

Disentangled Neural Relational Inference for Interpretable Motion Prediction

Victoria M. Dax¹, Jiachen Li¹, Enna Sachdeva², Nakul Agarwal², and Mykel J. Kochenderfer¹

Abstract—Effective interaction modeling and behavior prediction of dynamic agents play a significant role in interactive motion planning for autonomous robots. Although existing methods have improved prediction accuracy, few research efforts have been devoted to enhancing prediction model interpretability and out-of-distribution (OOD) generalizability. This work addresses these two challenging aspects by designing a variational auto-encoder framework that integrates graph-based representations and time-sequence models to efficiently capture spatio-temporal relations between interactive agents and predict their dynamics. Our model infers dynamic interaction graphs in a latent space augmented with interpretable edge features that characterize the interactions. Moreover, we aim to enhance model interpretability and performance in OOD scenarios by disentangling the latent space of edge features, thereby strengthening model versatility and robustness. We validate our approach through extensive experiments on both simulated and real-world datasets. The results show superior performance compared to existing methods in modeling spatio-temporal relations, motion prediction, and identifying time-invariant latent features.

I. INTRODUCTION

Understanding and modeling complex interactions among dynamic agents is important to various applications and tasks, including robotics [1], traffic modeling and management [2], and social network analysis [3]. In the field of robotics, one essential downstream task is multi-agent trajectory prediction, which serves as a prerequisite for safe and high-quality decision-making and motion planning in complex and crowded scenarios. Modeling inter-agent interactions is crucial to understanding the joint dynamic behavior of the agents. For instance, the joint prediction of two vehicles approaching an intersection requires modeling and reasoning about their potential interactions, such as yielding, overtaking, or not interacting at all.

Recent approaches [4]–[7] have focused on modeling interactions among agents by inferring latent interaction graphs, where edges represent different types of interactions. However, these methods are limited to inferring categorical relations and are not equipped to capture more nuanced characteristics. For example, while these models can identify whether or not a pair of particles is connected by a spring, they are unable to infer its elastic coefficient. Furthermore, these methods focus on minimizing distance-based prediction errors in test cases that align with the observed data distribution. By primarily evaluating their methods using in-distribution

¹V. M. Dax, J. Li, and M. J. Kochenderfer are with the Stanford Intelligent Systems Laboratory (SISL), Stanford University, USA. {vmdax, jiachen_li, mykel}@stanford.edu.

²E. Sachdeva and N. Agarwal are with the Honda Research Institute, USA. {enna_sachdeva, nakul_agarwal}@honda-ri.com




Fig. 1: The encoder evaluates edge features, a section of which is used to increase interpretability through disentanglement, such as restricted labeling or pair matching.

samples, they overlook out-of-distribution situations, where models may encounter unseen and challenging interaction patterns between agents or diverse environmental contexts. This limits both the interpretability and generalizability of the existing works, which undermines the reliability of prediction models and proves inadequate for safety-critical applications such as autonomous driving.

In this work, we propose a model (dG-VAE) built upon a variational auto-encoder framework for discovering interpretable interactions from observations. Our model achieves interpretability by incorporating additional edge features that capture latent characteristics, such as the elastic coefficient of a spring. It combines interpretable graph structures and edge feature learning with multi-relational decoders within an unsupervised training framework. To enhance the interpretability of the learned edge features, we employ two types of disentanglement techniques, either supervised or unsupervised, depending on the application. By using an inference-based approach that disentangles time-invariant features, our model demonstrates interpretability and improved generalizability in out-of-distribution settings. Integrating interpretable components and inference-based learning allows our model to capture and understand more complex interactions, leading to more reliable and robust predictions.

The prospective applications of dG-VAE span various fields, opening up new avenues for exploration. In robotics, it has the potential to enhance safety by effectively forecasting the motion of other participants. In healthcare, it holds promise in refining gait analysis and contributing to the development of smart prosthetics and exoskeletons that sync with user intentions. In security, transportation, and traffic management, it could be used for identifying anomalies, strategizing crowd management, and optimizing traffic flows.

The main contributions of this paper are as follows:

- We propose a VAE-based model architecture that goes beyond inferring dynamic interaction graphs by incorporating edge features to characterize the interactions.

- We leverage disentanglement techniques to partition the embedding space into time-invariant and temporal components, boosting prediction accuracy, enhancing out-of-distribution generalization, and occasionally yielding human-interpretable embeddings.
- We extensively evaluate our method on multiple benchmark datasets to show its effectiveness. Our approach outperforms baseline techniques in modeling spatio-temporal relations and accurately predicting interactions, trajectories, and underlying behavioral determinants.

II. RELATED WORK

Trajectory prediction is important in fields like robotics, autonomous vehicles, and human interaction. Social LSTM [8] uses long short-term memory networks to model pedestrian interactions and predict their trajectories in crowded spaces. Building upon this framework, Trajectron [9] uses dynamic spatio-temporal graphs and graph-structured LSTM networks to model and predict the trajectories of multiple agents in a scene. Generative adversarial networks have been used to predict trajectories compliant with social and physical constraints [10], [11]. Recent works proposed more advanced graph- or transformer-based approaches to model the spatio-temporal relations between interactive agents, thereby enhancing prediction performance [6], [12]–[16].

Spatio-temporal graph modeling has been widely studied to capture the spatio-temporal relations between dynamic interacting agents. Social-STGCNN [12] introduced a spatio-temporal graph convolutional neural network for human trajectory prediction. EvolveGraph [6] and dNRI [5] are dynamic relational reasoning methods to identify underlying relations based on the information encoded in given sets of trajectories. A spatio-temporal attention mechanism [17] and a spatio-temporal graph transformer network [13] were proposed to model the inter-graph temporal dependencies.

