Human Motion Prediction Under Social Grouping Constraints

Andrey Rudenko[1,2], Luigi Palmieri[1], Achim J. Lilienthal[2] and Kai O. Arras[1]

***Abstract*— Accurate long-term prediction of human motion in populated spaces is an important but difficult task for mobile robots and intelligent vehicles. What makes this task challenging is that human motion is influenced by a large variety of factors including the person's intention, the presence, attributes, actions, social relations and social norms of other surrounding agents, and the geometry and semantics of the environment. In this paper, we consider the problem of computing human motion predictions that account for such factors. We formulate the task as an MDP planning problem with stochastic policies and propose a weighted random walk algorithm in which each agent is locally influenced by social forces from other nearby agents. The novelty of this paper is that we incorporate social grouping information into the prediction process reflecting the soft formation constraints that groups typically impose to their members' motion. We show that our method makes more accurate predictions than three state-of-the-art methods in terms of probabilistic and geometrical performance metrics.**

I. Introduction

Long-term prediction of human motion is an important task for applications such as robot navigation in crowded environments, autonomous driving, video surveillance or human-robot collaboration. Particularly for service robots operating amidst humans, predicting future trajectories of surrounding people over longer periods of time has the potential to considerably improve the robot's capability to look and plan ahead and to avoid excessively reactive or overly conservative motion behavior in densely populated spaces.

Making accurate predictions of future pedestrian trajectories is challenging due to the many factors that influence human motion. Such factors include other agents with their intentions, actions, attributes or social rules, and the environment with its geometry, semantics or affordances. Prior art has addressed this challenging task using different approaches based on physical dynamics modeling, learning and planning methods, considering both the single-agent case and the multi-agent case, in which predictions are made jointly. A little-explored topic, addressed in this paper, is the consideration of social grouping information for motion prediction. This is motivated by the insight that social relations among people are an important factor for predicting future motion, as individuals in groups typically form and maintain certain spatial patterns, which e.g. in [1] is described by a model based on social communication between group members. Example of the group motion in real-world data is given in Fig. 1.

Fig. 1. Illustrating example of group motion in a shopping mall. Colored points show ground truth positions of people from the ATC dataset at different points in time, groups are encircled. It can be seen, that people in groups stay together and move in relatively stable formations. Social grouping is thus an important cue for long-term motion prediction.

Thus, in this paper, we present a novel planning-based approach for long-term motion prediction that accounts for social interactions and grouping of observed agents. Following our previous work [2], the presented method formulates the task as a set of Markov Decision Processes (MDP) to produce goal-directed global motion policies. For online prediction, those policies are locally modified based on social interactions between agents and soft formation constraints for agents in groups modeled by the group social force by Moussaïd et al. [1]. As we will show in the experiments with simulated and real data, this extension improves prediction accuracy. To the best of our knowledge, this is the first paper to incorporate social grouping information into long-term prediction of human motion.

The paper is structured as follows: in Sec. II we discuss the related work and in Sec. III we describe our approach. Experiments and results are presented in Sec. IV and Sec. V respectively, and Sec. VI concludes the paper.

[1]A. Rudenko, L. Palmieri and K.O. Arras are with Bosch Corporate Research, Stuttgart, Germany. {andrey.rudenko, luigi.palmieri, kaioliver.arras} @de.bosch.com.

[2]A. Rudenko and A.J. Lilienthal are with the Center of Applied Autonomous Sensor Systems (AASS), Örebro University, Sweden. achim.lilienthal@oru.se

This work has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 732737 (ILIAD).

II. Related Work

In the following we give a brief review of human motion prediction methods, group modeling approaches, and describe how they relate to this paper.

Physics-based approaches [3], [4] forward-simulate a set of dynamics equations to compute pedestrian motion. The popular social force model [5] belongs to this class which is used for motion prediction by Elfring et al. [4] and in the context of tracking by Luber et al. [6]. Other common techniques for motion prediction include learning-based and planning-based methods. Learning-based methods comprise various data-driven approaches to long-term motion prediction. There are methods that learn prototypical trajectories, “motion patterns", in a particular environment [7]. Other approaches learn typical spatial behaviors of humans navigating in social spaces [8], [9], [10]. Planning-based methods are based on the assumption that humans follow paths through the environment in an near-optimal, goal-directed manner. Such methods may use a cost function to model navigation through a known environment [11], [2] or recover costs from observed trajectories using inverse reinforcement learning [12].

Research in computational social science and human crowd dynamics has found that up to 70% of people move in groups of two and more members and that they maintain rather stable formations depending on crowd density [1], [13]. These findings motive our hypothesis that social grouping is an important cue for long-term prediction of human motion.

Related work in modeling group structure, crowd simulation or behavior analyzing of pedestrian groups include [14], [15], [16] with applications e.g. for building design or mass event planning [17], [18]. Detecting groups has also applications in video surveillance [19] and tracking [20], [21], [3], [22] where group-informed motion modeling was shown to improve data association. The detection task has been addressed using clustering of geometrically similar trajectories, estimating inter- and intra-group forces among individuals in crowds or so called coherent motion indicators [1]. Common techniques for modeling group motion include imposing attraction forces to other group members [15], [21], [18], to the geometrical center of the group [1] or to the group's leader [17], or imposing a certain relative formation in which the group is assumed to be moving [14], [16]. An extension of the social force model that uses group information was proposed by Moussaïd et al. [1].

