

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
<https://rluke22.github.io/FJMP>

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 INTERACTION 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, 46], 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

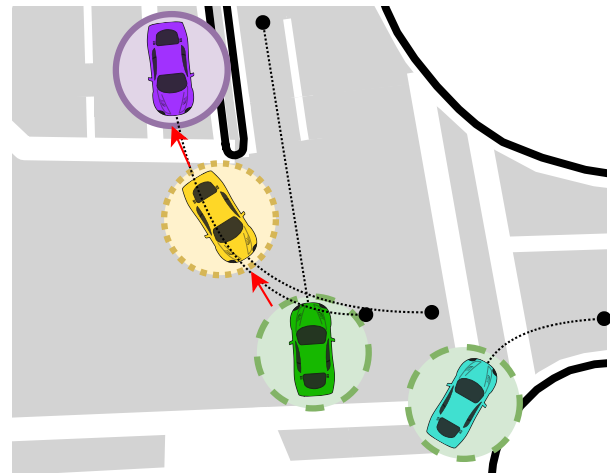


Figure 1. An illustration of the directed acyclic interaction graph, comprised of colored nodes and red arrows. The dotted black lines denote the ground-truth futures over a short time horizon. FJMP first produces marginal predictions for the green (dashed) nodes, followed by a conditional prediction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

likely. As scene-consistent future predictions are critical for downstream planning, recent work has shifted toward generating 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 pre-

*Corresponding author.

dictions. 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 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 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 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, 44, 49] 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 [46] 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, 50]. 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, 50] or graph neural networks (GNNs) [4, 5, 14, 21, 25, 29, 31, 48] 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, 43, 46–48, 50]; however, marginal

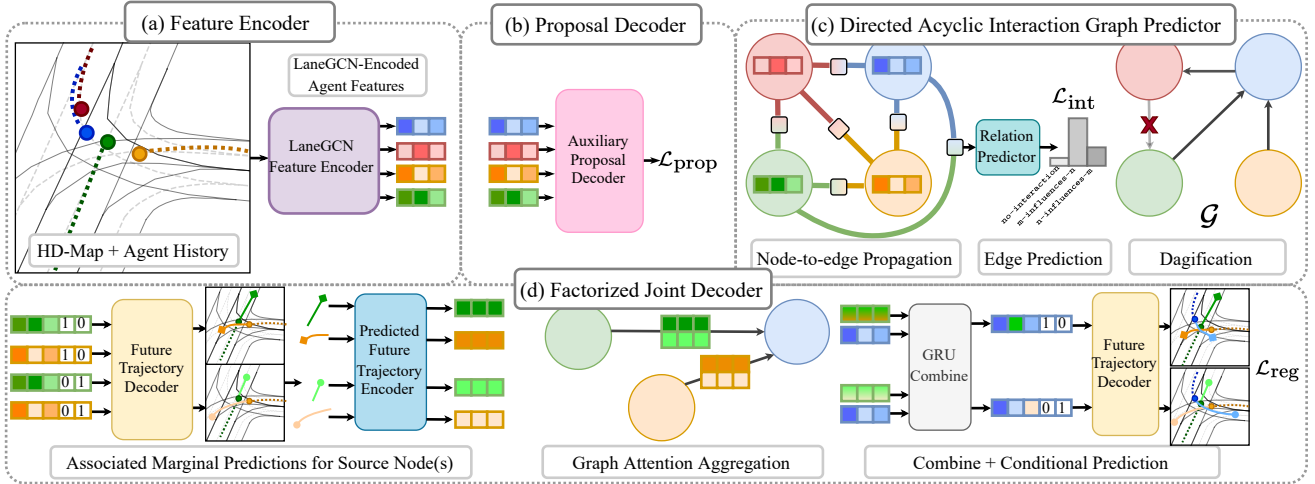


Figure 2. Illustration of the proposed FJMP framework. (a) Agent histories and the HD-Map are first processed by a LaneGCN-inspired feature encoder. (b) During training, the LaneGCN-encoded features are fed into an auxiliary future proposal decoder trained with a regression loss to encourage the LaneGCN 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, 45], 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:

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

where $\text{pa}_{\mathcal{G}}(n)$ 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 , $t \in [T_{\text{obs}}]$, 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, $\{\mathbf{A}_i\}_{i \in \{\text{pre}, \text{suc}, \text{left}, \text{right}\}}$, $\mathbf{A}_i \in \mathbb{R}^{M \times M}$, 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 $\{\mathbf{A}_i\}_{i \in \{\text{pre}, \text{suc}, \text{left}, \text{right}\}}$, and the preprocessed agent history states $[\hat{\mathbf{p}}_t^n, \mathbf{v}_t^n, \psi_t^n]$ for all $n \in [N]$, $t \in [T_{\text{obs}}]$.