The significance of *interpretability* is increasingly recognized in various fields such as autonomous vehicles [18] or security systems [19] for example. Specifically, in motion prediction, this emphasis on interpretability has given rise to methods like GRIT [20], which use decision trees to better understand learned spaces. Kothari, Stringer, and Alahi [21] combine rule-based and neural network models to predict high-level intents and scene-specific residuals. For multi-agent interaction modeling, studies such as Grounded Relational Inference [22], which learns reward functions, and GRIN [23], disentangling interactions from agent intentions, are relatively recent. Collaborative Prediction Units [24], learning a weighted aggregate of individual prediction units, consider inter-agent influences and adjust in real-time to incoming data. Other works [25], [26] annotate ground truth or pseudo-labels to interactions for better understanding. Despite these advancements, most methods view interactions as static, categorical, or both. In contrast, our research uncovers disentangled and interpretable sequences of dynamic interactions from multi-agent observations.

Our work lies at the intersection of these three fields. We build on DESIRE [27], which leverages a conditional

VAE to forecast future trajectories amid multiple interacting agents, an approach later enhanced by NRI [4] and dNRI [5] through the integration of Graph Neural Networks (GNNs) to deduce a latent graph. These works rely on the proficiency of VAEs to comprehend intricate data distributions and produce new samples resembling original data. Our work builds on these models, enhancing the predicted latent graph with edge features, thereby boosting expressivity, and introducing “disentanglement” of the embedding space to promote the discovery of interpretable factorized latent representations. Disentanglement, initially an adjustable hyperparameter in the loss function, has shown great success in computer vision [28] and reinforcement learning [29].

III. PRELIMINARIES

Variational auto-encoders (VAEs) are a type of deep generative model that blends the concepts from deep learning and probabilistic graphical models. VAEs consist of an encoder and a decoder. The encoder maps the input data x to a lower-dimensional latent variable z , represented by a probability distribution $q_\theta(z | x)$, typically a Gaussian distribution with a learnable mean and covariance. It learns to approximate the true posterior distribution of the latent variable given the input data. The decoder maps the latent variable back to the original data space, i.e., it learns the conditional probability distribution $p_\phi(x | z)$ of the input data given the latent variable.

During training, we optimize the *reconstruction* objective, which ensures that the generated samples resemble the original data, and the *regularization* objective, which encourages the learned latent space to have a specific structure $p(z)$ (e.g., a standard Gaussian distribution), simultaneously. The resulting loss function is written as

$$l_i(\theta, \phi) = - \mathbb{E}_{z \sim q_\theta(z|x)} [\log p_\phi(x | z)] + \text{KL}(q_\theta(z | x) \| p(z)), \quad (1)$$

a combination of the reconstruction error and the Kullback-Leibler (KL) divergence between the approximate and true posterior distributions, which is also known as the evidence lower bound (ELBO).

Graph neural networks (GNNs) are a class of deep learning models designed to handle data represented as graphs. They are particularly well-suited for problems where the data has a complex, irregular structure. GNNs process graph-structured data by iteratively passing and aggregating messages between neighboring nodes.

Let V denote the set of nodes in the graph. During each message-passing iteration, a hidden embedding \mathbf{x}_u corresponding to node $u \in V$ is updated according to information aggregated from u 's neighborhood nodes $\mathcal{N}(u)$. The message-passing update can be expressed as

$$\mathbf{x}_u^{(k+1)} = \text{UP}^{(k)} \left(\mathbf{x}_u^{(k)}, \text{AGG}^{(k)} \left(\{ \mathbf{x}_v^{(k)} \mid v \in \mathcal{N}(u) \} \right) \right), \quad (2)$$

where UP, short for UPDATE, and AGG denote arbitrary differentiable functions (e.g., MLP), and k denotes the index of message-passing layers. This mechanism can be applied

iteratively to capture information from a broader neighborhood, and the final node representations or aggregated graph representations can be used for various tasks, such as node classification, link prediction, and graph classification.

IV. METHOD

Problem Definition. Our goal is to learn accurate trajectory distributions for a variable number of interacting agents $N(t)$ based on historical observations. For each agent i present at time step t , we consider its history $x_{t-H+1:t}^{(i)}$ where H denotes the history horizon. We aim to predict the distribution of all entities' future states $p(y^{(i)} | \{x_{t-H+1:t}^{(j)} | \forall j \in \{1, \dots, N(t)\}\})$ for the upcoming T steps. Here $y^{(i)} = \hat{x}_{t+1:t+T}^{(i)} \in R^{T \times N(t) \times d}$, with d the state dimension, represents the predicted future trajectory of agent i .

Overview. Our proposed model architecture is inspired by NRI [4] and dNRI [5]. While these previous iterations primarily focused on edge classification, our method uniquely emphasizes learning edge features, which can be leveraged further. This modification increases the expressivity of the network and allows for the incorporation of disentanglement in the edge features to specific characteristics of the underlying interactions. We hypothesize that this added regularized expressivity increases the out-of-distribution generalization ability and interpretability of the learned embeddings.

We provide the model with historical input $X \in \mathbb{R}^{N \times H \times d}$. The input is passed through the encoder, which computes edge logits $p(z_t | X_{t-H+1:t})$, i.e., the likelihood of this edge existing in the latent space, alongside an edge feature matrix E_t . We then sample z_t from the posterior distribution and pass z_t and $E_{z,t}$ to the decoder, which predicts the next state $X' \in \mathbb{R}^{N \times d}$. This process is shown in Figure 1.

