

# A Neuro-Symbolic Approach for Enhanced Human Motion Prediction

Sariah Mghames<sup>1</sup>, Luca Castri<sup>1</sup>, Marc Hanheide<sup>1</sup>, Nicola Bellotto<sup>1,2</sup>

*Abstract*—Reasoning on the context of human beings is crucial for many real-world applications especially for those deploying autonomous systems (e.g. robots). In this paper, we present a new approach for context reasoning to further advance the field of human motion prediction. We therefore propose a neuro-symbolic approach for human motion prediction (NeuroSyM), which weights differently the interactions in the neighbourhood by leveraging an intuitive technique for spatial representation called Qualitative Trajectory Calculus (QTC). The proposed approach is experimentally tested on medium and long term time horizons using two architectures from the state of art, one of which is a baseline for human motion prediction and the other is a baseline for generic multivariate time-series prediction. Six datasets of challenging crowded scenarios, collected from both fixed and mobile cameras, were used for testing. Experimental results show that the NeuroSyM approach outperforms in most cases the baseline architectures in terms of prediction accuracy.

## I. INTRODUCTION

Human motion prediction has been the area of focus of many researchers to date, ranging from single human motion prediction (i.e. with no context) to the most developed frameworks in context-aware (dynamic and static context) human motion prediction. The importance given to this area of study traces back to the crucial impact it has on many real-world applications including but limited to video surveillance, anomaly detection, action and intention recognition, autonomous driving, and robot navigation. While many studies on human motion prediction have been relying on datasets collected from a fixed camera to enhance the accuracy and time complexity of their frameworks, very few have studied the field from a mobile camera perspective where the problem becomes more challenging with restriction on the global observability of the scene context and interactions. Hence, we focus in this work on studying the field from both fixed and mobile camera perspective targeting the autonomous systems (e.g. robotic) application.

Reasoning on the context (e.g multi-agents interactions, static key objects) of humans is crucial primarily for safe autonomous systems navigation where human-robot co-existence, for example, is increasingly taking part in domestic, healthcare, warehouse, and transportation domains. A robot tasked to deliver an order to a table in a restaurant, bring a medicine to a patient in a hospital, or clean a road side-walks, needs to update on-the-fly its internal state representation of dynamic agents in the scene and, therefore, update its target plan.

<sup>1</sup>School of Computer Science, University of Lincoln, UK.

<sup>2</sup>Dept. of Information Engineering, University of Padua, Italy. This project has received funding from the EU’s Horizon 2020 Research and Innovation programme under grant agreement No 101017274




Fig. 1: Conceptual illustration of the social cafe-bar scenario used for reasoning on the context and interactions of humans in dense environments. In the area of context-aware human motion prediction, a research gap remains for embedding all the neighbourhood interactions in the learning process without reasoning on the ones that are more or less stable (i.e. reliable) than others, and hence can be more or less important in affecting the future states of a single agent.

In addition to safety in navigation, reasoning on the motion of multi-agents presents also the advantage of implicit intent communication. For example, a social robot detecting a conversational group in the environment and predicting that the group will hold on its current interaction for a while, makes sure to not unnecessarily interfere with the group. On the other hand, if the robot predicts that a person is coming towards itself, e.g to handle a box in a warehouse, it can select an action that prioritizes the responsiveness to the human’s intent.

Context reasoning is presented over the literature in the form of human-human and human-objects interactions reasoning. The authors in [1] have jointly modeled human-robot and human-human interactions in a deep reinforcement learning framework to drive robot navigation. In [2], instead, the authors learn an optimal local trajectory from a global plan by fusing human trajectories, Lidar features, global path and odometry features in an attention layer. Context-awareness methods have also been proposed to deal with the challenges faced by the long-term prediction of single human motion [3]–[8]. In the previous works, interactions are processed either in a grid-based pooling approach or in a global pooling mechanism to deal with the problem of dynamic neighbourhood size.

Though context-awareness has proven better accuracy in predicting human motion, one problem remains unnoticed.

Indeed, when interactions are defined spatially and hence retrieved from the relative motion between pairs of agents, not all neighbourhood interactions are of equal reliability and hence of equal importance to the prediction of future states of a single human being, as can be the case for the spatial interactions interconnecting the queue line and the table in Fig. 1. In this work, we address the problem of context-aware human motion prediction by injecting a-priori information on the interactions in a neuro-symbolic approach. Among spatial interaction representations, the qualitative approach presents an intuitive way for interaction description. Qualitative spatial interactions are defined as symbolic representations of interactions between a pair of agents in the spatial domain, i.e. 2D navigation. One way to model qualitative spatial interactions in multi-agent scenarios is by using the qualitative trajectory calculus (QTC) [9], [10]. QTC-based models of moving agent pairs can be described by different combinations of QTC symbols that represent spatial relations between pairs of interacting agents, like relative distance (i.e. moving towards/away), velocity (i.e. moving faster/slower), and orientation (i.e. moving to left/right).

The contribution of this paper is therefore three-fold: (i) proposing a novel neuro-symbolic approach for enhancing human motion prediction (denoted NeuroSyM) using a-priori information on the spatial interactions between couple of agents, (ii) experimentally evaluating the proposed framework on two architectures from the state of art, one of which is a baseline for human motion prediction and the other is a baseline for generic multivariate time-series prediction, and on different datasets collected from both fixed and mobile camera perspective, (iii) releasing the source code as a *GitHub* repository<sup>1</sup> with some qualitative results to help in the testing and integration of NeuroSyM on other baselines for human motion prediction.