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 $\{\hat{\mathbf{y}}_{\text{prop}, k}^n\}_{k \in [K]} = f_{\text{prop}}(\mathbf{h}_n)$ and it is supervised by a joint regression loss $\mathcal{L}_{\text{prop}}$. 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 collision

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 H. 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:

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

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 $e_{m,n}$:

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

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:

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

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:

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

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

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

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:

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

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 | 0.179 | 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 |
| HGT-Joint | - | 0.307 | 1.056 | 0.186 | 0.016 | 0.190 |
| Traj-MAE | - | 0.307 | 0.966 | <u>0.183</u> | 0.021 | <u>0.188</u> |
| FJMP (Ours) | - | 0.275 | 0.922 | 0.185 | <u>0.005</u> | 0.187 |

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 underlined. Lower is better for all metrics.

parameterize AGG using graph attention [42]:

$$\mathbf{m}_n := \text{AGG}(\{\mathbf{b}_{mn} | m \in \text{pa}_{\mathcal{G}}(n)\}) = \sum_{m \in \text{pa}_{\mathcal{G}}(n)} \alpha_{mn} \mathbf{W}_1 \mathbf{b}_{mn}, \quad (8)$$

$$\alpha_{mn} = \frac{\exp(\text{LeakyReLU}(\mathbf{a}^\top [\mathbf{W}_1 \mathbf{b}_{mn} || \mathbf{W}_2 \mathbf{h}_n]))}{\sum_{k \in \text{pa}_{\mathcal{G}}(n)} \exp(\text{LeakyReLU}(\mathbf{a}^\top [\mathbf{W}_1 \mathbf{b}_{kn} || \mathbf{W}_2 \mathbf{h}_n]))}. \quad (9)$$

COMB is parameterized by a GRU recurrent module:

$$\mathbf{h}_n \leftarrow \text{COMB}(\mathbf{m}_n, \mathbf{h}_n) = \text{GRU}(\mathbf{m}_n, \mathbf{h}_n). \quad (10)$$

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}_{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:

$$\mathcal{L}_1 = \mathcal{L}_{\text{int}} + \mathcal{L}_{\text{prop}}. \quad (11)$$

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:

$$\mathcal{L}_2 = \mathcal{L}_{\text{reg}} + \mathcal{L}_{\text{prop}}. \quad (12)$$

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 F.

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

| Dataset | Actors Evaluated | Model | minFDE | minADE | SCR | SMR | iminFDE | iminADE | iminFDE ₃ | iminADE ₃ | iminFDE ₅ | iminADE ₅ |
|-------------|------------------|----------------|--------|--------|-------|-------|---------|---------|----------------------|----------------------|----------------------|----------------------|
| Interaction | - | Non-Factorized | 0.643 | 0.199 | 0.004 | 0.088 | 0.688 | 0.210 | 0.784 | 0.240 | 0.854 | 0.261 |
| | | FJMP | 0.630 | 0.194 | 0.003 | 0.084 | 0.672 | 0.206 | 0.758 | 0.232 | 0.826 | 0.252 |
| | | Δ | 0.013 | 0.005 | 0.001 | 0.004 | 0.016 | 0.004 | 0.026 | 0.008 | 0.028 | 0.009 |
| Argoverse 2 | Scored | Non-Factorized | 1.965 | 0.834 | - | 0.349 | 2.957 | 1.223 | 3.276 | 1.340 | 3.436 | 1.399 |
| | | FJMP | 1.921 | 0.819 | - | 0.343 | 2.893 | 1.204 | 3.205 | 1.320 | 3.356 | 1.377 |
| | | Δ | 0.044 | 0.015 | - | 0.006 | 0.064 | 0.019 | 0.071 | 0.020 | 0.080 | 0.022 |
| | All | Non-Factorized | 1.995 | 0.825 | - | 0.340 | 3.302 | 1.309 | 3.759 | 1.477 | 3.952 | 1.545 |
| | | FJMP | 1.963 | 0.812 | - | 0.337 | 3.204 | 1.273 | 3.652 | 1.439 | 3.839 | 1.504 |
| | | Δ | 0.032 | 0.013 | - | 0.003 | 0.098 | 0.036 | 0.107 | 0.038 | 0.113 | 0.041 |