Encoder. The architecture of the encoder builds upon dNRI's method. We use a GNN block composed of a linear embedding followed by an edge convolutional layer [30]:

$$\mathbf{x}'_{i,t} = \sum_{j \in \mathcal{N}_t(i)} h_{\Theta}(\mathbf{x}_{i,t} \parallel \mathbf{x}_{j,t} - \mathbf{x}_{i,t}), \quad (3)$$

where h_{Θ} denotes a neural network (i.e., an MLP) and \parallel indicates the concatenation of two node embeddings $\mathbf{x}_{i,t}$ and $\mathbf{x}_{j,t}$. We recall from Section III that GNNs aggregate messages from neighboring nodes to update the node embeddings at each iteration. Because we are learning edge features instead of node features, we adapt the last *edge convolutional* layer to edge-level embeddings, which is written as

$$\mathbf{e}'_{ij,t} = h_{\Theta}(\mathbf{x}_{i,t} \parallel \mathbf{x}_{j,t}). \quad (4)$$

The intermediate embedding we obtain from the GNN block is then forwarded to an RNN block composed of a single LSTM layer with ELU activation and a dropout layer. Next, we store the RNN hidden state and give the embeddings to two MLP heads each composed of two linear layers with ELU activations. The first returns the edge posterior logits and the second computes the edge features E_t .

The dNRI method [5] suggested *learning the prior*, an idea we incorporated into our encoder shown in Figure 2. This alteration requires a forward and a backward RNN to




Fig. 2: Encoder architecture that learns the prior.

generate two embeddings. The forward embedding is used to a) generate a prior, and b) in conjunction with the backward embedding, return the posterior $p(z_t | X_{t-H+1:t})$. As a result of using the backward RNN, the posterior is computed using future information. During training, the posterior is used the way it is currently in our model, and the prior is trained to match it. During testing, the latent variable z_t is sampled from the learned prior.

Sampling. The encoder returns a distribution over possible relations. Because traditional categorical sampling is not differentiable, required for backpropagating weight updates, we adopt the concrete distribution [31], a continuous proxy to the discrete categorical distribution. This sampling technique uses reparameterization by, first, sampling a vector \mathbf{g} from $\text{GUMBEL}(0, 1)$ and then calculating

$$z_{ij,t} = \text{SOFTMAX}((\hat{z}_{ij,t} + \mathbf{g})/\tau), \quad (5)$$

where $\hat{z}_{ij,t}^{(ij)}$ are the posterior logits at time t and τ adjusts distribution smoothness. This method approximates discrete sampling in a gradient-friendly manner, allowing the encoder to receive feedback from decoder reconstruction.

Decoder. The decoder takes as inputs graph z_t , node features X_t , and edge features $E_{z,t}$. These are processed by an adapted *edge convolutional* layer:

$$\mathbf{x}'_{i,t} = \sum_{j \in \mathcal{N}_t(i)|z_t} h_{\Theta}(\mathbf{x}_{i,t} \parallel \mathbf{x}_{j,t} \parallel \mathbf{e}_{ij,t}), \quad (6)$$

to obtain node-level embeddings, which are then forwarded to a GRU layer with ELU activation and a linear readout.

Loss function and training strategy. The loss function is composed of a generative loss, which measures the difference between the model's input and output, and a latent loss, which compares the latent vector to a Gaussian distribution with zero mean and unit variance. We use the negative log-likelihood loss:

$$\text{NLL} = \frac{1}{T} \sum_{t \in [T]} \sum_{j \in [d]} \frac{1}{2\sigma} (z_{t,j} - y_{t,j})^2, \quad (7)$$

which is widely used when training probabilistic models. Here, $\sigma = 5.0 \times 10^{-5}$ is a hyperparameter. We used the learned prior when evaluating the KL-divergence (i.e., the latent loss) of $p(z | X)$ and $p(z)$:

$$\text{KL} = \sum_{t \in [T]} \left(- \sum_z p(z_t | X) \log p(z_t | X, z_{1:t-1}) \right). \quad (8)$$

Disentanglement. To encourage interpretable edge feature learning, we introduce two types of disentanglement: *restricted labeling* and *pair matching*, as shown in Figure 3.

(a) Restricted Labeling $p(x, s_I)$ (b) Pair Matching $p(x, x')$

Fig. 3: Variations of disentanglement.

The former, a supervised method from style-content disentanglement [32], matches distributions based on observed x and a subset of ground truth features s_I . As the encoder processes the edge feature matrix E from Figure 2, certain learned variables are directed to align with specified attributes using an additional loss. For example, in the Spring dataset (Section V-B), we apply regression to parts of the predicted feature matrix E guided by the known spring constants during training. Although the actual spring constants are unknown during testing, the model has been conditioned to infer them. By basing the model on a feature with known semantics and predictive implications, we can further enhance the interpretability and performance of our model.

Pair Matching is an unsupervised approach, focusing on paired data (x, x') with a subset of common feature values. It can be deemed a weakly supervised approach as it does not rely on the underlying values of s_I but correlates with its indices. Although various strategies exist for its implementation, this method typically involves sampling twice from the encoder and equating dimensions between E_1 and $E_2 \in \mathbb{R}^{m \times T \times d}$, where m is the number of edges and d is the hidden dimension, by averaging their differences. Given our model’s T sequential samples, we forgo double sampling. We average the partial matrix $E_h \in \mathbb{R}^{m \times T \times d_h}$ across time, resulting in $\hat{E}_h \in \mathbb{R}^{m \times d_h}$, and populate E with these time-invariant features, whereby the model learns time-invariant latent characteristics of each interaction. Our strategy thereby renders dimensions $[d_h] \subset [d]$ virtually static.

V. EXPERIMENTS

Datasets and Baselines. We evaluate our method on four benchmark datasets: NBA, Spring, Motion Capture, and inD datasets through comparisons with the following baselines: a) A two-layer MLP with a hidden dimension of 256 and ReLU activation, b) an LSTM layer, with a hidden dimension of 128 between two fully connected layers with hidden dimensions of 64, that processes all agents simultaneously, c) IMMA [16], a forward prediction model that uses a multiplex latent graph to represent different types of interactions and attention, and d) dNRI [5], a VAE model with recurrent GNN modules that re-evaluates the predicted latent graph dynamically.