In this paper we present a global planning-based approach for motion prediction that accounts for local collision avoidance and social interaction behavior of humans. We build on past work [2] using an MDP planning approach and extend the random walk policy sampling procedure with the group social force model according to Moussaïd et al. [1]. In contrast to related work, this approach does not require a training phase (like [7], [9]), is map-aware and predictions do not get stuck in local minima like in physics-based approaches [3], [4]. Similarly to [8], [9], [10], we predict motion of multiple agents jointly but incorporate, as the major novelty of this paper, social grouping information into the prediction process.

III. Joint Sampling MDP for Motion Prediction

In this section we present our approach. After introducing the MDP formulation in Sec. III-A, we describe our method to generate predictions from the global motion policies in Sec. III-B with a random sampling procedure which accounts for group social forces. In Sec. III-C we analyze the complexity of our algorithm.

A. MDP for Global Motion Prediction

In this section we briefly describe the model of global agent motion towards a goal, originally presented in [2]. We use the MDP-based formulation of the optimal path planning problem in a known environment. Given a 2D static map M of the environment representing occupied and free space, and a set of goal states $\mathcal{G}$, we formulate a separate MDP path finding problem for each goal $g \in \mathcal{G}$ to obtain the cost-to-go state values $V_g^*(s)$ as well as the optimal policy $\pi_g^*(s)$ in each state $s = (s_x, s_y) \in M$. Each MDP is constructed with the *absorbing zero* state in the goal position. We describe actions as orientation-velocity pairs: $a = \langle\theta, v\rangle,\ \theta \in [0, 2\pi), v \in [0, v_{\max}]$. An action $a = \langle\theta, v\rangle$ defines the deterministic transition between states $s \xrightarrow{a} s'$, calculated as $s'_x = s_x + v\cos(\theta),\ s'_y = s_y + v\sin(\theta)$. The reward function $\mathcal{R}_g(s,a)$ is constructed as a weighted sum of Euclidean distance covered by a, and the unitary cost of the target state $C(s')$, provided by the optional input semantic map $C(s)$.

To predict also alternative paths to the goal and allow deviations from the optimal policy, we relax the obtained π_g^* with the *stochastic Boltzmann policy* that assigns to each action a probability to be executed in state s proportional to its value $\hat{Q}_g^*(s,a)$. Temperature parameter α controls the level of stochasticity, i.e. the probability that sub-optimal actions are chosen by the agent. We denote the stochastic policy as π_g and compute it as in Eq. 1, where $\hat{Q}_g^*(s,a)$ is the value of action a, and $V_g^*(s)$ is the value of the optimal action.

$$a \sim \pi_g(s) \text{ with prob. } \propto \ \exp(\alpha(\hat{Q}_g^*(s,a) - V_g^*(s))) \quad (1)$$

The obtained policy π_g allows actions up to a pre-defined very large velocity $v_{\max}$. For handling individual observed velocities $v_{\text{obs}} < v_{\max}$, we use a simple *policy cutting* technique that incorporates information about v_{obs} into the obtained policy. For each person i, the action space is redefined with $v \in [0, 2\, v_{\text{obs}}^i]$. The individual stochastic policy $\hat{\pi}_g^i$ is then computed as in Eq. 2. In $\hat{\pi}_g^i$ the probability of faster actions $a = \langle\theta, v\rangle$ with $v > v_{\text{obs}}^i$ is set the same as for the symmetrically slower actions with $v < v_{\text{obs}}^i$.

$$p(a) \text{ in } \hat{\pi}_g^i \ \propto \ \begin{cases} p(\langle\theta, v\rangle) \text{ in } \pi_g, \text{ if } v \le v_{\text{obs}}^i, \\ p(\langle\theta, 2\, v_{\text{obs}}^i - v\rangle) \text{ in } \pi_g, \text{ if } v > v_{\text{obs}}^i \end{cases} \quad (2)$$

B. Joint Human Motion Prediction with Group Social Forces

In this subsection we present our method for jointly predicting trajectories of all agents in the scene. We assume that a person tracking system delivers short sequences of observed agent positions, called tracklets, and that this system also provides group detection as partitionings of individual agents into groups. These are both realistic assumptions as many tracking systems, for example [23], are able to robustly track people also across misdetection and occlusions using e.g. advanced data association techniques. Such systems have also been extended with the ability to detect and reason about social grouping hypotheses as discussed in Sec. II.

Given N people in the scene, the observed track of length $l(i)$, associated with person i, is denoted as $\mathcal{T}^i = \langle s_1^i, s_2^i, ..., s_{l(i)}^i \rangle$, where $s_t^i = (s_{x,t}^i, s_{y,t}^i)$ is the state where the person was observed at time t, and $i \in [1, \ldots, N]$. The tracklet's end $s_{l(i)}^i = s^i(t_0)$ is the position of person i at the current time t_0 and $\mathcal{T}$ is the set of all observed tracks. Membership in one and only one of the groups $Gr_h \in \mathbf{Gr}$ is assigned to each person: $i \in Gr_h$, $Gr_h \cap Gr_{h'} = \emptyset \; \forall h' \neq h$, $\cup_h Gr_h = \{1, \ldots, N\}$.

