

Intention-Aware Motion Planning Using Learning Based Human Motion Prediction

Jae Sung Park and Chonhyon Park and Dinesh Manocha
Department of Computer Science, UNC Chapel Hill, NC, USA
{jaesungp, chpark, dm}@cs.unc.edu
<http://gamma.cs.unc.edu/SafeMP> (video included)

Abstract—We present a motion planning algorithm to compute collision-free and smooth trajectories for high-DOF robots interacting with humans in a shared workspace. Our approach uses offline learning of human actions along with temporal coherence to predict the human actions. Our intention-aware online planning algorithm uses the learned database to compute a reliable trajectory based on the predicted actions. We represent the predicted human motion using a Gaussian distribution and compute tight upper bounds on collision probabilities for safe motion planning. We highlight the performance of our planning algorithm in complex simulated scenarios and real world benchmarks with 7-DOF robot arms operating in a workspace with a human performing complex tasks. We demonstrate the benefits of our intention-aware planner in terms of computing safe trajectories in such uncertain environments.

I. INTRODUCTION

Motion planning algorithms are used to compute collision-free paths for robots among obstacles. As robots are increasingly used in workspace with moving or unknown obstacles, it is important to develop reliable planning algorithms that can handle environmental uncertainty and the dynamic motions. In particular, we address the problem of planning safe and reliable motions for a robot that is working in environments with humans. As the humans move, it is important for the robots to predict the human actions and motions from sensor data and to compute appropriate trajectories.

In order to compute reliable motion trajectories in such shared environments, it is important to gather the state of the humans as well as predict their motions. There is considerable work on online tracking and prediction of human motion in computer vision and related areas [32]. However, the current state of the art in gathering motion data results in many challenges. First of all, there are errors in the data due to the sensors (e.g., point cloud sensors) or poor sampling [4]. Secondly, human motion can be sudden or abrupt and this can result in various uncertainties in terms of accurate representation of the environment. One way to overcome some of these problems is to use predictive or estimation techniques for human motion or actions, such as using various filters like Kalman filters or particle filters [40]. Most of these prediction algorithms use a motion model that can predict future motion based on the prior positions of human body parts or joints, and corrects the error between the estimates and actual measurements. In practice, these approaches only work well when there is sufficient information about prior motion that can be accurately modeled by the underlying




Fig. 1: A 7-DOF Fetch robot is moving its arm near a human, avoiding collisions. (a) While the robot is moving, the human tries to move his arm to block the robot’s path. The robot arm trajectory is planned without human motion prediction, which may result in collisions and a jerky trajectory, as shown with the red circle. This is because the robot cannot respond to the human motion to avoid collisions. (b) The trajectory is computed using our human motion prediction algorithm; it avoids collisions and results in smoother trajectories. The robot trajectory computation uses collision probabilities to anticipate the motion and compute safe trajectories.

motion model. In some scenarios, it is possible to infer high-level human intent using additional information, and thereby perform a better prediction of future human motions [1, 2]. These techniques are used to predict the pedestrian trajectories based on environmental information in 2D domains.

Main Results: We present a novel high-DOF motion planning approach to compute collision-free trajectories for robots operating in a workspace with human-obstacles or human-robot cooperating scenarios. Our approach is general, and doesn’t make much assumptions about the environment or the human actions. We track the positions of the human using depth cameras and present a new method for human action prediction using combination of classification and regression methods. Given the sensor noises and prediction errors, our online motion planner uses probabilistic collision checking to compute a high dimensional robot trajectory that tends to compute safe motion in the presence of uncertain human motion. As compared to prior methods, the main benefits of our approach include:

- 1) A novel data-driven algorithm for intention and motion prediction, given noisy point cloud data. Compared to prior methods, our formulation can account for big noise in skeleton tracking in terms of human motion prediction.
- 2) An online high-DOF robot motion planner for efficient completion of collaborative human-robot tasks that uses upper bounds on collision probabilities to compute safe trajectories in challenging 3D workspaces. Furthermore, our trajectory optimization based on probabilistic collision checking results in smoother paths.