The first two are deterministic, non-variational baselines, while IMMA and dNRI are generative models. We further evaluate our method with and without disentanglement. As restricted labeling requires ground truth labels to be provided for each interaction, it is reserved for certain datasets, such as the Spring dataset in Section V-B. Otherwise, we resort to pair matching, an unsupervised approach, which we use to learn a set of time-invariant variables of the feature space that characterize specific aspects of certain interactions.

Ablation Studies. dVAE is a variant of dG-VAE, characterized by its lack of edge features. When referring to dG-VAE, unless otherwise specified, we are addressing the version without disentanglement. Therefore, ablation studies are intrinsically incorporated and have not been neglected.

Evaluation. As in prior work [6], [9], we evaluate the standard metrics used in trajectory forecasting: 1) The minimum average displacement error (minADE), which refers to the mean Euclidean distance between the ground truth and predicted trajectories, 2) the minimum final displacement error (minFDE), which refers to the Euclidean distance between the predicted final position and the ground truth at the prediction horizon, and 3) the root mean squared error (RMSE). We also evaluate graph accuracy, the percentage overlap between edges present in the ground truth graph and those present in the inferred graph, for the Spring dataset only, as it is the only one that has ground truth graphs. For minADE and minFDE, we evaluate $k = 20$ randomly sampled predicted trajectories and choose the best value in the set. We report the mean and standard error across 20 test subsets for each metric.

Implementation. dG-VAE consists of a GNN component with a hidden dimension of 128, a RNN with a hidden dimension of 64, and read-out heads with a hidden dimension of 256. The dimension of edge feature embedding is set to 16 for the Spring dataset, 32 for the NBA dataset, 32 for the Motion Capture dataset, and 64 for the inD dataset, respectively. We allocated 16 nodes to temporal pair matching, when applicable. The decoder RNN has a hidden dimension of 256. The sampler uses a Gumbel temperature of 0.5, which is the same as dNRI. The batch size and learning rate are set to 128, 2×10^{-4} for NBA, Spring, and Motion Capture datasets, and 1, 5×10^{-4} for the inD dataset. While training multiple edge types is possible, we only evaluate binary edge prediction to evaluate the impact of learned edge features without correlation bias. We trained each model with the Adam optimizer for 150 epochs for NBA, Spring, and Motion Capture, and 400 epochs for inD.

A. NBA

The NBA dataset contains 100K examples with an 80/10/10 train/validation/test split. It features the trajectories of all ten players and the ball. Each trajectory has 50 time steps at a frequency of 3.0 Hz, which translates to a prediction horizon of 3.6 s. We normalize and mean-shift the position and velocity to the range of $[-1, 1]$. The preprocessed data can be downloaded here. Unlike some other works that used this dataset, our study is not limited to the offensive team’s coordination on a half-court, ignoring defensive players and the ball, but analyzes the performance of semantically understanding the whole move.

Figure 4 shows trajectories generated by each model, highlighting notable differences. It is evident that our approach demonstrates a smaller deviation from the ground truth trajectories compared to dNRI and LSTM. Table I summarizes the testing statistics for different models where the best values for each metric are highlighted in bold. We note that

TABLE I: Performance comparison for NBA dataset

Model	RMSE [m]	min ADE [m]	min FDE [m]	Connectivity
Linear	$1.027 \pm 1.00 \times 10^{-2}$	$1.038 \pm 9.45 \times 10^{-3}$	$4.118 \pm 1.52 \times 10^{-2}$	–
LSTM	$0.884 \pm 3.55 \times 10^{-3}$	$0.830 \pm 2.02 \times 10^{-3}$	$2.895 \pm 1.03 \times 10^{-2}$	–
IMMA	$1.036 \pm 9.41 \times 10^{-3}$	$0.965 \pm 9.99 \times 10^{-3}$	$2.922 \pm 1.75 \times 10^{-2}$	–
dNRI	$1.140 \pm 5.09 \times 10^{-3}$	$0.802 \pm 3.18 \times 10^{-3}$	$4.080 \pm 2.05 \times 10^{-2}$	0.678
dG-VAE (ours)	$0.876 \pm 4.19 \times 10^{-3}$	$0.547 \pm 1.84 \times 10^{-3}$	$3.199 \pm 1.39 \times 10^{-2}$	1.000
dG-VAE + pair matching	$0.579 \pm 4.06 \times 10^{-3}$	$0.381 \pm 1.78 \times 10^{-3}$	$2.166 \pm 1.48 \times 10^{-2}$	0.999

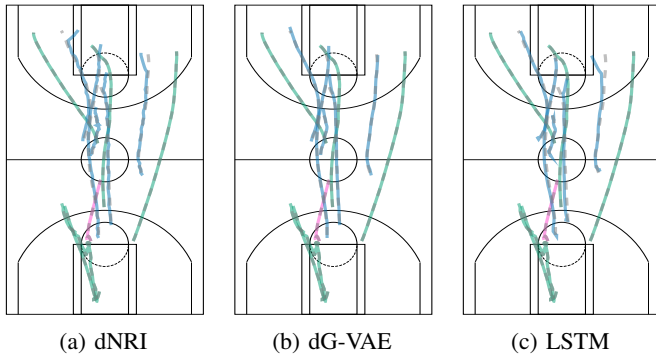


Fig. 4: (NBA) Trajectory samples predicted by different models. The grey lines represent ground truth trajectories and the blue and green lines show the predictions for home and visiting teams. The purple lines represent the basketball.

our method outperforms the strongest baseline (dNRI) in all metrics by around 25% when edge features are learned and around 45% when these features are disentangled into temporal and static features, i.e., pair matched.