From each tracklet we derive the observed speed v_{obs}^i, orientation θ_{obs}^i and the discrete probability distribution $p^i(\mathcal{G})$ over destinations $\mathcal{G}$. We predict the final destination of person i based on the observed tracklet. Similarly to [12] and [24], for each goal $g \in \mathcal{G}$ we estimate the gradient of the cost-to-go $V_g^*(s)$ along $\mathcal{T}^i$ as the difference between the costs at s_1^i and $s_{l(i)}^i$ using a softmax function:

$$p(g) \;\propto\; \exp\left(\beta\left(V_g^*(s_{l(i)}^i) - V_g^*(s_1^i)\right)\right). \tag{3}$$

Temperature parameter β defines to what extent alternative goals are considered. Members of the same group Gr_h share the goal probability vector, computed as the average of individual vectors: $p_{Gr}^h(\mathcal{G}) = |Gr_h|^{-1} \sum_i p^i(\mathcal{G}),\; i \in Gr_h$.

1) Local Interaction and Group Motion Modeling: the social force model [5] describes how the intended motion of a person changes according to the influence the repulsive forces from other people. Formally, the social force $\mathbf{f}_{i,k}^{\text{soc}}$, emitted by person k in the direction of person i is defined as

$$\mathbf{f}_{i,k}^{\text{soc}} = a_k e^{\left(\frac{r_{i,k} - d_{i,k}}{b_k}\right)} \boldsymbol{n}_{i,k} \left(\lambda + (1-\lambda)\frac{1+\cos(\varphi_{i,k})}{2}\right), \tag{4}$$

where $a_k \geq 0$ specifies the magnitude and $b_k > 0$ the range of the force, $d_{i,k}$ is the distance between people and $r_{i,k}$ is the sum of their radii. The normalized vector $\boldsymbol{n}_{i,k}$ pointing from k to i defines the direction of the repulsive force. An anisotropic factor $\lambda \in [0,1]$ scales the force in the person's direction of motion: the force reaches its full magnitude when the angle $\varphi_{i,k}$ between the intended motion direction of person i and $\mathbf{n}_{k,i}$ is zero, and has minimal effect when $\varphi_{i,k} = \pi$. Social forces, cast on the person i by the surrounding people, are accumulated and used to change the desired direction of motion $\mathbf{F}_i^{\text{pers}}$, which in our case is the action $a = \langle \theta, v \rangle$ sampled from the stochastic policy.

An extension of the social force model to include group interaction was proposed by Moussaïd et al. [1]. Several new

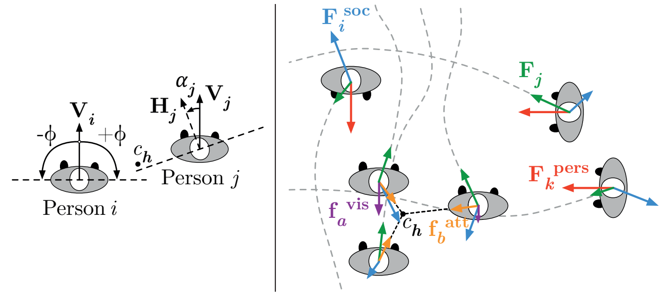

Fig. 2. **Left:** illustration of the group social force parameters. **Right:** in this example, three people a, b and c in the bottom are walking upwards as a group Gr_h. Three individual pedestrians i, j and k are opposing them from different directions. Intended directions $\mathbf{F}^{\text{pers}}$ are shown with **red** arrows, omitted for the group members for the sake of clarity. Person i is influenced by a strong social force $\mathbf{F}^{\text{soc}}$, depicted in **blue**, and has to halt and adjust the motion trajectory, shown as a **gray dotted line**. Person k stops and lets the group pass, while j attempts to cross in front of the group. Resulting motion directions $\mathbf{F}$ are shown in **green**. Intra-group social forces $\mathbf{F}^{\text{vis}}$ and $\mathbf{F}^{\text{att}}$ are shown in **blue** and **orange** respectively.

forces are defining attraction of people walking in groups to other members of the group (attraction term) and imposing soft constraints on the walking formation that resembles typical patterns of humans in groups (visibility term). For each member i of the group Gr_h, the visibility term $\mathbf{f}_i^{\text{vis}}$ is defined as

$$\mathbf{f}_i^{\text{vis}} = -\beta_1 \alpha_i \mathbf{V}_i, \tag{5}$$

where β_1 is a model parameter describing the strength of the social interaction between group members, and $\mathbf{V}_i$ is the current velocity vector of person i. This deceleration component $\mathbf{f}_i^{\text{vis}}$ is oriented in the opposite direction of current movement $\mathbf{V}_i$, and it is proportional to the angle α_i between the gazing direction $\mathbf{H}_i$ of person i and the group center of mass c_h, given the person's field of view ϕ. An illustration of the parameters is given in Fig. 2, left.

Formulation of $\mathbf{f}_i^{\text{vis}}$ imposes a *line formation*, perpendicular to the direction of motion, as the preferred walking pattern of a group. However, in order to facilitate intra-group social interactions, members of larger groups of 4 or more people often switch to the more compact *V-formation*. The same happens in cluttered spaces, as well as in crowded environments, where the members have to balance between comfortable interaction and efficient movement. To model this behavior, the attraction term $\mathbf{f}_i^{\text{att}}$ to the geometrical center of the group is introduced as

$$\mathbf{f}_i^{\text{att}} = \beta_2 q_A \mathbf{U}_i, \tag{6}$$

where β_2 is the strength of the group attraction effect, and $\mathbf{U}_i$ is the unit vector pointing from pedestrian i to the center of masses c_h of Gr_h. This force is only activated if the distance between person i and c_h exceeds a certain threshold q_A, otherwise the attraction force is zero.