We highlight the performance of our algorithms in a simulator with a 7-DOF KUKA arm operating and in a real world setting with a 7-DOF Fetch robot arm in a workspace with a moving human performing cooperative tasks. We have evaluated its performance in some challenging or cluttered 3D environments where the human is close to the robot and moving at varying speeds. We demonstrate the benefits of our intention-aware planner in terms of computing safe trajectories in these scenarios.

The rest of paper is organized as follows. In Section II, we give a brief survey of prior work. Section III presents an overview of our human intention-aware motion planning algorithm. Offline learning and runtime prediction of human motions are described in Section IV, and these are combined with our optimization based motion planning algorithm in Section V. The performance of the intention-aware motion planner is analyzed in Section VI. Finally, we demonstrate the performance of our planning framework for a 7-DOF robot in Section VII.

II. RELATED WORKS

In this section, we give a brief overview of prior work on human motion prediction, task planning for human-robot collaborations, and motion planning in environments shared with humans.

A. Intention-Aware Motion Planning and Prediction

Intention-Aware Motion Planning (IAMP) denotes a motion planning framework where the uncertainty of human intention is taken into account [1]. The goal position and the trajectory of moving pedestrians can be considered as human intention and used so that a moving robot can avoid pedestrians [37].

In terms of robot navigation among obstacles and pedestrians, accurate predictions of humans or other robot positions are possible based on crowd motion models [11, 39] or integration of motion models with online learning techniques [15] for 2D scenarios and they are orthogonal to our approach.

Predicting the human actions or the high-DOF human motions has several challenges. Estimated human poses from recorded videos or realtime sensor data tend to be inaccurate or imperfect due to occlusions or limited sensor ranges [4]. Furthermore, the whole-body motions and their complex dynamics with many high-DOF makes it difficult to represent them with accurate motion models [14]. There has been a considerable literature on recognizing human actions [36]. Machine learning-based algorithms using Gaussian Process

Latent Variable Models (GP-LVM) [8, 38] or Recurrent neural network (RNN) [10] have been proposed to compute accurate human dynamics models. Recent approaches use the learning of human intentions along with additional information, such as temporal relations between the actions [22, 13] or object affordances [16] to improve the accuracy. Inverse Reinforcement Learning (IRL) has been used to predict 2D motions [42, 18] or 3D human motions [5].

B. Robot Task Planning for Human-Robot Collaboration

In human-robot collaborative scenarios, robot task planning algorithms have been developed for the efficient distribution of subtasks. One of their main goal is reducing the completion time of the overall task by interleaving subtasks of robot with subtasks of humans with well designed task plans. In order to compute the best action policy for a robot, Markov Decision Processes (MDP) have been widely used [3]. Nikolaidis et al. [22] use MDP models based on mental model convergence of human and robots. Koppula and Saxena [17] use Q-learning to train MDP models where the graph model has the transitions corresponding to the human action and robot action pairs. Pérez-D’Arpino and Shah [28] used a Bayesian learning algorithm on hand motion prediction and tested the algorithm in a human-robot collaboration tasks. Our MDP models extend these approaches, but also take into account the issue of avoiding collisions between the human and the robot.

C. Motion Planning in Environments shared with Humans

Prior work on motion planning in the context of human-robot interaction has focused on computing robot motions that satisfy cognitive constraints such as social acceptability [33] or being legible to humans [6].

In human-robot collaboration scenarios where both the human and the robot perform manipulation tasks in a shared environment, it is important to compute robot motions that avoid collisions with the humans for safety reasons. Dynamic window approach [9] (which searches the optimal velocity in a short time interval) and online replanning [29, 24, 25] (which interleaves planning with execution) are widely used approaches for planning in such dynamic environments. As there are uncertainties in the prediction model and in the sensors for human motion, the future pose is typically represented as the *Belief* state, which corresponds to the probability distribution over all possible human states. Mainprice and Berenson [20] explicitly construct an occupied workspace voxel map from the predicted Belief states of humans in the shared environment and avoid collisions.

III. OVERVIEW

In this section, we first introduce the notation and terminology used in the paper and give an overview of our motion planning algorithms.