The remainder of the paper is as follows: Sec. II presents an overview of the related works; Sec. III explains the approach adopted to reason on the context of humans in dense scenes; Sec. IV illustrates and discusses the results from experiments conducted on open-source datasets for human motion prediction and social navigation; finally, Sec. V concludes by summarising the main outcomes and suggesting future research work.

## II. RELATED WORKS

**Context-aware human motion prediction:** The state of the art shows extensive works in the area of context-aware human motion prediction. Among those works, some incorporate spatio-temporal dependencies [11]–[13] of interactions and others are limited to only spatial [3]–[8], [14], [15]. Another sub-category differentiates related works into those dealing with both dynamic and static context [4], [7], [11], [14], [15], those neglecting the dynamic context [8], while others focus more on the dynamic context of interactions only [3], [5], [6], [12], [13].

<sup>1</sup><https://github.com/sariahmghames/NeuroSyM-prediction>

The Dynamic and Static Context-aware Motion Predictor (DSCMP) in [11] integrates dynamic interactions between agents in a Social-aware Context Module (SCM), whereas the static context is incorporated in a latent space with a semantic scene mapping. The two most common baseline architectures used over the literature for human motion prediction are the Social-LSTM (S-LSTM) [3] and the Social Generative Adversarial Networks (SGAN) [5]. They use a spatially-aware pooling mechanism for incorporating the hidden states of proximal dynamic agents as a way to overcome the problem of variable and (potentially) large number of people in a scene. The SGAN, however, has outperformed S-LSTM in terms of both accuracy and time complexity by avoiding the grid-based pooling mechanism technique. In parallel, SGAN outperformed Stgat [13] with time complexity and parameters consumption.

In this work, the fundamental SGAN architecture from the literature is used to evaluate our NeuroSyM approach for motion prediction, leaving room for potential integration of other architectures with static context awareness (e.g. the image-driven static context of [11]). Here we consider only raw trajectories (or metric coordinates) of the context (dynamic and/or static) as possible input to the deployed model architectures.

**Human-human interactions modeling:** The methods for interactions modeling with nearby dynamic agents can be classified into two types of problem: (a) one-to-one modeling, and (b) crowd modeling [16]–[18]. A one-to-one interaction modeling was presented in the literature in the form of quantitative or qualitative representations. Quantitative representation of interactions leverages a multi-layer perceptron to embed relative pose (positions or velocities) between pairs of agents as in [1], [3], [5]. While qualitative representation of interactions was used in [19] and [20] to model human-robot spatial interactions using QTC. In [19] the use of qualitative rather than quantitative representations for analysing human-robot spatial interactions (HRSI) was motivated by the need of a more intuitive understanding of the observed interactions. In [21] and [22] similar models are used to implement human-aware robot navigation strategies. The prediction of interactions in [21] is based on a Bayesian temporal model limited to single human-robot pairs, without considering nearby static or dynamic objects, which limits the prediction performance.

In our study of single human motion, we rely on one-to-one (i.e. pairwise) interactions modeling due to the different nature of interactions that may occur in the neighbourhood of a single agent. Hence, we build on previous works from qualitative representation of interactions [19] to weight the quantitative embedding of neighbourhood interactions.

## III. NEUROSYM PREDICTION APPROACH

### A. Problem Definition

While most, if not all, works in the area of context-aware human motion prediction embeds equally all kind of interactions in the pre-defined neighbourhood size of

a single agent, in this work we formulate the problem of context-aware human motion prediction in terms of weighted interactions embedding between pairs of agents. We show that a-priori information on the kind of interactions helps the network to predict motion with a better accuracy. In the following, we present the formulation of spatial interactions which will be used later on to label (or weight) the interactions as a symbolic reasoning called by the neural model.

### B. Spatial Interactions: a Qualitative Formulation

A qualitative spatial interaction is defined by a vector of  $m$  QTC relations [9], which consist of qualitative symbols  $(q_i, i \in \mathbb{Z})$  in the domain  $U = \{-, 0, +\}$ . We can distinguish between four types of QTC: (a)  $QTC_B$  basic, (b)  $QTC_C$  double-cross, (c)  $QTC_N$  network, and (d)  $QTC_S$  shape. Here, we focus on the use of  $QTC_C$ , since it better represents the dynamics of the agents in our application scenario. Two types of  $QTC_C$  exist in the literature:  $QTC_{C_1}$ , with four symbols  $\{q_1, q_2, q_3, q_4\}$ , and  $QTC_{C_2}$ , with six symbols  $\{q_1, q_2, q_3, q_4, q_5, q_6\}$ . The symbols  $q_1$  and  $q_2$  represent the towards/away (relative) motion between a pair of agents;  $q_3$  and  $q_4$  represent the left/right relation;  $q_5$  indicates the relative speed, faster or slower; finally,  $q_6$  depends on the (absolute) angle with respect to the reference line joining a pair of agents. The  $QTC_{C_1}$  type is illustrated in Fig. 2 for a case of interaction between three body points. Given the time series of two moving points,  $P_k$  and  $P_l$ , the qualitative interaction between them is expressed by the symbols  $q_i$  as follows:

$$\begin{aligned}
(q_1) \quad & -: d(P_k|t^-, P_l|t) > d(P_k|t, P_l|t) \\
& 0: d(P_k|t^-, P_l|t) = d(P_k|t, P_l|t) \\
& +: d(P_k|t^-, P_l|t) < d(P_k|t, P_l|t) \\
(q_2) \quad & \text{same as } q_1, \text{ but swapping } P_k \text{ and } P_l \\
(q_3) \quad & -: \|P_k^{\vec{t}} P_k^t \wedge P_l^{\vec{t}} P_l^t\| < 0 \\
& 0: \|P_k^{\vec{t}} P_k^t \wedge P_l^{\vec{t}} P_l^t\| = 0 \\
& +: \text{all other cases} \\
(q_4) \quad & \text{same as } q_3, \text{ but swapping } P_k \text{ and } P_l \\
(q_5) \quad & -: \|\vec{V}_k^t\| < \|\vec{V}_l^t\| \\
& 0: \|\vec{V}_k^t\| = \|\vec{V}_l^t\| \\
& +: \text{all other cases} \\
(q_6) \quad & -: \theta(\vec{V}_k^t, P_k^{\vec{t}} P_l^t) < \theta(\vec{V}_l^t, P_l^{\vec{t}} P_k^t) \\
& 0: \theta(\vec{V}_k^t, P_k^{\vec{t}} P_l^t) = \theta(\vec{V}_l^t, P_l^{\vec{t}} P_k^t) \\
& +: \text{all other cases.}
\end{aligned}$$

where  $d(\cdot)$  is the euclidean distance between two positions,  $V(\cdot)$  the velocity vector of a body point,  $\theta(\cdot)$  is the absolute angle between two vectors, and  $\wedge$  is the cross-product notation between two vectors. In this paper, we propose a neuro-symbolic approach for motion prediction (NeuroSyM) that can be implemented to every related work in the field to enhance the accuracy of the motion. In order to narrow down the study, we take advantage of  $QTC_{C_1}$  to




Fig. 2: A case of  $QTC_{C_1}$  representation of interactions between three body points  $P_k$ ,  $P_l$ , and  $P_q$ .

label the interactions as described in the following section. We leave therefore the investigation into the additional information provided by  $QTC_{C_2}$  to our future work.

### C. Data Labeling

We leverage our labeling technique for pairwise spatial interactions on the concept of Conceptual Neighbourhood Diagram (CND) presented in [23] and in the original work of qualitative spatial interactions in [9]. As per [9], the construction of a CND (as in Fig. 3) for QTC is based on the notion of conceptual distance ( $\mathbf{d}$ ), which is used to define the closeness of two QTC states at time  $t$  and  $t'$ , respectively, and can be calculated as follows:

$$\mathbf{d}_{QTC^t}^{QTC^{t'}} = \sum_{q_i} |q_i^{QTC^t} - q_i^{QTC^{t'}}|, \quad (1)$$

where, for practical reasons, the symbols “+” and “-” are associated to the numerical values “+1” and “-1”, as in [9]. In Fig. 3 (left), for each link (i.e. edge) between conceptual neighbours (the nodes) the conceptual distance between the adjacent relations is indicated. In a CND, and due to the laws of continuity, the conceptual neighbours of each particular relation constitute only a subset of the base relations. For example,  $QTC_{C_1}$  has 81 basic states or relations (each symbol  $q_i$  has 3 different possible types of transitions from domain  $U$ ) but the conceptual neighbours of  $\{-, -, -, -\}$  relation as illustrated in Fig. 3 (right) reduce from 80 to 15 for the following reasons [23]:

- Transition from “+” to “-” (and vice versa) is impossible without passing through 0, hence transition from  $\{-, -, +, +\}$  to  $\{-, +, +, +\}$  is impossible without passing through  $\{-, 0, +, +\}$ .
- “0” Dominates “+” and “-”, hence a transition from  $\{+, -, -, 0\}$  to  $\{+, -, 0, +\}$  is impossible without passing through  $\{+, -, -, +\}$  or  $\{+, -, 0, 0\}$ .
- The combination of both former rules.

For the sake of labeling, we omit the need for information on the conceptual distance between states, and we focus on the possible transitional states for each QTC relation given the state at time  $t$ . The CND for  $QTC_{C_1}$  is not completely shown in Fig. 3 (right) as it is too complex to visualise on a two-dimensional medium. The label ( $\alpha_{cnd}$ ) for each of the 81 states of a  $QTC_{C_1}$  type of qualitative calculus is formulated as follows:

$$\alpha_{cnd} = \Pr(QTC^{t'} | QTC^t) = \frac{1}{N_{Tr}} \quad (2)$$




Fig. 3: (left) The complete CND for  $QTC_{B1}$  in  $n$ -dimensional space. The  $QTC_{B1}$  has only  $q_1$  and  $q_2$  symbols. Straight edges represent a conceptual distance of 1 while dashed edges represent a conceptual distance of 2. (right) one part of the  $QTC_{C1}$  illustrating the possible transitions of the QTC relation  $\{-, -, -, -\}$ , resulting in 15 possible transitions and therefore in  $\alpha_{cnd}$  of 0.067 weighting that same QTC relation at the next time step.

where  $N_{Tr}$  represents the number of transitional states.  $\alpha_{cnd}$  represents the level of stability or reliability of a transitional state. The higher the number of possible transitional states, the lower the likelihood to transition into a single state and vice versa. In Fig. 3, the likelihood to transition from  $\{-, -, -, -\}$  into  $\{0, 0, 0, 0\}$  is 0.067, however the likelihood increases to 0.2 if the 15 possible transitional states reduces to 5, rendering the  $\{0, 0, 0, 0\}$  state more reliable in the learning process. Given an interaction at time  $t$ , we associate its label to the interaction (observed or predicted) at  $t + 1$ . In most related works, an interaction between agents A and B is calculated as an embedding of the relative pose between them, as follows:

$$Inter_{AB} = Dense(X_B - X_A) \quad (3)$$

where  $Dense()$  is the embedding layer. Given the pose  $X$  of each agent at time  $t$ , a QTC state can be formulated and the corresponding label “ $\alpha_{cnd}$ ” will be loaded from a dictionary. Hence, the symbolic reasoning transforms Eq. 3 into the form of  $\alpha_{cnd} Inter_{AB}$ . From a practical point of view, the symbolic knowledge of interactions between two moving body-points can be readily exploited by any neural architecture for context-aware motion prediction, since the CND dictionary (associating QTC states with their corresponding  $\alpha_{cnd}$ ), generated for a specific QTC configuration, remains the same regardless of the data distribution domain.

