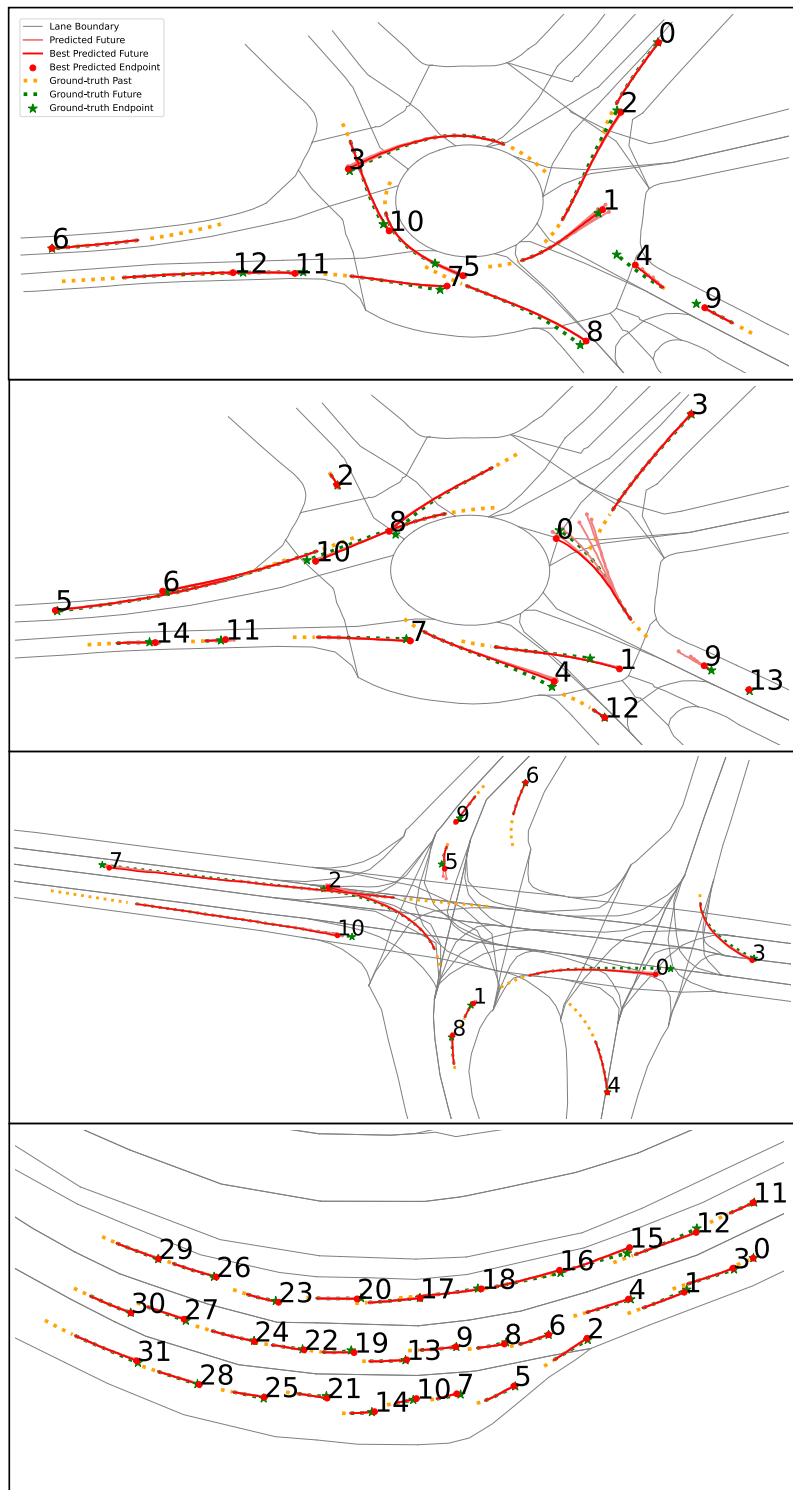
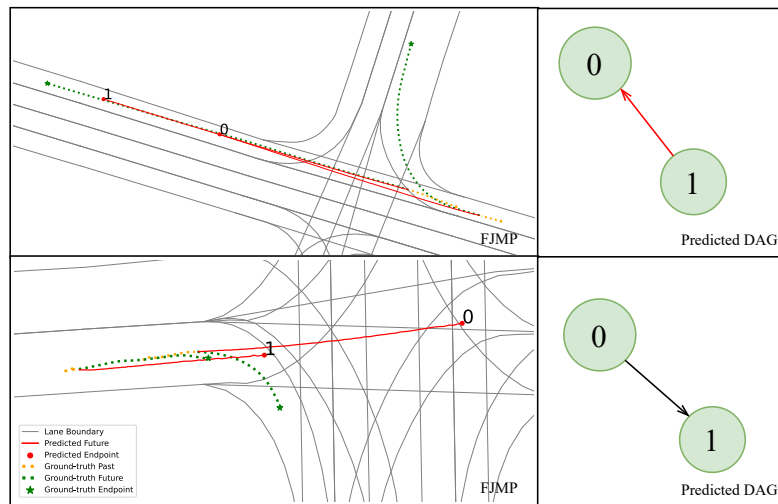Figure 11. Qualitative examples of FJMP on agent-dense scenes in the INTERACTION dataset.

Figure 12. Qualitative examples of failure cases of the FJMP model. The predicted DAGs are shown on the right, where true positive edges are indicated in black and false positive edges are indicated in red. In the top row, the model erroneously predicts an edge from 1 to 0, which biases the model prediction towards leader-follower behaviour between 1 and 0. This causes the model to miss the right-turn mode for agent 0. In the bottom row, the interaction graph is correctly predicted; however, the influencer trajectory is erroneously predicted as continuing straight, which causes the reactor's prediction to follow 0 and also continue straight. Consequently, the prediction for both agents misses the right-turn-mode.