A. Notation and Assumptions

As we need to learn about human actions and short-term motions, a large training dataset of human motions is needed. We collect N demonstrations of how human joints typically move while performing some tasks and in which order sub-tasks are performed. Each demonstration is represented using $T^{(i)}$ time frame of human joint motion, where the superscript (i) represents the demonstration index. The motion training dataset is represented as following:

- ξ is a matrix of tracked human joint motions. $\xi^{(i)}$ has $T^{(i)}$ columns, where a column vector represents the different human joint positions during each time frame.
- F is a feature vector matrix. $F^{(i)}$ has $T^{(i)}$ columns and is computed from $\xi^{(i)}$.
- \mathbf{a}^h is a human action (or subtask) sequence vector that represents the action labels over different time frames. For each time frame, the action is categorized into one of the m^h discrete action labels, where the action label set is $A^h = \{a_1^h, \dots, a_{m^h}^h\}$.
- $L = \{(\xi^{(i)}, F^{(i)}, \mathbf{a}^{h,(i)})\}_{i=1}^N$ is the motion database used for training. It consists of human joint motions, feature descriptors and the action labels at each time frame.

During runtime, MDP-based human action inference is used in our task planner. The MDP model is defined as a tuple (P, A^r, T) :

- $A^r = \{a_1^r, \dots, a_{m^r}^r\}$ is a robot action (or subtask) set of m^r discrete action labels.
- $P = \mathcal{P}(A^r \cup A^h)$, a power set of union of A^r and A^h , is the set of states in MDP. We refer to the state \mathbf{p} as a *progress state* because each state represents which human and robot actions have been performed so far. We assume that the sequence of future actions for completing the entire task depends on the list of actions completed. \mathbf{p} has $m^h + m^r$ binary elements, which represent corresponding human or robot actions have been completed ($\mathbf{p}_j = 1$) or not ($\mathbf{p}_j = 0$). For cases where same actions can be done more than once, the binary values can be replaced with integers, to count the number of actions performed.
- $T : P \times A^r \rightarrow \Pi(P)$ is a transition function. When a robot performs an action a^r in a state \mathbf{p} , $T(\mathbf{p}, a^r)$ is a probability distribution over the progress state set P . The probability of being state \mathbf{p}' after taking an action a^r from state \mathbf{p} is denoted as $T(\mathbf{p}, a^r, \mathbf{p}')$.
- $\pi : P \rightarrow A^r$ is the action policy of a robot. $\pi(\mathbf{p})$ denotes the best robot action that can be taken at state \mathbf{p} , which results in maximal performance.

We use the Q-learning [35] to determine the best action policy during a given state, which rewards the values that are induced from the result of the execution.

B. Robot Representation

We denote a single configuration of the robot as a vector \mathbf{q} that consists of joint-angles. The n -dimensional space of configuration \mathbf{q} is the configuration space \mathcal{C} . We represent each link of the robot as R_i . The finite set of bounding spheres for

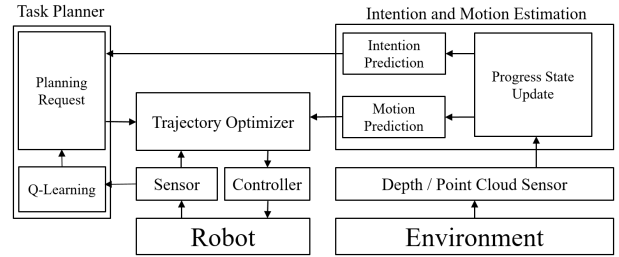


Fig. 2: **Overview of our Intention-Aware Planner:** Our approach consists of three main components: task planner, trajectory optimization, and intention and motion estimation.

link R_i is $\{B_{i1}, B_{i2}, \dots\}$, and is used as a bounding volume of the link, i.e., $R_i \subset \cup_j B_{ij}$. The links and bounding spheres at a configuration \mathbf{q} are denoted as $R_i(\mathbf{q})$ and $B_{ij}(\mathbf{q})$, respectively. In our benchmarks, where the robot arms are used, these bounding spheres are automatically generated using the medial axis of robot links. We also generate the bounding spheres $\{C_1, C_2, \dots\}$ for humans and other obstacles.