#### IV. EXPERIMENTS

As anticipated in Sec. II, we evaluated our approach for enhancing human motion prediction on two architectures from the state of the art, using raw trajectories as input. The first one is a well-known baseline architecture, the socially-acceptable trajectories with generative adversarial networks (SGAN [5]), used in the literature to enhance the accuracy and speed of human motion prediction in crowds. It relies on datasets collected from a fixed top-down camera



(a) Zara dataset

(b) University dataset



(c) JackRabbit cafe *bytes-cafe-2019-02-07.0* scene

Fig. 4: Examples from UCY and JackRabbit datasets.

in public spaces, capturing the entire scene dynamics (ETH dataset [24] – sequences ETH and Hotel), and on datasets collected from a mobile stereo rig mounted on a car (UCY dataset [25] – sequences Zara01-02 and Univ). In order to generalise our evaluation to robotics application, we chose another well-known dataset, the JackRabbit (JRDB [26]), which provides multi-sensor data of human behaviours from a mobile robot perspective in populated indoor and outdoor environments (Fig. 4c). JackRabbit was never exploited in the literature for human motion prediction, although it clearly benefits applications of social robot navigation, where local interactions can be extracted from the on-board 360° Lidar (Velodyne) and Fisheye camera sensors. To this end, we chose to use JackRabbit on a generic network architecture for time series prediction and where the following features can be incorporated: (a) the ability to integrate a dynamic context; (b) the ability to integrate key static objects of potential interactions (e.g. door, table, bar), differently from S-LSTM and SGAN; (c) the ability to test our neuro-symbolic approach on prediction architectures that, instead of using a pooling mechanism to overcome the size problem of dynamic input series (representing the neighbourhood in social scenarios), weights every single input (i.e. neighbour) by giving special attention to each one separately. One of the recent architectures that satisfy the last features is the dual-stage attention mechanism (DA-RNN) developed for time-series forecasting in [27].

##### A. Neuro-Symbolic SGAN

**NeuroSym SGAN Architecture:** The original SGAN architecture (SGAN-20VP-20 in [5]) has proven good performance in terms of accuracy, collision avoidance, and time complexity with respect to its precedent baselines, as the social LSTM [3]. The core of SGAN is a generator and a discriminator trained adversarially. The generator model  $G$  has the role of generating candidate trajectories, while the discriminator model  $D$  estimates the probability that a sample comes from the training data (i.e. real) rather than from the generator output samples. The generator consists of an encoder and a decoder, separated by a pooling mechanism, while the discriminator is mainly an encoder. In SGAN, a variety loss is introduced on top

of the adversarial (min-max) loss in order to encourage the generator to output diverse samples, thanks to a noise distribution injected to the pooling mechanism output. For details on the SGAN architecture, the reader is advised to refer to the original work [5]. The performance measures used in SGAN for the evaluation process are the absolute displacement error (ADE) and the final displacement error (FDE) of the predicted trajectory ( $\tilde{X}$ ). The measures are calculated as follows:

$$ADE = \frac{\sum_{i=1}^N \sum_{t=1}^{T_{pred}} \|\tilde{X}_t^i - X_t^i\|_2}{N * T_{pred}} \quad (4)$$

$$FDE = \frac{\sum_{i=1}^N \|\tilde{X}_{T_{pred}}^i - X_{T_{pred}}^i\|_2}{N} \quad (5)$$

where  $N$  is the total number of training trajectories.

The neuro-symbolic version of SGAN proposed in this paper is illustrated in Fig. 5, highlighting the difference to the original pooling mechanism of SGAN [5]. NeuroSyM acts mainly on the pooling mechanism of the predictive models, where it represents human-human interactions by (a) embedding first their relative pose in all the observed states of each agent through a dense layer, then (b) weighing the embedded relative pose based on the CND-inspired label ( $\alpha_{cnd}$ ) associated to the interaction at a previous time step, and finally (c) max-pooling the weighted embedding across neighbours in the global scene. On the contrary, the original SGAN considers relative poses at the final observed state only, with no attention given to the reliability or stability level the interactions might have to help inferring future states of the agent under consideration.

**Results:** For a reliable comparison between SGAN and NeuroSyM SGAN, we trained again the former on our computing system (11th Gen Intel® Core™ i7-11800H processor and NVIDIA GeForce RTX 3080 16GB GPU), which was able to replicate almost the same hyper-parameters of the original work on SGAN, except for the batch size, in our case limited to 10 instead of 64. A comprehensive list of the hyper-parameters used to train and validate all model architectures is reported in the appendix Sec. V. The ADE and FDE results for both architectures are reported together with their standard deviations (DE-STD and FDE-STD) in Table I for  $T_{pred} = 8$  steps (i.e. 3.2 seconds) and 12 steps (i.e. 4.8 seconds), and on the five sequences from the publicly available datasets ETH and UCY. The results show a better ADE, FDE, DE-STD, and FDE-STD for the NeuroSyM approach compared to the original SGAN. The relative gain in terms of error drop is represented in Table I by a positive percentage for all the four measures with NeuroSyM with respect to SGAN on each dataset. The average relative gain for ADE, FDE, DE-STD, and FDE-STD, over the 5 datasets, is 60.84%, 58.4%, 28%, and 33.68%, respectively, for  $T_{pred} = 8$ ; and 78.58%, 76.97%, 43.5%, 46.3%, respectively, for  $T_{pred} = 12$ .