The added intra-group forces $\mathbf{f}_i^{\text{vis}}$ and $\mathbf{f}_i^{\text{att}}$ yield a decelerating effect on pedestrians, whose stochastic motions often lead them in front of the group. In reality this effect is not present as humans by nature are able to better coordinate their motion within the group. To counterbalance the deceleration effect

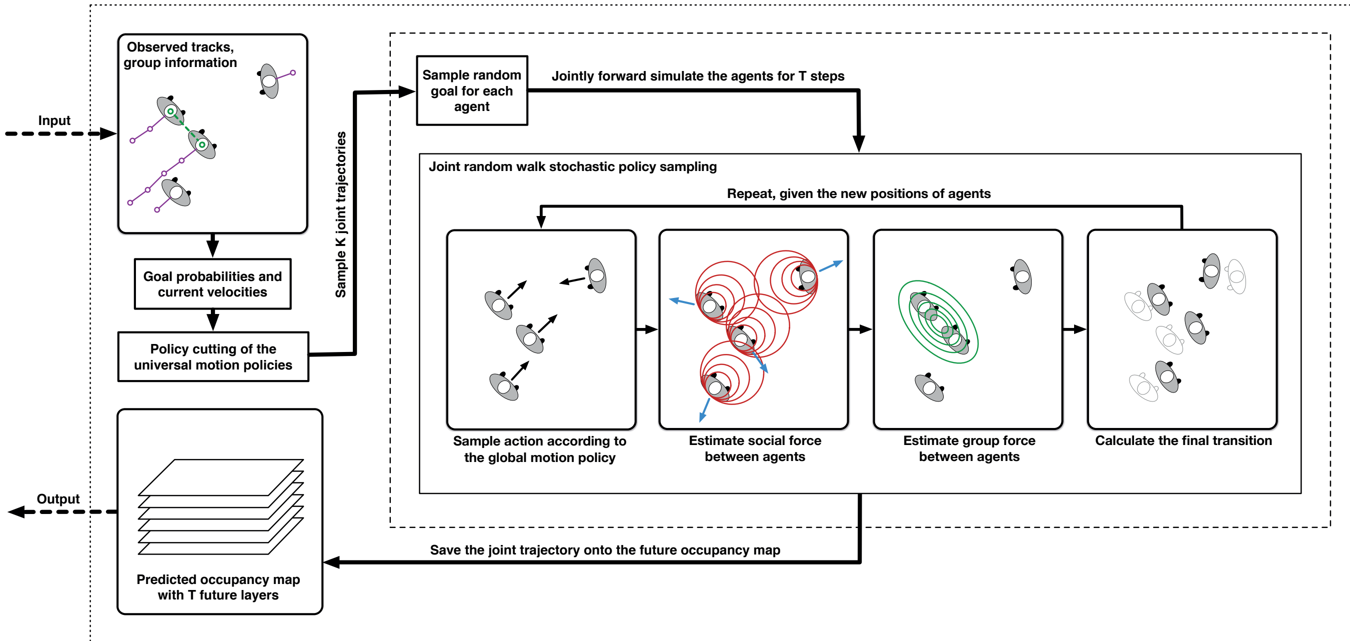

Fig. 3. Summary and flow diagram of the approach. Based on the observed tracklets and group detections, our prediction method samples K joint trajectories, each time drawing a random goal for the people and groups. The random walk stochastic policy sampling procedure estimates the social interactions and forward-simulates the agent positions T steps ahead. Joint trajectories are then saved onto the future occupancy map.

and get more precise predictions on average, we simply scale the observed speed v^i_{obs} of each human i by a factor $q_S > 1$.

The final direction of motion for person i is computed as

$$\mathbf{F}_i = \mathbf{F}_i^{\text{pers}} + \mathbf{F}_i^{\text{soc}} + \mathbf{F}_i^{\text{group}} = \mathbf{F}_i^{\text{pers}} + \sum_{k \neq i}^{N} \mathbf{f}_{i,k}^{\text{soc}} + \mathbf{f}_i^{\text{vis}} + \mathbf{f}_i^{\text{att}}. \quad (7)$$

An example of the social forces affecting the motion of people in a social scenario is given in Fig. 2, right.

2) Stochastic Policy Sampling Using Random Walks: To make predictions using the stochastic policy π_g, we utilize the random walk algorithm from our prior work [2] that samples K joint paths for all people in the scene. Each joint path is representing a possible future interaction given the observed tracklets and available group information. In each of the K samples we randomly draw a goal $g(i)$ for person i from the distribution $p^i(\mathcal{G})$ and randomly generate actions $a^i = (\theta^i, v^i)$ from the policy corresponding to $g(i)$. Group members share the same goal, sampled from $p^h_{Gr}(\mathcal{G})$. During the random walk, we evaluate the social interactions among the agents that affect each agent's instantaneous stochastic policy according to the group social force model. The position of each person at time t is then saved in the corresponding layer L^i_t of the probabilistic occupancy map L, that is shared among the K samples. Each layer L^i_t is normalized to represent the probability distribution of the person's location.