For a planning task with start and goal configurations \mathbf{q}_s and \mathbf{q}_g , the robot's trajectory is represented by a matrix \mathbf{Q} ,

$$\mathbf{Q} = \begin{bmatrix} \mathbf{q}_s & \mathbf{q}_1 & \dots & \mathbf{q}_{n-1} & \mathbf{q}_g \\ t_0 & t_1 & \dots & t_{n-1} & t_n \end{bmatrix},$$

where robot trajectory passes through the $n+1$ waypoints. We denote the i -th waypoint of \mathbf{Q} as $\mathbf{x}_i = [\mathbf{q}_i^T \ t_i]$.

C. Online Motion Planning

The main goals of our motion planner are: (1) planning high-level tasks for a robot by anticipating the most likely next human action and (2) computing a robot trajectory that reduces the probability of collision between the robot and the human or other obstacles, by using motion prediction.

At the high-level task planning step, we use MDP, which is used to compute the best action policies for each state. A state of an MDP graph denotes the progress of the whole task. The best action policies are determined through reinforcement learning with Q-learning. Then, the best action policies are updated within the same state. The probability of choosing the action increases or decreases according to the reward function. Our reward computation function is affected by the prediction of intention and the delay caused by predictive collision avoidance.

We also estimate the short-term future motion from learned information in order to avoid future collisions. From the joint position information, motion features are extracted based on human poses and surrounding objects related to human-robot interaction tasks, such as joint positions of humans, relative positions from a hand to other objects, etc. The motions are classified over the human action set A^h . For classifying the motions, we use Import Vector Machine (IVM) [41] for classification and a Dynamic Time Warping (DTW) [21] kernel function for incorporating the temporal information. Given the human motions database of each action type, we train future motions using Sparse Pseudo-input Gaussian Process (SPGP) [34]. Combining these two prediction results, the

final future motion is computed as the weighted sum over different action types weighed by the probability of each action type that could be performed next. For example, if the action classification results in probability 0.9 for action *Move forward* and 0.1 for action *Move backward*, the future motion prediction (the results of SPGP) for *Move forward* dominates. If the action classification results in probability 0.5 for both actions, the predicted future motions for each action class will be used in avoiding collisions but with weights 0.5. However, in this case, the current motion does not have specific features to distinguish the action, meaning that the future motion in a short term will be similar and there will be an overlapped region in 3D space, working as a future motion of weight 1. More details are described in Section IV.

After deciding which robot task will be performed, the robot motion trajectory is then computed that tends to avoid collisions with humans. An optimization-based motion planner [24] is used to compute a locally optimal solution that minimizes the objective function subject to many types of constraints such as robot related constraints (e.g., kinematic constraint), human motion related constraints (e.g., collision free constraint), etc. Because future human motion is uncertain, we can only estimate the probability distribution of the possible future motions. Therefore, we perform probabilistic collision checking to reduce the collision probability in future motions. We also continuously track the human pose and update the predicted future motion to re-plan safe robot motions. Our approach uses the notion of online probabilistic collision detection [23, 26, 27] between the robot and the point-cloud data corresponding to human obstacles, to compute reactive costs and integrate them with our optimization-based planner.

IV. HUMAN ACTION PREDICTION

In this section, we describe our human action prediction algorithm, which consists of offline learning and online inference of actions.

A. Learning of Human Actions and Temporal Coherence

We collect N demonstrations to form a motion database L . The 3D joint positions are tracked using OpenNI library [30], and their coordinates are concatenated to form a column vector $\xi^{(i)}$. For full-body motion prediction, we used 21 joints, each of which has 3D coordinates tracked by OpenNI. So, $\xi^{(i)}$ is a 63-dimensional vector. For upper-body motion prediction, there are 10 joints and thus $\xi^{(i)}$ is length 30. Then, feature vector $F^{(i)}$ is derived from $\xi^{(i)}$. It has joint velocities and joint accelerations, as well as joint positions.