Further analysis showed that dNRI, which only predicts binary edges converges to a 67% graph connectivity, while dG-VAE converges to a fully connected graph. This implies that all agents are to be considered when predicting a player’s next moves. Our method allows for this “all interactions are relevant” state, as the augmented edge features add the expressivity necessary to distinguish between interactions.

B. *k*-Spring

The Spring dataset is composed of 70K total rollouts of simulated systems with N particles, with 50K for training and 10K for testing and validation each. In our experiments, we use $N = 5$. No external forces are applied, except elastic collisions with the box boundaries. With a probability of 0.7, we connect each pair of particles with a spring. The interaction between particles linked by springs is governed by Hooke’s law, i.e., $F_{ij} = -k \cdot (r_i - r_j)$. Here, F_{ij} represents the force exerted by particle v_j on particle v_i . The spring constant is denoted by k and is uniformly sampled between 0.5 and 2, and r_i indicates the 2D coordinate of particle v_i . The out-of-distribution data is sampled with a 0.5 probability of connection and spring constants of 1, 2, or 3. Given the initial locations and velocities, which are sampled from a multivariate Gaussian, we simulate 50 time steps at a frequency of 6.0 Hz. We provide the preprocessed data here.

Figure 5 shows trajectory examples generated by each

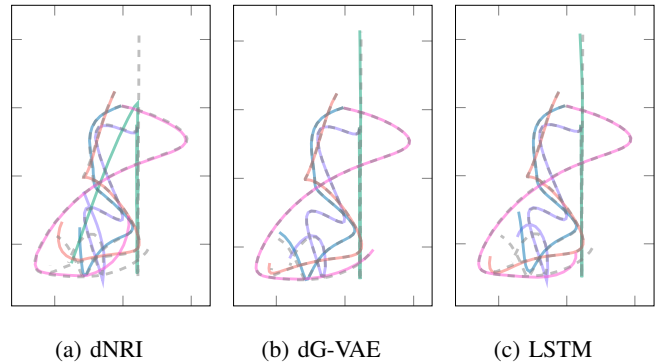


Fig. 5: (*k*-Spring) Trajectory samples predicted by different models. Each color represents a different point mass and the grey lines represent the ground truth trajectories.




Fig. 6: Comparison of inferred graphs for dNRI and dG-VAE. The edge labels refer to spring constants k , which dG-VAE is learning through restricted labeling.

model. While the green point mass is headed upwards, all other predictions cluster in the lower half of the visualization. We can clearly see an improvement when comparing our method to other baselines, the predictions are much cleaner and closer to the ground truth than dNRI and LSTM. This observation is also supported by our numerical results in Table II, where our method outperforms all other approaches by almost an order of magnitude. We note that the addition of restricted labeling yields an additional 5% improvement in RMSE and ADE and, more importantly, improves the graph accuracy to near perfect: dNRI has an accuracy of 0.877 and dG-VAE with restricted labeling has 0.984.

Figure 6 provides the inferred graphs for the same sample scenario. While dNRI infers a fully connected graph, dG-VAE shows a more adequate inference, such as the green node’s predicted trajectory to be independent of any other point mass. Restricted labeling leads to an additional improvement by inferring spring constants. While our method’s predictions are promising, inferring stronger spring constants where needed, they could be improved further.

TABLE II: Performance comparison for 5-Spring dataset

Model	RMSE [m]	min ADE [m]	min FDE [m]	OOD	
				min ADE [m]	min FDE [m]
Linear	$0.196 \pm 9.36 \times 10^{-4}$	$0.202 \pm 9.81 \times 10^{-4}$	$0.446 \pm 2.31 \times 10^{-3}$	$2.158 \pm 1.03 \times 10^{-2}$	$3.541 \pm 2.02 \times 10^{-2}$
LSTM	$0.037 \pm 5.12 \times 10^{-3}$	$0.024 \pm 6.54 \times 10^{-3}$	$0.159 \pm 1.33 \times 10^{-2}$	$0.406 \pm 1.39 \times 10^{-3}$	$1.485 \pm 1.13 \times 10^{-2}$
IMMA	$0.162 \pm 1.15 \times 10^{-3}$	$0.176 \pm 1.25 \times 10^{-3}$	$0.350 \pm 2.66 \times 10^{-3}$	$1.389 \pm 5.75 \times 10^{-3}$	$1.934 \pm 1.16 \times 10^{-2}$
dNRI	$0.114 \pm 4.69 \times 10^{-4}$	$0.058 \pm 2.32 \times 10^{-4}$	$0.469 \pm 2.29 \times 10^{-3}$	$0.379 \pm 1.11 \times 10^{-3}$	$1.372 \pm 9.17 \times 10^{-3}$
dG-VAE (ours)	$0.019 \pm 1.16 \times 10^{-4}$	$0.011 \pm 5.71 \times 10^{-5}$	$0.048 \pm 3.50 \times 10^{-4}$	$0.302 \pm 1.02 \times 10^{-3}$	$1.063 \pm 6.82 \times 10^{-3}$
dG-VAE + restricted lab.	$0.016 \pm 7.25 \times 10^{-5}$	$0.010 \pm 3.65 \times 10^{-5}$	$0.063 \pm 3.49 \times 10^{-4}$	$0.309 \pm 1.35 \times 10^{-3}$	$1.032 \pm 1.09 \times 10^{-2}$




Fig. 7: (Motion Capture #35) Samples from prediction models at time steps 25 and 50. The grey skeletons represent the ground truth, and the blue ones are the predictions.

C. Motion Capture

We also evaluate the efficacy of our proposed method using motion capture recordings sourced from the CMU Motion Capture database.¹ These recordings were obtained through a motion capture system featuring 12 infrared cameras at a frequency of 120 Hz. The dataset includes recordings that track movements across 31 distinct joints of the subject’s body. Specifically, we analyze sequence 35, which involves walking, and use sequence 39, which tracks another test subject walking on uneven grounds, as OOD samples. It is difficult to define OOD for any application, but, for the purpose of this paper, we consider these sequences to provide data for evaluating the effectiveness of our method in practical motion analysis.