### B. Neuro-Symbolic DA-RNN

**NeuroSyM DA-RNN Architecture:** The original DA-RNN architecture [27] implements a dual-stage attention

mechanism for time-series forecasting. The dual-stage network consists of an encoder with an input attention module weighing the  $n^*$  time-series data spatially, each of length  $T_h$ , where  $T_h$  is the observed time history. The encoder is then followed by a decoder with a temporal attention layer, capturing the temporal dependencies in the input series. The encoder and decoder are based on an LSTM recurrent neural network. The network outputs the prediction of one time-series data of length  $T_f$ , where  $T_f$  is the predictive time horizon. The reader can refer to [27] for a detailed explanation of the network components, where  $T_f$  was limited to 1.

The NeuroSyM version of DA-RNN we propose in this paper takes advantage of the symbolic knowledge of the spatial interactions between pairs of agents. In DA-RNN, the encoder attention weights (“ $\alpha$ ” in Fig. 6) highlights the importance of each input series at time  $t$  on the output prediction at  $t + 1$ . The input attention weights in DA-RNN are calculated as follows:

$$\alpha_t^k = \frac{\exp(e_t^k)}{\sum_{i=1}^n \exp(e_t^i)} \quad (6)$$

where  $e_t^k$  is the embedding of the  $k^{th}$  input series at time  $t$ . It is implemented as:

$$e_t^k = \text{dense}[\text{tanh}(\text{dense}(\mathbf{h}_{t-1}; \mathbf{s}_{t-1}) + \text{dense}(\mathbf{x}_{1..T_h}^k))] \quad (7)$$

where  $h_{t-1}$  and  $s_{t-1}$  are the hidden and cell state of the encoder LSTM at a previous time step. The NeuroSyM DA-RNN acts on the input series embedding  $e_t^k$  before the softmax function (Eq. 6) is applied on it. Hence, the NeuroRoSyM approach transforms Eq. 7 into  $\alpha_{cnd,t}^k e_t^k$ , updating the encoder attention weights with an a-priori knowledge of the reliability or stability of each input series. For applications of human motion prediction in crowds (i.e. with context),  $\alpha_{cnd,t}^k$  is generated from Eq. 2. Each input series represents the motion history of a neighbour agent, whereas the first time series is the motion history of the considered person and the output is the predicted motion of that specific agent. Fig. 6 illustrates schematically where the NeuroSyM module intervenes on the original DA-RNN architecture with the injection of a CND layer at the interface between the embedding and the softmax layers.

**Data Processing:** Social dense scenarios as the ones presented in the JackRabbit dataset often have an unpredictable number of people entering ( $P_e$ ) and leaving ( $P_l$ ) the environment, possibly leading to a combinatorial explosion in the input size of the predictive model and in its number of training parameters (i.e. when  $P_e \gg P_l$  and  $P_e$  is very large). Indeed,  $n$  individuals in a scene results in  $n(n-1)/2$  pairwise data points. That is in addition to the difficulty of deploying an online model with variable input size. As a consequence, we implement a crowd clustering approach on JRDB for local interactions embedding (as shown in Figs. 1 and 6). For each agent  $i$  in a given scene, we generate a cluster with a fixed interaction radius  $R = 3.7m$ . The latter is selected based on the proxemics’ literature [28], where the social distance




Fig. 5: The neuro-symbolic SGAN pooling mechanism. The difference with the original SGAN pooling mechanism can be seen from the mixed arrows colour added within and outside the red grid to represent different types of spatial relations or interactions with the central agent standing on the red spot, which can be inferred from the NeuroSyM SGAN architecture.

Measure	Model	Zara1	Zara2	Hotel	Univ	ETH	Mean Gain
ADE	Baseline (SGAN)	0.7 / 2.29	0.44 / 0.95	1.76 / 2.45	1.25 / 2.96	0.88 / 3.8	—
	NeuroSyM (SGAN)	<b>0.21 / 0.34</b>	<b>0.2 / 0.3</b>	<b>0.35 / 0.5</b>	<b>0.36 / 0.62</b>	<b>0.63 / 0.73</b>	—
	Relative Gain (%)	+70 / +85.15	+54.5 / +68.4	+80.10 / +79.6	+71.2 / +79	+28.4 / +80.78	+60.84/+78.58
FDE	Baseline (SGAN)	1.31 / 4.33	0.84 / 1.85	3.33 / 4.55	2.31 / 5.79	1.63 / 6.71	—
	NeuroSyM (SGAN)	<b>0.41 / 0.7</b>	<b>0.4 / 0.61</b>	<b>0.67 / 0.99</b>	<b>0.74 / 1.31</b>	<b>1.25 / 1.44</b>	—
	Relative Gain (%)	+68.7 / +83.8	+52.38 / +67	+79.88 / +78.2	+67.9 / +77.37	+23.3 / +78.5	+58.4/+76.97
DE-STD	Baseline (SGAN)	0.35 / 0.9	0.26 / 0.52	0.44 / 0.9	0.51 / 0.84	0.37 / 1.02	—
	NeuroSyM (SGAN)	<b>0.22 / 0.4</b>	<b>0.2 / 0.35</b>	<b>0.32 / 0.56</b>	<b>0.24 / 0.38</b>	<b>0.37 / 0.64</b>	—
	Relative Gain (%)	+37.14 / +55.5	+23 / +32.7	+27.27 / +37.7	+52.9 / +54.7	+0 / +37.2	+28/+43.5
FDE-STD	Baseline (SGAN)	1.1 / 2.86	0.8 / 1.61	1.51 / 2.88	1.6 / 2.63	1.15 / 3.29	—
	NeuroSyM (SGAN)	<b>0.64 / 1.2</b>	<b>0.58 / 1.05</b>	<b>0.96 / 1.72</b>	<b>0.68 / 1.14</b>	<b>1.09 / 1.91</b>	—
	Relative Gain (%)	+41.8 / +58	+27.5 / +34.78	+36.4 / +40.27	+57.5 / +56.6	+5.2 / +41.9	+33.68/+46.3