To learn the temporal coherence between the actions, we deal with only the human action sequences $\{\mathbf{a}^{h,(i)}\}_{i=1}^N$. Based on the progress state representation, for any time frame s , the prefix sequence of $\mathbf{a}^{h,(i)}$ of length s yields a progress state $\mathbf{p}_s^{(i)}$ and the current action $c_s^{h,(i)} = \mathbf{a}_s^{h,(i)}$. The next action label $n_s^{h,(i)}$ performed after frame s can also be computed, at which the action label differs at the first time while searching in the increasing order from frame $s + 1$. Then, for all possible pairs of demonstrations and frame index (i, s) , the




Fig. 3: **Motion uncertainty and prediction:** (a) A point cloud and the tracked human (blue spheres). The joint positions are used as feature vectors. (b) Prediction of next human action and future human motion, where 4 locations are colored according to their probability of next human action from white (0%) to black (100%). Prediction of future motion after 1 second (red spheres) from current motion (blue spheres) is shown as performing the action: move right hand to the second position which has the highest probability associated with it.

tuples $(\mathbf{p}_s^{(i)}, c_s^{h,(i)}, n_s^{h,(i)})$ are collected to compute histograms $h(n^h; \mathbf{p}, c^h)$, which counts the next action labels at each pair (\mathbf{p}, c^h) that have appeared at least once. We use the normalized histograms to estimate the next future action for the given \mathbf{p} and c^h . i.e.,

$$p(n^h = a_j^h | \mathbf{p}, c^h) = \frac{h(a_j^h; \mathbf{p}, c^h)}{\sum_{k=1}^m h(a_k^h; \mathbf{p}, c^h)}. \quad (1)$$

In order to train the human motion, the motion sequence $\xi^{(i)}$ and the feature sequence $F^{(i)}$ are learned, as well as the action sequences $\mathbf{a}^{(i)}$. Because we are interested in short-term human motion prediction for collision avoidance, we train the learning module from multiple short periods of motion. Let n_p be the number of previous consecutive frames and n_f be the number of future consecutive frames to look up. n_f and n_p are decided so that the length of motion is short-term motion (about 1 second) that will be used for short-term future collision detection in the robot motion planner. At the time frame s , where $n_p \leq s \leq T^{(i)} - n_f$, the columns of feature matrix $F^{(i)}$ from column index $s - n_p + 1$ to s are denoted as $F_{prev,s}^{(i)}$. Similarly, the columns of motion matrix from index $s + 1$ to $s + n_f$ are denoted as $\xi_{next,s}^{(i)}$.

Tuples $(F_{prev,s}^{(i)}, \mathbf{p}_s^{(i)}, c_s^{h,(i)}, \xi_{next,s}^{(i)})$ for all possible pairs of (i, s) are collected as the training input. They are partitioned into groups having the same progress state \mathbf{p} . For each progress state \mathbf{p} and current action c^h , the set of short-term motions are regressed using SPGP with the DTW kernel function [21], considering $\{F_{prev}\}$ as input and $\{\xi_{next}\}$ as multiple channeled outputs. We use SPGPs, a variant of Gaussian Processes, because it significantly reduces the running time for training and inference by choosing M pseudo-inputs from a large number of an original human motion inputs. The final

learned probability distribution is

$$p(\xi_{next}|F_{prev}, \mathbf{p}, c^h) = \prod_{c:\text{channels}} p(\xi_{next,c}|F_{prev}, \mathbf{p}, c^h),$$

$$p(\xi_{next,c}|F_{prev}, \mathbf{p}, c^h) \sim \mathcal{GP}(m_c, K_c), \quad (2)$$

where $\mathcal{GP}(\cdot, \cdot)$ represents trained SPGPs, c is an output channel (i.e., an element of matrix ξ_{next}), and m_c and K_c are the learned mean and covariance functions of the output channel c , respectively.

The current action label $c_s^{h,(i)}$ should be learned to estimate the current action. We train $c_s^{h,(i)}$ using Tuples $(F_{prev,s}^{(i)}, \mathbf{p}_s^{(i)}, c_s^{h,(i)})$. For each state \mathbf{p} , we use Import Vector Machine (IVM) classifiers to compute the probability distribution:

$$p(c^h = a_j^h | F_{prev}, \mathbf{p}) = \frac{\exp(f_j(F_{prev}))}{\sum_{a_k^h} \exp(f_k(F_{prev}))}, \quad (3)$$

where $f_j(\cdot)$ is the learned predictive function [41] of IVM.