In a comparative overview of our method in Table III, we notice that it significantly outperforms dNRI on all three metrics, with a 45% improvement. These findings are also reflected qualitatively in Figure 7, where dG-VAE matched the longer prediction (i.e., at $t = 50$) much closer than dNRI and LSTM, especially around the legs and feet.

While disentanglement appears to not significantly affect performance quantitatively as seen in the previous two test cases, it introduces a new level of expressivity, as illustrated in Figure 8. We first note that a very limited number of features are associated with nodes along the centerline, such as the core, shoulders, and head. This minimal correlation indicates that other factors play a more significant role in predicting future motion. Overall, the learned features seem to correlate with relatively distinct relations:

- 1) Feature 1 showcases the main relation of the upper body to the front leg, which suggests this feature might encode

¹<http://mocap.cs.cmu.edu>




Fig. 8: A subset of stationary features learned by dG-VAE.

weight shifting when a step is taken.

- 2) Feature 3 encodes the relation from the left foot and knee to the right hand and shoulder, and feature 6 encodes the opposite, i.e., the correlation between the right foot and the left hand and shoulder. These features highlight a common walking pattern where our arms swing synchronously with the opposite leg.
- 3) Feature 7 strongly relates all extremities, i.e., hands and feet, and weakly links more stationary relations. This encoding, therefore, suggests a variance in distance, e.g., while the distance between shoulders remains relatively constant, the distance between feet increases and decreases like a pendulum as each step is taken.

The features discussed here are only a subset to show how feature matrices can be interpreted like filters. We refer the reader to Section VI for a more complete overview of static and dynamic edge features.

D. inD

The inD dataset [33] is composed of vehicle tracks extracted from 33 drone recordings at four German intersections. For each track, metadata such as agent type (i.e., pedestrian, car, truck, bike), width, height, and the number of frames it is present for, is given in addition to its trajectory. The feature vector comprises x, y position, heading, and lateral and longitudinal velocity and acceleration. The positional error from the drone recordings is expected to be less than 10.0 cm. The OOD generalization is measured by deploying a model trained on inD data and evaluating its performance on the roundD dataset. As the name suggests, the roundD dataset is a sister dataset of inD but for recordings in roundabouts. Frames were taken at a 0.2s interval, resulting in a 4.0s prediction horizon when evaluating $T = 20$ steps.

The experimental results summarized in Table IV highlight that while learning edge features produces significant improve-

TABLE III: Performance comparison for Motion Capture (#35) dataset

Model	#35			#39	
	RMSE [m]	min ADE [m]	min FDE [m]	min ADE [m]	min FDE [m]
Linear	$2.726 \pm 4.74 \times 10^{-2}$	$3.989 \pm 6.16 \times 10^{-2}$	$4.493 \pm 1.25 \times 10^{-1}$	$5.466 \pm 2.91 \times 10^{-1}$	$5.750 \pm 6.55 \times 10^{-1}$
LSTM	$0.208 \pm 7.04 \times 10^{-3}$	$0.167 \pm 4.77 \times 10^{-3}$	$0.829 \pm 2.85 \times 10^{-2}$	$0.291 \pm 1.86 \times 10^{-2}$	$1.398 \pm 1.07 \times 10^{-1}$
IMMA	$0.418 \pm 9.99 \times 10^{-3}$	$0.583 \pm 1.36 \times 10^{-1}$	$1.081 \pm 2.45 \times 10^{-1}$	$1.274 \pm 2.00 \times 10^{-2}$	$2.328 \pm 3.24 \times 10^{-2}$
dNRI	$0.161 \pm 2.60 \times 10^{-3}$	$0.145 \pm 1.19 \times 10^{-3}$	$0.699 \pm 1.29 \times 10^{-2}$	$0.228 \pm 5.89 \times 10^{-3}$	$1.124 \pm 6.88 \times 10^{-2}$
dG-VAE (ours)	$0.068 \pm 1.33 \times 10^{-3}$	$0.075 \pm 1.36 \times 10^{-3}$	$0.294 \pm 5.68 \times 10^{-3}$	$0.168 \pm 8.85 \times 10^{-3}$	$0.790 \pm 4.45 \times 10^{-2}$
dG-VAE + pair matching	$0.088 \pm 3.80 \times 10^{-3}$	$0.083 \pm 2.04 \times 10^{-3}$	$0.344 \pm 1.12 \times 10^{-2}$	$0.167 \pm 6.89 \times 10^{-3}$	$0.786 \pm 4.91 \times 10^{-2}$

TABLE IV: Performance comparison for inD dataset

Model	inD			round	
	RMSE [m]	min ADE [m]	min FDE [m]	min ADE [m]	min FDE [m]
dNRI	$1.080 \pm 3.14 \times 10^{-2}$	$0.561 \pm 3.14 \times 10^{-2}$	$4.290 \pm 1.67 \times 10^{-1}$	4.831 ± 0.086	37.077 ± 0.966
dG-VAE (ours)	$1.056 \pm 2.83 \times 10^{-2}$	$0.486 \pm 1.44 \times 10^{-2}$	$4.285 \pm 1.04 \times 10^{-1}$	4.958 ± 0.076	38.985 ± 1.146
dG-VAE + pair matching	$1.104 \pm 5.55 \times 10^{-2}$	$0.558 \pm 1.57 \times 10^{-2}$	$4.441 \pm 3.36 \times 10^{-1}$	4.384 ± 0.111	34.848 ± 1.279