TABLE I: Performance comparison between the baseline architecture SGAN and its neuro-symbolic approach across all datasets. We report results in the format of 8/12 prediction time steps. ADE, FDE, DE-STD, and FDE-STD measures are in meters and in bold is highlighted the better measure among the two approaches. The lower error the better. The mean gain represents the mean of the relative gains over the 5 datasets, hence it only applies to the relative gain rows.

for interactions among acquaintances is indeed between 2.1m and 3.7m. Each cluster includes  $n$  input series, with  $n$  being the maximum number of agents entering the cluster of agent  $i$  in a time interval  $T$ . The maximum number of input series among all clusters,  $n^*$ , is fixed for practical (training) purposes. Each cluster is then post-processed to include  $(n^* - n)$  input series with complementary “fake” values. We make use of the open-source annotated 3D point clouds from JRDB, provided as metric coordinates of human (dynamic context) bounding boxes centroid, and extracted from the upper Velodyne sensor, as raw data and ground truth to our network architecture. The raw data are further processed to extract QTC representations of spatial interactions between pairs of agents in a cluster. We then make use of the CND dictionary to associate each local QTC representation to its corresponding weight  $\alpha_{cnd,i}^k$ , which is then called by the NeuroSyM architecture as an a-priori information on the reliability of each input series and hence their importance on the predicted output. The a-priori information is then used to weight the embedding of

the inputs. The environments considered in JRDB are fairly crowded. Among them, we selected a cafe shop (*bytes-cafe-2019-02-07\_0*). DA-RNN embeds input series using dense layers, facilitating the possibility to integrate static context as series of constant metric coordinates. In the cafe scenario, the static context includes objects such as bar order and check-out points, exit door, and drinking water station, as illustrated in Fig. 1. These objects were manually selected based on a previous investigation to identify the most common ones used by people in the scenario, although in the future we plan to learn them automatically in order to adapt to different environments. The spatial coordinates of the selected key objects are incorporated in the network architecture as any other dynamic agent.

**Results:** We evaluated the DA-RNN and the NeuroSyM DA-RNN for the cafe scenario over to medium (i.e. 48 time steps, or 3.2 seconds) and long (i.e. 80 time steps, or 5.33 seconds) term horizons. The parameters for medium and long term horizon prediction were chosen based on relevant literature of human motion prediction [5], [11]. The results




Fig. 6: A neuro-symbolic approach for attention-based time-series prediction models. Differently from SGAN-like architectures, attention-based mechanisms have no pooling modules. The diagram is extended from [27] and modified for multi-steps attention-based context-aware human motion prediction in crowded environments. The input are  $n^*$  time series of agents, within a cluster centered at the first time series, while the output is the prediction of the cluster center’s agent. The vector  $\mathbf{e}$  denotes the input embeddings normalised to  $\alpha$  after passing through the CND layer, which adds a-priori knowledge to them in the form of  $\alpha_{cnd,t}^k$ . The CND layer weights differently the spatial relations (represented by mixed arrows colour) of the neighbour agents with the central one. The vector  $\mathbf{l}$  denotes the temporal attention weights of the encoder’s hidden states output, normalised to  $\beta$ , while  $\mathbf{c}$  represents the context.  $\mathbf{X} = \{x, y\}$  is the input driving vector;  $\mathbf{Y} = \{x', y'\}$  is the label vector;  $\mathbf{h}$  and  $\mathbf{d}$  are the encoder and decoder hidden states, respectively. The input and temporal attention layers are constructed from dense layers.

of DA-RNN architecture on the JRDB dataset, and the NeuroSyM DA-RNN, are illustrated in Table II, showing root mean square error (RMSE) and mean absolute error (MAE) between the predicted  $(x', y')$  coordinates and their true labels  $(x, y)$ . We can clearly see that the NeuroSyM version of the architecture succeeds in decreasing the RMSE metric by 22% on the 48 steps prediction horizon, while influencing the performance negatively by 4% on the longer 80 steps horizon. At the same time, the neuro-symbolic approach decreased the MAE by 21% on the 80 steps horizon, while influencing negatively by 3% the 48 time steps prediction. Although the improvement percentage of NeuroSyM DA-RNN is superior, by a large extent, to its counterpart on one of the two horizon windows, the unequal performance suggests that the prediction is affected by some outliers on the long term. This issue will be addressed in our future work by exploiting the influence of cluster radius selection on the prediction, a factor that we foresee to affect the hidden states embedding of the relative context. For a complete performance evaluation, another fundamental point that we will address in the future is the difference in computational cost between neuro-symbolic architectures and their neural counterparts.