B. Runtime Human Intention and Motion Inference

Based on the learned human actions, at runtime we infer the next most likely short-term human motion and human subtask for the purpose of collision avoidance and task planning, respectively. The short-term future motion prediction is used for the collision avoidance during the motion planning. The probability of future motion is given as:

$$p(\xi_{next}|F, \mathbf{p}) = \sum_{c^h \in A^h} p(\xi_{next}, c^h | F, \mathbf{p}).$$

By applying the Bayes theorem, we get

$$p(\xi_{next}|F, \mathbf{p}) = \sum_{c^h \in A^h} p(c^h | F, \mathbf{p}) p(\xi_{next} | F, \mathbf{p}, c^h). \quad (4)$$

The first term $p(c^h | F, \mathbf{p})$ is inferred through the IVM classifier in (3). To infer the second term, we use the probability distribution in (2) for each output channel.

We use Q-learning for training the best robot action policy π in our MDP-based task planner. We first define the function $Q: P \times A^r \rightarrow \mathbb{R}$, which is iteratively trained with the motion planning executions. Q is updated as

$$Q_{t+1}(\mathbf{p}_t, a_t^r) = (1 - \alpha_t) Q_t(\mathbf{p}_t, a_t^r) + \alpha_t (R_{t+1} + \gamma \max_{a^r} Q_t(\mathbf{p}_{t+1}, a^r)),$$

where the subscripts t and $t+1$ are the iteration indexes, R_{t+1} is the reward function after taking action a_t^r at state \mathbf{p}_t , and α_t is the learning rate, where we set $\alpha_t = 0.1$ in our experiments. A reward value R_{t+1} is determined by several factors:

- Preparation for next human action: the reward gets higher when the robot runs an action before a human action which can be benefited by the robot's action. We define this reward as $R_{prep}(\mathbf{p}_t, a_t^r)$. Because the next human subtask depends on the uncertain human decision, we predict the likelihood of the next subtask from the learned

actions in (1) and use it for the reward computation. The reward value is given as

$$R_{prep}(\mathbf{p}_t, a_t^r) = \sum_{a^h \in A^h} p(n^h = a^h | \mathbf{p}_t) H(a^h, a_t^r), \quad (5)$$

where $H(a^h, a^r)$ is a prior knowledge of reward, representing the amount of how much the robot helped the human by performing the robot action a^r before the human action a^h . If the robot action a^r has no relationship with a^h , the H value is zero. If the robot helped, the H value is positive, otherwise negative.

- Execution delay: There may be a delay in the robot motion's execution due to the collision avoidance with the human. To avoid collisions, the robot may deviate around the human and make it without delay. In this case the reward function is not affected, i.e. $R_{delay,t} = 0$. However, there are cases that the robot must wait until the human moves to another pose because the human can block the robot's path, which causes delay d . We penalize the amount of delay to the reward function, i.e. $R_{delay,t} = -d$. Note that the delay can vary during each iteration due to the human motion uncertainty.

The total reward value is a weighted sum of both factors:

$$R_{t+1} = w_{prep} R_{prep}(\mathbf{p}_t, a_t^r) + w_{delay} R_{delay,t}(\mathbf{p}_t, a_t^r),$$

where w_{prep} and w_{delay} are weights for scaling the two factors. The preparation reward value is predefined for each action pairs. The delay reward is measured during runtime.

V. INTENTION-AWARE MOTION PLANNING

Our motion planner is based on an optimization formulation, where $n+1$ waypoints in the space-time domain \mathbf{Q} define a robot motion trajectory to be optimized. Specifically, we use an optimization-based planner, ITOMP [24], that repeatedly refines the trajectory while interleaving the execution and motion planning for dynamic scenes. We handle three types of constraints: smoothness constraint, static obstacle collision-avoidance, and dynamic obstacle collision avoidance. To deal with the uncertainty of future human motion, we use probabilistic collision detection between the robot and the predicted future human pose.

Let s be the current waypoint index, meaning that the motion trajectory is executed in the time interval $[t_0, t_s]$, and let m be the replanning time step. A cost function for collisions between the human and the robot can be given as:

$$\sum_{i=s+m}^{s+2m} p \left(\bigcup_{j,k} B_{jk}(\mathbf{q}_i) \cap C_{dyn}(t_i) \neq \emptyset \right) \quad (6)$$

where $C_{dyn}(t)$ are the workspace volumes occupied by dynamic human obstacles at time t . The trajectory being optimized during the time interval $[t_s, t_{s+m}]$ is executed during the next time interval $[t_{s+m}, t_{s+2m}]$. Therefore, the future human poses are considered only in the time interval $[t_{s+m}, t_{s+2m}]$.