The inputs of our algorithm are the map M, goals $\mathcal{G}$, tracklets $\mathcal{T}$, groups $\mathbf{Gr}$ and the prediction horizon T. The algorithm has the following parameters: stochasticity level α, goal uncertainty β, human motion inertia coefficients I_v and I_θ, social force parameters $SF_p = (a_k, b_k, \lambda)$, group social force parameters $GSF_p = (\beta_1, \beta_2, q_A, \phi, q_S)$ and K joint trajectory samples. The summary of the approach is shown in Fig. 3. For more information on the algorithm's parameters and implementation details, see [2].

C. Complexity Analysis

Alg. 1 summarizes the operations required to obtain predictions with our algorithm. We assume that K joint random paths are requested, N people are in the scene and T prediction steps are made. The complexity of the goal sampling operation for every human (line 2) depends on the number of goals $|\mathcal{G}|$. Group center calculation is done only once for each time step (line 4). The random action sampling procedure (line 6) depends on the action space discretization (A angles and V velocities) and has the worst-case complexity of $O(AV)$. This happens when the agent is moving with velocity close to $v_{\max}$. The social force in the direction of agent i (line 7) is computed for each surrounding agent within a certain radius. In the worst-case, when all agents are densely located, the complexity is $O(N)$. The group social force computation (line 8) is a constant time operation.

The overall complexity of our prediction algorithm is then $O(K(N|\mathcal{G}| + T(N(AV + N))))$. Runtime measurements with comparison to the considered baselines are given in Sec. V.

IV. Experiments

In this section we present several experiments conducted to evaluate qualitatively and quantitatively our *Group Social Force MDP* (GSF-MDP) approach and compare its predictive capabilities with several baselines. All algorithms are implemented in C++ and running on a laptop with a 2.8

Algorithm 1 Joint Random Walk Stochastic Policy Sampling

```
for k = 1, ..., K do
    Sample a goal for each person: O(N|G|)
    for t = 1, ..., T do
        Calculate group center for each group: O(N)
        for i = 1, ..., N do
            Sample a random action: O(AV)
            Calculate social force: O(N)
            Calculate group social force: O(1)
```

GHz Xeon processor and 32 GB RAM. The action space of the MDP is discretized with $\pi/20$ increments of θ; 0.1 m/s increments of v, $v \in [0,3]$ m/s. Cell sizes of the grid maps are 0.05 m in Experiment 1 and 0.15 m in Experiment 2. The frequency of prediction is 4 Hz, the number of random walk samples $K = 200$.

A. Experiment 1: Predicting Social Interactions

This experiment includes several qualitative demonstrations of the predicted group collision avoidance behavior of people. To this end we use define maps of two environments and simulate observed trajectories in those maps to see the predicted development of interactive scenarios. The first scenario (Fig. 4) stages an experiment with 5 people in a narrow corridor. The second scenario (Fig. 5) sets up a challenging crowded environment with multiple non-convex obstacles and 21 people walking in 7 groups.

B. Experiment 2: Prediction Evaluation

Quantitative evaluation of GSF-MDP is conducted using the ATC dataset[1] recorded in a shopping center with 15 most common goals. We extract 21 social scenarios with trajectories of 172 people, including 90 pedestrians walking in groups, observed for long periods of time (see Fig. 1 for an example scenario). Static obstacles, motion stochasticity, observation noise and extensive social interaction involving many groups makes this dataset a challenging one, particularly for methods that do not model group motion. As a baseline for predictive performance evaluation, we compare GSF-MDP to a planning-based method by Karasev et al. [11] and the social force-based approach by Elfring et al. [4]. For the sake of a fair comparison, our own goal estimation technique, that requires no training data, is applied to both baselines. Finally, we include our previous *Joint Sampling MDP* (JS-MDP) method from [2] in the comparison to heuristically evaluate the benefit from considering group information.

We evaluate the predictions provided by all methods based on the NLP and MHD metrics. *Negative Log-Probability* (NLP) is a probabilistic measure, that computes the average predicted probability, measured at each point i of the ground truth path $\mathcal{T}$ for T steps into the future: $\text{NLP}(\mathcal{T}) = -\frac{1}{T}\sum_{i=1}^{T} \log p(\mathcal{T}_i|t_i)$. *Modified Hausdorff Distance* (MHD) is a geometric measure of distance between the ground truth path and the most probable path in the predicted probability distribution. For both metrics, lower values corresponds to better prediction accuracy or smaller geometric deviation, respectively. Metric values are calculated for each trajectory in the 21 interactive scenarios and averaged across 20 experiments for each scenario. We use 1.5 seconds as observation period, and predictions are obtained for $T = 2.5 - 12.5$ seconds ahead. We also measure the average time to compute predictions using our algorithm and the baselines.

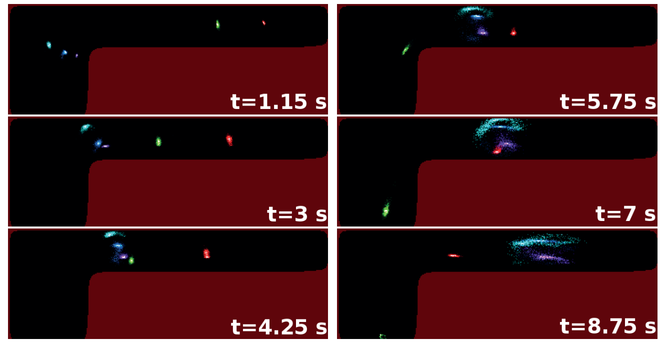

Fig. 4. Prediction results in a simulated scenario. Predicted distributions are color-coded. At $t = 1.15$ seconds a group of three people, depicted in **blue, cyan** and **purple** is walking upwards and then turns into the corridor to the right without losing its formation. At $t = 3$, $t = 4.25$ seconds the group is handling a hindrance caused by the **green** pedestrian, at $t = 5.75$, $t = 7$ seconds the group is handling the another hindrance with the **red** pedestrian.