## V. CONCLUSION

In this work, we presented a neuro-symbolic approach for context-aware human motion prediction (NeuroSyM) in dense scenarios, leveraging a qualitative representation of interactions between dynamic agents to assess the type of neighbourhood interactions and weight them accordingly. We formulate the spatial interactions in terms of a qualitative trajectory calculus (QTC) and we use the conceptual neighbourhood diagram (CND) to anticipate possible interactions that might influence the future state of an individual agent. The likelihood of the next-step interaction state is used for labeling it given the current

Architecture	RMSE	MAE
DA-RNN (Baseline)	3.61 / <b>3.572</b>	<b>2.097</b> / 2.753
NeuroSyM DA-RNN	<b>2.815</b> / 3.728	2.162 / <b>2.166</b>
Relative Gain (%)	+22 / -4.37	-3.1 / +21.32

TABLE II: Performance comparison between the baseline architecture DA-RNN and the NeuroSyM approach on the JackRabbit dataset. The results’ format refers to the 48/80 prediction time steps. RMSE and MAE values are in meters, and the best results are highlighted in bold (i.e. the lower error, the better).

state of each agent (position and interaction). We tested the NeuroSyM approach on a fundamental baseline architecture for context-aware human motion prediction, i.e. SGAN, and on another baseline architecture for multivariate time-series prediction, i.e. DA-RNN. Differently from SGAN, where interactions are pooled altogether, DA-RNN includes an attention mechanism on the input time series. We show that, in most of the cases, our neuro-symbolic approach outperforms the baseline architectures in terms of prediction accuracy on the medium- and long-term horizons. We plan in our future work to test the NeuroSyM approach for motion prediction on other architectures, as S-LSTM and those incorporating static context in addition to the dynamic one. Also, we will exploit the proposed neuro-symbolic approach for human motion prediction in social robot navigation environments, incorporating causal (symbolic) models from literature as [29], [30].

## APPENDIX: HYPER-PARAMETERS

The hyper-parameters used to train and validate each of the network architectures deployed in this work are specified in Tab III. For a complete list of SGAN hyper-

HP <sup>SGAN</sup>	Val <sup>SGAN</sup>	HP <sup>DA-RNN</sup>	Val <sup>DA-RNN</sup>
Batch size	10	Batch size	5
num_iterations	8512	observed_timesteps	5
num_epochs	200	num_epochs	80
noise_dim	8	learning_rate	0.001
noise_type	Gaussian	train_ratio	0.8
pool_every_timestep	0	decay_frequency	1000
observed_timesteps	8	decay_rate	0.99
embedding_dim	16	num_labels	1
encoder_h_dim_g	32	encoder_h_dim	256
bottleneck_dim	8	decoder_h_dim	256
decoder_h_dim_g	32	validation_ratio	0.1
encoder_h_dim_d	48	test_ratio	0.1
g_learning_rate	0.0001		
d_learning_rate	0.001		
mlp_dim	64		
noise_mix_type	global		

TABLE III: (Left) SGAN and NeuRoSyM SGAN hyper-parameter values. (Right) DA-RNN and NeuRoSyM DA-RNN hyper-parameter values.

parameters and a better understanding of their roles, the reader should refer to the original open-source repository at <https://github.com/agrim Gupta92/sgan>.