The collision probability between the robot and the dynamic obstacle at time frame i in (6) can be computed as a maximum between bounding spheres:

$$\max_{j,k,l} p(B_{jk}(\mathbf{q}_i) \cap C_l(t_i) \neq \emptyset), \quad (7)$$

where $C_l(t_i)$ denotes bounding spheres for a human body at time t_i whose centers are located at line segments between human joints. The future human poses ξ_{future} are predicted in (4) and the bounding sphere locations $C_l(t_i)$ are derived from it. Note that the probabilistic distribution of each element in ξ_{future} is a linear combination of current action proposal $p(c^h|F, \mathbf{p})$ and Gaussians $p(\xi_{future}|F, \mathbf{p}, c^h)$ over all c^h , i.e., (7) can be reformulated as

$$\max_{j,k,l} \sum_{c^h} p(c^h|F, \mathbf{p}) p(B_{jk}(\mathbf{q}_i) \cap C_l(t_i) \neq \emptyset).$$

Let \mathbf{z}_l^1 and \mathbf{z}_l^2 be the probability distribution functions of two adjacent human joints obtained from $\xi_{future}(t_i)$, where the center of $C_l(t_i)$ is located between them by a linear interpolation $C_l(t_i) = (1-u)\mathbf{z}_l^1 + u\mathbf{z}_l^2$ where $0 \leq u \leq 1$. The joint positions follows Gaussian probability distributions:

$$\mathbf{z}_l^i \sim \mathcal{N}(\mu_l^i, \Sigma_l^i)$$

$$\mathbf{c}_l(t_i) \sim \mathcal{N}((1-u)\mu_l^1 + u\mu_l^2, (1-u)^2\Sigma_l^1 + u^2\Sigma_l^2) \quad (8)$$

$$= \mathcal{N}(\mu_l, \Sigma_l), \quad (9)$$

where $\mathbf{c}_l(t_i)$ is the center of $C_l(t_i)$. Thus, the collision probability between two bounding spheres is bounded by

$$\int_{\mathbb{R}^3} I(\|\mathbf{x} - \mathbf{b}_{jk}(\mathbf{q}_i)\|^2 \leq (r_1 + r_2)^2) f(\mathbf{x}) d\mathbf{x}, \quad (10)$$

where $\mathbf{b}_{jk}(\mathbf{q}_i)$ is the center of bounding sphere $B_{jk}(\mathbf{q}_i)$, r_1 and r_2 are the radius of $B_{jk}(\mathbf{q}_i)$ and $C_l(t_i)$, respectively, $I(\cdot)$ is an indicator function, and $f(\mathbf{x})$ is the probability distribution function. The indicator function restricts the integral domain to a solid sphere, and $f(\mathbf{x})$ is the probability density function of $\mathbf{c}_l(t_i)$, in (9). There is no closed form solution for (10), therefore we use the maximum possible value to approximate the probability. We compute \mathbf{x}_{max} at which $f(\mathbf{x})$ is maximized in the sphere domain and multiply it by the volume of sphere, i.e.

$$p(B_{jk}(\mathbf{q}_i) \cap C_l(t_i) \neq \emptyset) \leq \frac{4}{3}\pi(r_1 + r_2)^2 f(\mathbf{x}_{max}). \quad (11)$$

Since even \mathbf{x}_{max} does not have a closed form solution, we use the bisection method to find λ with

$$\mathbf{x}_{max} = (\Sigma^{-1} + \lambda I)^{-1}(\Sigma^{-1}\mathbf{p}_{lm} + \lambda \mathbf{o}_{jk}(\mathbf{q}_i)),$$

which is on the surface of sphere, explained in Generalized Tikhonov regularization [12] in detail.

The collision probability, computed in (10), is always positive due to the uncertainty of the future human position, and we compute a trajectory that is guaranteed to be nearly collision-free with sufficiently low collision probability. For a user-specified confidence level δ_{CL} , we compute a trajectory that its probability of collision is upper-bounded by $(1 - \delta_{CL})$.