Prior to the main experiment, we perform hyperparameter optimization using the SMAC3 optimization toolbox [25] for each algorithm. Optimization criteria is to minimize the sum of NLP and MHD values. The optimal parameters are found to be as follows: $\alpha = 4.64$, $\beta = 18.65$, $I = (0.09,\ 0.02)$, $(a_k, b_k, \lambda) = (0.09,\ 0.32,\ 0)$, $(\beta_1,\ \beta_2,\ q_A,\ \phi,\ q_S) = (0.05,\ 1.18,\ 2.93,\ 0.38,\ 1.49)$ for GSF-MDP; $\alpha = 13.26$, $\beta = 9.12$, $I = (0.01,\ 0.19)$, $(a_k, b_k, \lambda) = (1.46,\ 0.11,\ 0)$ for JS-MDP; $(w_{g,t}, w_{s,t}) = (0.03,\ 0.14)$, $\alpha = 21.31$, $\beta = 18.68$ for [11]; $(q_w, f_w, c_w) = (1.44,\ 0.23,\ 3.1)$, $\zeta_\rho = 83.74$ for [4].

V. Results

Fig. 4 and Fig. 5 show the results of Experiment 1. The first simulated scenario (Fig. 4) demonstrates a collision avoidance maneuver, performed by a group of three pedestrians in a narrow corridor. The group is able to keep its "social" linear walking formation that facilitates intra-group interaction. In the end, however, spreading of samples indicates the predicted possibility of re-grouping into a more compact V-formation – a behavioral pattern observed in real crowds [1]. In the second scenario (Fig. 5) our method predicts realistic behavior of group members. In particular, they are able to wait for the passage to clear before continuing their motion as a group, keeping the broad V-shape walking pattern when the available space allows it, and not lose its members behind in the dense crowd. Predicted results are visually compared with a baseline, where the group motion is not modeled.

Fig. 6 presents the quantitative results of Experiment 2, displaying the mean of the NLP and MHD metrics over the prediction horizon of 2.5–12.5 seconds. The NLP re-

[1] http://www.irc.atr.jp/crest2010_HRI/ATC_dataset/

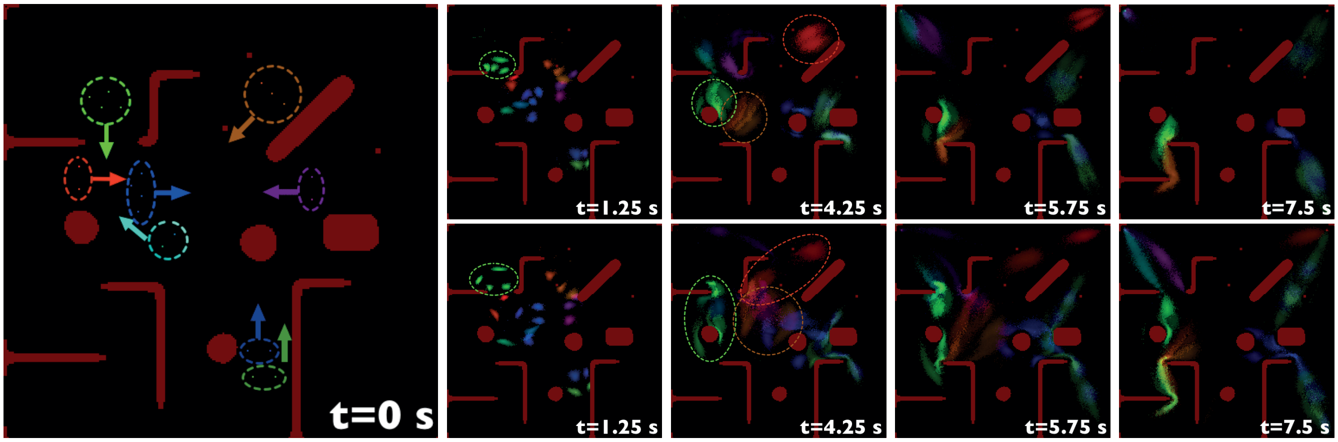

Fig. 5. Prediction results in a simulated scenario with obstacles and 21 people walking in 7 groups. Goals are placed in the four corners of the map. **Left:** initial positions of people are shown in colored circles, each color corresponding to one group. **Right, top row:** predicted positions with GSF-MDP for several points in time. Consider e.g. the green group that waits until the passage is cleared by the red and blue groups without losing its formation. Then it gives way for the faster orange group. People in the red group are correctly predicted to maintain a side-by-side walking formation. **Right, bottom row:** predicted positions with the JS-MDP baseline, where group motion is not modeled. The green group performs unnecessary maneuvers, then gets separated. The same happens with the red and orange groups, who lose their members in the crowd.