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.

the K joint predictions; **SMR** is the minimum proportion of agents whose predicted trajectories “miss” the ground-truth from the K joint predictions, where a miss is defined in Appendix I; and **SCR** is the proportion of modalities where two or more agents collide. The INTERACTION test set additionally reports two joint prediction metrics: **Cross-Col** 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. For all metrics, we evaluate $K = 6$. 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 J 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. As bounding-box information is not provided with Argoverse 2, the collision checker used to construct interaction labels uses a pre-

defined length/width for each agent type, as listed in Appendix F.2. Our INTERACTION and Argoverse 2 models train in 10 and 15 hours. See Appendix G for more details.

| Model | minFDE | minADE | SMR | Prop. Edges | Inf. Time (s) |
|----------------|--------------|--------------|--------------|-------------|---------------|
| Non-Factorized | 0.643 | 0.199 | 0.088 | - | 0.010 |
| FJMP (Dense) | 0.623 | 0.193 | 0.081 | 0.180 | 0.062 |
| FJMP | <u>0.626</u> | 0.193 | <u>0.083</u> | 0.045 | <u>0.038</u> |

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. FJMP performs the best on minFDE, minADE, and the official ranking metric CMR, while performing competitively on other metrics. Crucially, FJMP produces joint predictions that are both more accurate—as demonstrated by its superior performance on minADE and minFDE—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.

Table 2 reports validation results on the INTERACTION and Argoverse 2 datasets, where we compare FJMP against

| Model | Prop? | TF? | minFDE | minADE | iminFDE | iminADE |
|----------------|-------|-----|--------------|--------------|--------------|--------------|
| Non-Factorized | ✗ | ✗ | 1.995 | 0.825 | 3.302 | 1.309 |
| FJMP | ✗ | ✗ | 2.004 | 0.829 | 3.274 | 1.304 |
| FJMP | ✗ | ✓ | 2.001 | 0.827 | 3.300 | 1.312 |
| FJMP | ✓ | ✗ | <u>1.987</u> | <u>0.820</u> | <u>3.255</u> | <u>1.293</u> |
| FJMP | ✓ | ✓ | 1.963 | 0.812 | 3.204 | 1.273 |

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

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 M 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 K. 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 nearly 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 denser interaction graph; however, there may exist better heuristics for classifying 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.

6. Acknowledgements

This work was funded by Ontario Graduate Scholarship and NSERC. We thank Benjamin Thérien and Prarthana Bhattacharyya for their valuable insights and discussions.

References

- [1] Yutong Ban, Xiao Li, Guy Rosman, Igor Gilitschenski, Ozan 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. In *Proceedings of the Conference on Robot Learning (CoRL)*, 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. Spaggn: 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 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. *arXiv preprint arXiv: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 International Conference on Computer Vision (ICCV)*, 2021. 2
- [14] Jiyang Gao, Chen Sun, Hang Zhao, Yi Shen, Dragomir Anguelov, Congcong Li, and Cordelia Schmid. Vectormet: 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 Stanculescu, 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 Stanculescu, 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 Stanculescu, 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. Densentnt: 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. *arXiv preprint arXiv: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 multimodal 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. *arXiv preprint arXiv: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. *arXiv preprint arXiv: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, 13
- [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, Jinghui 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. In *Advances in Neural Information Processing Systems (NeurIPS)*, 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] Petar Velickovic, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Liò, and Yoshua Bengio. Graph attention networks. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2018. 6
- [43] 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
- [44] 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
- [45] David Wu and Yunnan Wu. Air² for interaction prediction. *arXiv preprint arXiv:2111.08184*, 2021. 3
- [46] 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
- [47] Maosheng Ye, Jiamiao Xu, Xunnong Xu, Tongyi Cao, and Qifeng Chen. DCMS: motion forecasting with dual consistency and multi-pseudo-target supervision. *arXiv preprint arXiv:2204.05859*, 2022. 2

- [48] 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
- [49] Wei Zhan, Litng Sun, Di Wang, Haojie Shi, Aubrey Clause, Maximilian Naumann, Julius Kümmerle, Hendrik Königshof, Christoph Stiller, Arnaud de La Fortelle, and Masayoshi Tomizuka. INTERACTION dataset: An international, adversarial and cooperative motion dataset in interactive driving scenarios with semantic maps. *arXiv preprint arXiv:1910.03088*, 2019. 2
- [50] 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