## REFERENCES

- [1] C. Chen, Y. Liu, S. Kreiss, and A. Alahi, "Crowd-robot interaction: Crowd-aware robot navigation with attention-based deep reinforcement learning," in *2019 Int. Conf. on Robotics and Automation (ICRA)*. IEEE, 2019, pp. 6015–6022.
- [2] A. Pokle, R. Martín-Martín, P. Goebel, V. Chow, H. M. Ewald, J. Yang, Z. Wang, A. Sadeghian, D. Sadigh, S. Savarese, et al., "Deep local trajectory replanning and control for robot navigation," in *2019 Int. Conf. on Robotics and Automation (ICRA)*. IEEE, 2019, pp. 5815–5822.
- [3] A. Alahi, K. Goel, V. Ramanathan, A. Robicquet, L. Fei-Fei, and S. Savarese, "Social Istm: Human trajectory prediction in crowded spaces," in *Proceedings of the IEEE Conf. on Computer Vision and Pattern Recognition*, 2016, pp. 961–971.
- [4] M. Lisotto, P. Coscia, and L. Ballan, "Social and scene-aware trajectory prediction in crowded spaces," in *Proceedings of the IEEE/CVF International Conf. on Computer Vision Workshops*, 2019, pp. 0–0.
- [5] A. Gupta, J. Johnson, L. Fei-Fei, S. Savarese, and A. Alahi, "Social gan: Socially acceptable trajectories with generative adversarial networks," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2018, pp. 2255–2264.
- [6] K. Mangalam, H. Girase, S. Agarwal, K.-H. Lee, E. Adeli, J. Malik, and A. Gaidon, "It is not the journey but the destination: Endpoint conditioned trajectory prediction," in *European conference on computer vision*. Springer, 2020, pp. 759–776.
- [7] J. Liang, L. Jiang, J. C. Niebles, A. G. Hauptmann, and L. Fei-Fei, "Peeking into the future: Predicting future person activities and locations in videos," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2019, pp. 5725–5734.
- [8] Z. Cao, H. Gao, K. Mangalam, Q.-Z. Cai, M. Vo, and J. Malik, "Long-term human motion prediction with scene context," in *European Conference on Computer Vision*. Springer, 2020, pp. 387–404.
- [9] M. Delafontaine, S. H. Chavoshi, A. G. Cohn, and N. Van de Weghe, "A qualitative trajectory calculus to reason about moving point objects," in *Qualitative spatio-temporal representation and reasoning: Trends and future directions*. IGI Global, 2012, pp. 147–167.
- [10] N. Bellotto, M. Hanheide, and N. Van de Weghe, "Qualitative design and implementation of human-robot spatial interactions," in *Int. Conf. on Social Robotics*. Springer, 2013, pp. 331–340.
- [11] C. Tao, Q. Jiang, L. Duan, and P. Luo, "Dynamic and static context-aware lstm for multi-agent motion prediction," in *European Conference on Computer Vision*. Springer, 2020, pp. 547–563.
- [12] C. Yu, X. Ma, J. Ren, H. Zhao, and S. Yi, "Spatio-temporal graph transformer networks for pedestrian trajectory prediction," in *European Conf. on Computer Vision*. Springer, 2020, pp. 507–523.
- [13] Y. Huang, H. Bi, Z. Li, T. Mao, and Z. Wang, "Stgat: Modeling spatial-temporal interactions for human trajectory prediction," in *Proceedings of the IEEE/CVF int. conf. on computer vision*, 2019, pp. 6272–6281.
- [14] N. Lee, W. Choi, P. Vernaza, C. B. Choy, P. H. Torr, and M. Chandraker, "Desire: Distant future prediction in dynamic scenes with interacting agents," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017, pp. 336–345.
- [15] N. Bisagno, C. Saltori, B. Zhang, F. G. De Natale, and N. Conci, "Embedding group and obstacle information in lstm networks for human trajectory prediction in crowded scenes," *Computer Vision and Image Understanding*, vol. 203, p. 103126, 2021.
- [16] K. Ijaz, S. Sohail, and S. Hashish, "A survey of latest approaches for crowd simulation and modeling using hybrid techniques," in *Proceedings of the 2015 17th UKSIM-AMSS Int. Conf. on Modelling and Simulation*. IEEE Computer Society, 2015, pp. 111–116.
- [17] H. Hedayati, A. Muehlbradt, D. J. Szafrin, and S. Andrist, "Reform: Recognizing f-formations for social robots," in *2020 IEEE/RSSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2020, pp. 11 181–11 188.
- [18] S. Thompson, A. Gupta, A. W. Gupta, A. Chen, and M. Vázquez, "Conversational group detection with graph neural networks," in *Proceedings of the 2021 International Conference on Multimodal Interaction*, 2021, pp. 248–252.
- [19] C. Dondrup, N. Bellotto, and M. Hanheide, "A probabilistic model of human-robot spatial interaction using a qualitative trajectory calculus," in *2014 AAAI Spring Symposium Series*, 2014.
- [20] M. Hanheide, A. Peters, and N. Bellotto, "Analysis of human-robot spatial behaviour applying a qualitative trajectory calculus," in *2012 IEEE RO-MAN: The 21st IEEE Int. Symposium on Robot and Human Interactive Communication*. IEEE, 2012, pp. 689–694.
- [21] C. Dondrup and M. Hanheide, "Qualitative constraints for human-aware robot navigation using velocity costmaps," in *2016 25th IEEE Int. Symposium on Robot and Human Interactive Communication (RO-MAN)*. IEEE, 2016, pp. 586–592.
- [22] N. Bellotto, "Robot control based on qualitative representation of human trajectories," in *2012 AAAI Spring Symposium Series*, 2012.
- [23] N. Van de Weghe and P. De Maeyer, "Conceptual neighbourhood diagrams for representing moving objects," in *Perspectives in Conceptual Modeling: ER 2005 Workshops AOIS, BP-UML, CoMoGIS, eCOMO, and QoIS, Klagenfurt, Austria, October 24-28, 2005. Proceedings 24*. Springer, 2005, pp. 228–238.
- [24] A. Ess, B. Leibe, and L. Van Gool, "Depth and appearance for mobile scene analysis," in *2007 IEEE 11th international conference on computer vision*. IEEE, 2007, pp. 1–8.
- [25] A. Lerner, Y. Chrysanthou, and D. Lischinski, "Crowds by example," in *Computer graphics forum*, vol. 26, no. 3. Wiley Online Library, 2007, pp. 655–664.
- [26] R. Martín-Martín, H. Rezatofighi, A. Sheno, M. Patel, J. Gwak, N. Dass, A. Federman, P. Goebel, and S. Savarese, "Jrdb: A dataset and benchmark for visual perception for navigation in human environments," *arXiv preprint arXiv:1910.11792*, 2019.
- [27] Y. Qin, D. Song, H. Chen, W. Cheng, G. Jiang, and G. Cottrell, "A dual-stage attention-based recurrent neural network for time series prediction," *arXiv preprint arXiv:1704.02971*, 2017.
- [28] E. T. Hall and T. Edward, "Hall, the hidden dimension," *Anchor Books New York*, vol. 20, p. 71, 1969.
- [29] M. Seitzer, B. Schölkopf, and G. Martius, "Causal influence detection for improving efficiency in reinforcement learning," *Advances in Neural Information Processing Systems*, pp. 22 905–22 918, 2021.
- [30] L. Castri, S. Mghames, M. Hanheide, and N. Bellotto, "Causal discovery of dynamic models for predicting human spatial interactions," in *Social Robotics: 14th International Conference, ICSR 2022*. Springer, 2023, pp. 154–164.