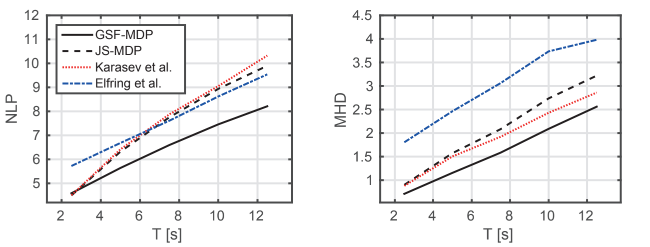

Fig. 6. **Left:** Mean of the Negative Log-Probability (NLP) metric in the ATC dataset. Our approach outperforms the baselines along the entire prediction horizon of up to 12.5 seconds. **Right:** Mean of the Modified Hausdorff Distance (MHD) metric. Our approach delivers more precise results on both short and long prediction horizons.

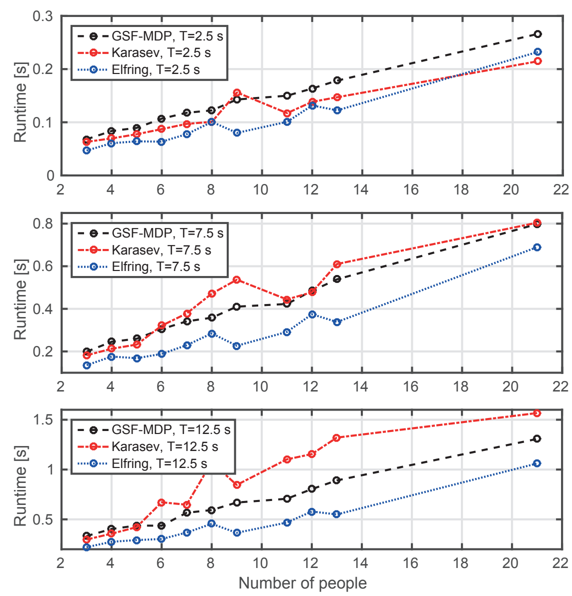

Fig. 7. Average runtime of our algorithm for prediction horizons $T = 2.5$, 7.5 and 12.5 seconds ahead on the ATC dataset with various numbers of people. On average, our method performs on par with the baselines.

sults suggest that our algorithm assigns higher probabilities to the ground truth states of the person's future location, outperforming all the baselines. The planning-based method of Karasev et al. [11] accumulates errors from non-predicted social interactions over the growing prediction horizon, while JS-MDP [2] suffers from the lack of the group awareness. The social force-based method of Elfring et al. [4] generates worse results due to the lack of global knowledge of the environment's structure. MHD evaluation results further confirm the improvement of our method over the state-of-the-art on both short and long-term prediction horizons.

In Fig. 7 we give the prediction runtime of GSF-MDP compared to the baselines. For example, our method is capable of computing 2.5 seconds of predictions for 5 people in less than 0.1 seconds, or predict 7.5 seconds of 10 people motion in 0.4 seconds. On average, our method performs on par with the state-of-the-art. Given that the range of the social forces is not large, and people are typically not agglomerated in a single region, the method most often scales linearly with the number of people, and not quadratically as in the worst-case, described in Section III-C.

A. Discussion

The evaluation results presented above are encouraging. Performing at similar runtime with the state-of-the-art, our method is capable of delivering more accurate predictions across the entire prediction horizon. Still, during our experiments we have encountered situations, generally challenging for long-term predictors, see e.g. Fig. 8. Our stochastic policy accounts for variations in paths and homotopy classes, but does not handle sudden velocity or motion intent changes – this limitation in a long-term setting is a common unexplored aspect in the literature. Predicting paths accurately in situations shown in Fig. 8 could be done with a dynamic

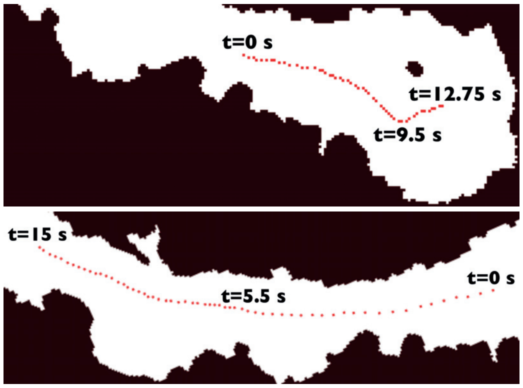

Fig. 8. Challenging motion trajectories in the ATC dataset. Pedestrian positions are measured at 4 Hz and plotted in **red**. We observe a change in motion intent (**top**, at $t = 9.5$ seconds) and motion velocity (**bottom**, at $t = 5.5$ seconds) not explainable by nearby people, group membership, environment geometry or other observable factors in the data.

α value, which increases uncertainty for more distant points in time. Learning relevant simuli for motion behavior in the environment and spatially incorporating them into the local behavior model could be another possibility to better foresee the uncertainty from sudden intention or velocity changes.

VI. Conclusions

In this paper, we presented a novel planning-based algorithm for predicting human motion that accounts for social grouping information. Our approach models the global long-term aspects using an MDP planner with stochastic policies, and the local aspects using a social-force based model to describe social interactions and group social forces. We use joint sampling of the individual global motion policies by a weighted random walk procedure in which each person is influenced by forces from other nearby agents and group members. Our approach outperforms alternative methods in terms of probabilistic and geometric measures on real data across the entire range of prediction horizons while being on par with respect to runtime performance.