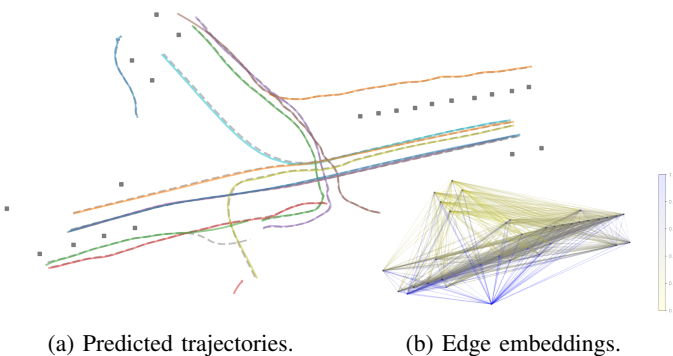


Fig. 9: (inD) (a) Trajectories sampled from dG-VAE. Grey lines represent ground truth trajectories and squares are stationary cars. (b) Edge embeddings at $t = 50$.

ments on smaller graphs, e.g., the Spring dataset with 5 nodes sees a performance increase of an order of magnitude and the Motion dataset featuring 31 nodes sees around 42%, the return is diminished on larger ones: inD with 500 nodes on average sees around 5-10% improvement. Similarly, disentangling methods have different advantages depending on the test case. While always improving interpretability and OOD generalizability, the improvements in distribution, beyond feature learning, are limited in specific test cases.

Figure 9a shows trajectory examples generated using dG-VAE. The trajectories are well matched with a few exceptions, e.g., the light green and purple trajectories deviate slightly towards the end. Figure 9b is a visualization of the embedded edge features learned by dG-VAE at the same time step. What makes this intriguing is how the model discerns between edges that connect two moving vehicles and edges linking two stationary entities, attributing high-value features to the former and near-zero feature weights to the latter. Edges that link stationary and moving agents demonstrate feature weights that lie somewhere in between. From this, we can infer that the network can successfully recognize connections significant for predicting the future positions of agents.

Limitations. Our implementation of dG-VAE has some limitations, in terms of data dependency, model generalization, and interpretability. The model generally requires large amounts of high-fidelity data for learning trends in latent relations. Generalizability in our model requires that the underlying dynamics remain unchanged, e.g., rules of the roads still apply. Although our model was designed for interpretability, additional postprocessing may be required to make the inferences more human-understandable.

VI. CONCLUSION

Our work addresses the critical challenges of interpretability and out-of-distribution generalizability in the context of interaction modeling and behavior prediction for dynamic agents. We designed a variational auto-encoder framework that integrates graph-based representations and recurrent neural networks, enabling efficient capture of spatio-temporal relations and higher prediction accuracy. Our model uses a latent space that infers dynamic interaction graphs, enriched with interpretable edge features characterizing the interactions and identifying latent factors. Furthermore, we employed two techniques to disentangle the latent space of edge features, thereby enhancing model interpretability and performance in out-of-distribution scenarios.

We demonstrate the effectiveness of learning edge features on multiple datasets, surpassing existing methods in predicting future interactions and motions. While feature space disentanglement further improved performance on smaller graphs, e.g., NBA and Spring datasets saw an additional 20% improvement, these benefits were limited on larger ones. Disentangling still yielded some gains in OOD generalization and, most importantly, created directly interpretable or latent static features that were correlated to semantically meaningful relations. In future work, we will explore possible interpretations of these disentangled embeddings.

APPENDIX

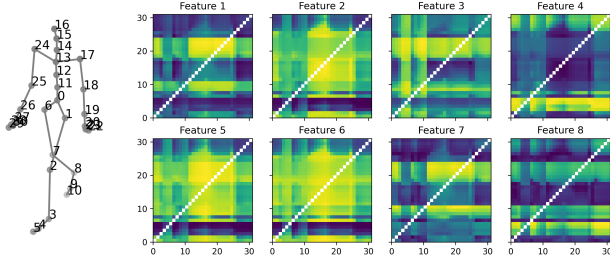


Fig. 10: Stationary features learned through pair matching.




Fig. 11: Dynamic features learned with dG-VAE at $t = 0$.

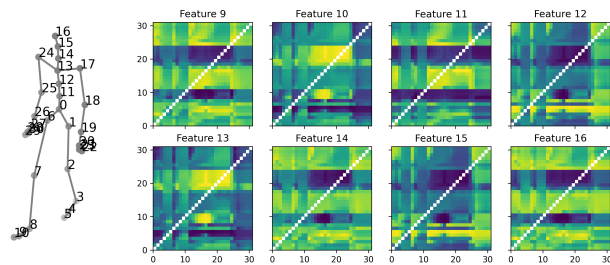


Fig. 12: Dynamic features learned with dG-VAE at $t = 50$.

ACKNOWLEDGMENT

This work is supported by Honda Research Institute USA.