In future work we plan an implementation on a real robot using a real tracking system. We are also interested in comparing several group motion models and study their impact on accuracy and efficiency metrics.

References

[1] M. Moussaïd, N. Perozo, S. Garnier, D. Helbing, and G. Theraulaz, "The walking behaviour of pedestrian social groups and its impact on crowd dynamics," *PloS one*, vol. 5, no. 4, 2010.

[2] A. Rudenko, L. Palmieri, and K. O. Arras, "Joint prediction of human motion using a planning-based social force approach," in *IEEE Int. Conf. on Robotics and Automation (ICRA)*, 2018.

[6] M. Luber, J. A. Stork, G. D. Tipaldi, and K. O. Arras, "People tracking with human motion predictions from social forces," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2010.

[3] K. Yamaguchi, A. C. Berg, L. E. Ortiz, and T. L. Berg, "Who are you with and where are you going?" in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2011.

[4] J. Elfring, R. Van De Molengraft, and M. Steinbuch, "Learning intentions for improved human motion prediction," *Robotics and Autonomous Systems*, vol. 62, no. 4, pp. 591–602, 2014.

[5] D. Helbing and P. Molnar, "Social force model for pedestrian dynamics," *Physical review E*, vol. 51, no. 5, p. 4282, 1995.

[7] M. Bennewitz, W. Burgard, G. Cielniak, and S. Thrun, "Learning motion patterns of people for compliant robot motion," *The International Journal of Robotics Research*, vol. 24, no. 1, 2005.

[8] P. Trautman and A. Krause, "Unfreezing the robot: Navigation in dense, interacting crowds," in *IEEE/RSJ Int. Conf. IROS*, Oct 2010.

[9] A. Alahi, K. Goel, V. Ramanathan, A. Robicquet, L. Fei-Fei, and S. Savarese, "Social LSTM: Human trajectory prediction in crowded spaces," in *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2016.

[10] A. Vemula, K. Mülling, and J. Oh, "Modeling cooperative navigation in dense human crowds," in *IEEE Int. Conf. ICRA*, 2017.

[11] V. Karasev, A. Ayvaci, B. Heisele, and S. Soatto, "Intent-aware long-term prediction of pedestrian motion," in *2016 IEEE International Conference on Robotics and Automation (ICRA)*, May 2016.

[12] B. D. Ziebart, N. Ratliff, G. Gallagher, C. Mertz, K. Peterson, J. A. Bagnell, M. Hebert, A. K. Dey, and S. Srinivasa, "Planning-based prediction for pedestrians," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, Piscataway, NJ, USA, 2009.

[13] L. Cheng, R. Yarlagadda, C. Fookes, and P. K. Yarlagadda, "A review of pedestrian group dynamics and methodologies in modelling pedestrian group behaviours," *World*, vol. 1, no. 1, pp. 002–013, 2014.

[14] I. Karamouzas and M. Overmars, "Simulating and evaluating the local behavior of small pedestrian groups," *IEEE Transactions on Visualization and Computer Graphics*, vol. 18, no. 3, 2012.

[15] F. Qiu and X. Hu, "Modeling group structures in pedestrian crowd simulation," *Simulation Modelling Practice and Theory*, 2010.

[16] H. Singh, R. Arter, L. Dodd, P. Langston, E. Lester, and J. Drury, "Modelling subgroup behaviour in crowd dynamics dem simulation," *Applied Mathematical Modelling*, vol. 33, no. 12, 2009.

[17] M. Seitz, G. Köster, and A. Pfaffinger, "Pedestrian group behavior in a cellular automaton," in *Pedestrian and Evacuation Dynamics*, 2012.

[18] S. Bandini, F. Rubagotti, G. Vizzari, and K. Shimura, "An agent model of pedestrian and group dynamics: experiments on group cohesion," *AI* IA 2011: Artificial Intelligence Around Man and Beyond*, 2011.

[19] J. Šochman and D. C. Hogg, "Who knows who-inverting the social force model for finding groups," in *IEEE International Conference on Computer Vision Workshops (ICCV Workshops)*, 2011.

[20] B. Lau, K. O. Arras, and W. Burgard, "Multi-model hypothesis group tracking and group size estimation," *International Journal of Social Robotics*, vol. 2, no. 1, March 2010.

[21] S. Pellegrini, A. Ess, and L. Van Gool, "Improving data association by joint modeling of pedestrian trajectories and groupings," in *European Conference on Computer Vision*. Springer, 2010, pp. 452–465.

[22] M. Luber and K. O. Arras, "Multi-hypothesis social grouping and tracking for mobile robots," in *Robotics: Science and Systems (RSS'13)*, Berlin, Germany, 2013.

[23] T. Linder, S. Breuers, B. Leibe, and K. O. Arras, "On multi-modal people tracking from mobile platforms in very crowded and dynamic environments," in *IEEE Int. Conf. on Robotics and Automation*, 2016.

[24] D. Vasquez, "Novel planning-based algorithms for human motion prediction," in *2016 IEEE Int. Conf. ICRA*, May 2016.

[25] M. Lindauer, K. Eggensperger, M. Feurer, S. Falkner, A. Biedenkapp, and F. Hutter, "Smac v3: Algorithm configuration in python," https://github.com/automl/SMAC3, 2017.