REFERENCES

- [1] W. Wang, L. Wang, C. Zhang, C. Liu, and L. Sun, “Social interactions for autonomous driving: A review and perspectives,” *Foundations and Trends in Robotics*, vol. 10, pp. 198–376, 3–4 2022.
- [2] E. Binshafout, H. Ghazzai, and Y. Massoud, “Graph neural networks for traffic pattern recognition: An overview,” in *IEEE International Conference on Smart Mobility*, 2023.
- [3] Y. Li, Y. Ji, S. Li, S. He, *et al.*, “Relevance-aware anomalous users detection in social network via graph neural network,” in *International Joint Conference on Neural Networks (IJCNN)*, 2021.
- [4] T. Kipf, E. Fetaya, K.-C. Wang, M. Welling, and R. Zemel, “Neural relational inference for interacting systems,” in *International Conference on Machine Learning (ICML)*, 2018.
- [5] C. Graber and A. Schwing, “Dynamic neural relational inference,” in *IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020.
- [6] J. Li, F. Yang, M. Tomizuka, and C. Choi, “Evolveg-raph: Multi-agent trajectory prediction with dynamic relational reasoning,” in *Advances in Neural Information Processing Systems (NeurIPS)*, 2020.
- [7] E. Sachdeva and C. Choi, “Dider: Discovering interpretable dynamically evolving relations,” *IEEE Robotics and Automation Letters*, vol. 7, pp. 1–8, Oct. 2022.
- [8] A. Alahi, K. Goel, V. Ramanathan, A. Robicquet, L. Fei-Fei, and S. Savarese, “Social LSTM: Human trajectory prediction in crowded spaces,” in *IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [9] B. Ivanovic and M. Pavone, “The Trajectron: Probabilistic multi-agent trajectory modeling with dynamic spatiotemporal graphs,” in *International Conference on Computer Vision (ICCV)*, 2019.
- [10] A. Gupta, J. Johnson, L. Fei-Fei, S. Savarese, and A. Alahi, “Social GAN: Socially acceptable trajectories with generative adversarial networks,” in *IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [11] A. Sadeghian, V. Kosaraju, A. Sadeghian, N. Hirose, and S. Savarese, “SoPhie: An attentive gan for predicting paths compliant to social and physical constraints,” in *IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.
- [12] A. Mohamed, K. Qian, M. Elhoseiny, and C. Claudel, “Social-stgcn: A social spatio-temporal graph convolutional neural network for human trajectory prediction,” in *IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020.
- [13] C. Yu, X. Ma, J. Ren, H. Zhao, and S. Yi, “Spatio-temporal graph transformer networks for pedestrian trajectory prediction,” in *European Conference on Computer Vision (ECCV)*, 2020.
- [14] J. Li, F. Yang, H. Ma, S. Malla, *et al.*, “Rain: Reinforced hybrid attention inference network for motion forecasting,” in *International Conference on Computer Vision (ICCV)*, 2021.
- [15] R. Zhou, H. Zhou, H. Gao, M. Tomizuka, J. Li, and Z. Xu, “Grouptron: Dynamic multi-scale graph convolutional networks for group-aware dense crowd trajectory forecasting,” in *IEEE International Conference on Robotics and Automation (ICRA)*, 2022.
- [16] F.-Y. Sun, I. Kauvar, R. Zhang, *et al.*, “Interaction modeling with multiplex attention,” in *Advances in Neural Information Processing Systems (NeurIPS)*, 2022.
- [17] J. Li, H. Ma, Z. Zhang, J. Li, and M. Tomizuka, “Spatio-temporal graph dual-attention network for multi-agent prediction and tracking,” in *IEEE Transactions on Intelligent Transportation Systems*, 2021.
- [18] E. Sachdeva, N. Agarwal, S. Chundi, *et al.*, “Rank2Tell: A multimodal driving dataset for joint importance ranking and reasoning,” *arXiv preprint arXiv:2309.06597*, 2023.

- [19] W. Samek, T. Wiegand, and K.-R. Müller, “Explainable artificial intelligence: Understanding, visualizing and interpreting deep learning models,” *ITU Journal: ICT Discoveries - The Impact of Artificial Intelligence on Communication Networks and Services*, pp. 1–10, 2017.
- [20] C. Brewitt, B. Gyevar, S. Garcin, and S. Albrecht, “Grit: Fast, interpretable, and verifiable goal recognition with learned decision trees for autonomous driving,” in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2021.
- [21] P. Kothari, B. Stringer, and A. Alahi, “Interpretable social anchors for human trajectory forecasting in crowds,” in *IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021.
- [22] C. Tang, N. Srishankar, S. Martin, and M. Tomizuka, “Grounded relational inference: Domain knowledge driven explainable autonomous driving,” *arXiv preprint*, 2021.
- [23] L. Li, J. Yao, L. K. Wenliang, T. He, *et al.*, “GRIN: Generative relation and intention network for multi-agent trajectory prediction,” in *Advances in Neural Information Processing Systems*, 2021.
- [24] M. Li, S. Chen, *et al.*, “Online multi-agent forecasting with interpretable collaborative graph neural networks,” *IEEE Transactions on Neural Networks and Learning Systems*, pp. 1–15, 2022.
- [25] L. Sur, C. Tang, Y. Niu, E. Sachdeva, *et al.*, “Domain knowledge driven pseudo labels for interpretable goal-conditioned interactive trajectory prediction,” in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2019.
- [26] D. Lee, Y. Gu, J. Hoang, and M. Marchetti-Bowick, “Joint interaction and trajectory prediction for autonomous driving using graph neural networks,” *arXiv preprint arXiv:1912.07882*, 2019.
- [27] N. Lee, W. Choi, P. Vernaza, C. B. Choy, *et al.*, “DESIRE: Distant future prediction in dynamic scenes with interacting agents,” in *IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.
- [28] I. Higgins, L. Matthey, A. Pal, *et al.*, “Beta-VAE: Learning basic visual concepts with a constrained variational framework,” in *International Conference on Learning Representations (ICLR)*, 2017.
- [29] I. Higgins, A. Pal, A. A. Rusu, L. Matthey, *et al.*, “Darla: Improving zero-shot transfer in reinforcement learning,” in *International Conference on Machine Learning (ICML)*, 2017.
- [30] Y. Wang, Y. Sun, Z. Liu, S. E. Sarma, *et al.*, “Dynamic graph CNN for learning on point clouds,” in *IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [31] C. J. Maddison, A. Mnih, and Y. W. Teh, “The concrete distribution: A continuous relaxation of discrete random variables,” in *International Conference on Learning Representations (ICLR)*, 2017.
- [32] A. Gabbay and Y. Hoshen, “Demystifying inter-class disentanglement,” in *International Conference on Learning Representations (ICLR)*, 2020.
- [33] J. Bock, R. Krajewski, T. Moers, S. Runde, *et al.*, “The inD dataset: A drone dataset of naturalistic road user trajectories at German intersections,” in *IEEE Intelligent Vehicles Symposium (IV)*, 2020.