If it is unable to compute a collision-free trajectory, a new waypoint \mathbf{q}_{new} is appended next to the last column of \mathbf{Q} to make the robot wait at the last collision-free pose until it finds a collision-free trajectory. This approach computes a guaranteed collision-free trajectory, but leads to delay, which is fed to the Q-learning algorithm for the MDP task planner. The higher the delay that the collision-free trajectory of a task has, the less likely the task planner selects the task again.

VI. ANALYSIS

The overall performance is governed by three factors: the predicted human motions described in Section IV; The optimization-based motion planner described in Section V; and the collision probability between the robot and predicted human motions for safe trajectory computation. In this section, we analyze the performance and accuracy of each factor.

A. Upper bound of Collision Probability

Using the predicted distribution and user-specified threshold δ_{CL} , we can compute an upper bound using the following lemma.

Lemma 1. *The collision probability is less than $(1 - \delta_{CL})$ if $\frac{4}{3}\pi(r_1 + r_2)^2 f(\mathbf{x}_{max}) < 1 - \delta_{CL}$.*

Proof: This bounds follows Equation (10) and (11).

$$p(B_{jk}(\mathbf{q}_i) \cap C_l(t_i) \neq \emptyset) \leq \frac{4}{3}\pi(r_1 + r_2)^2 f(\mathbf{x}_{max}) < 1 - \delta_{CL}. \quad \blacksquare$$

We use this bound in our optimization algorithm for collision avoidance.

B. Safe Trajectory Optimization

Our goal is to compute a robot trajectory that will either not collide with the human or reduces the probability of collision below a certain threshold. Sometimes, there is no feasible collision-free trajectory that the robot cannot avoid with the human. However, if there is any trajectory where the collision probability is less than a threshold, we seek to compute such a trajectory. Our optimization-based planner also generates multiple initial trajectories and finds the best solution in parallel. In this manner, it expands the search space and reduces the probability of the robot being stuck in a local minima configuration. This can be expressed as the following theorem:

Theorem 2. *An optimization-based planner with n parallel threads will compute the a global solution trajectory with collision probability less than $(1 - \delta_{CL})$, with the probability $1 - (1 - \frac{|A(\delta_{CL})|}{|S|})^n$, if it exists, where S is the entire search space. $A(\delta_{CL})$ corresponds to the neighborhood around the optimal solutions, with collision probability being less than $(1 - \delta_{CL})$, where the local optimization converges to one of the global optima. $|\cdot|$ is the measure of the search or configuration space.*

We give a brief overview of the proof. In this case, $\frac{|A(\delta_{CL})|}{|S|}$ measures the probability that a random sample lies in the

neighborhood of the global optima. In general, $|A(\delta_{CL})|$ will be smaller as the environment becomes more complex and has more local minima. Overall, this theorem provides a lower bound on the probability that our intention-aware planner with n threads. In the limit that n increases, the planner will compute the optimal solution if it exist. This can be stated as:

Corollary 2.1 (Probabilistic Completeness). *Our intention-aware motion planning algorithm with n trajectories is probabilistic complete, as n increases.*

Proof:

$$\lim_{n \rightarrow \infty} 1 - \left(1 - \frac{|A(\delta_{CL})|}{|S|}\right)^n = 1. \quad \blacksquare$$

VII. IMPLEMENTATION AND PERFORMANCE

We highlight the performance of our algorithm in a situation where the robot is performing a collaborative task with a human and computing safe trajectories. We use a 7-DOF KUKA-IIWA robot arm. The human motion is captured by a Kinect sensor operating with a 15Hz frame rate, and only upper body joints are tracked for collision checking. We use the ROS software [31] for robot control and sensor data communication. The motion planner has a 0.5s re-planning timestep, The number of pseudo-inputs M of SPGPs is set to 100 so that the prediction computation is performed online.

In the simulated benchmark scenario, the human is sitting in front of a desk. In this case, the robot arm helps the human by delivering objects from one position that is far away from the human to target position closer to the human. The human waits till the robot delivers the object. As different tasks are performed in terms of picking the objects and their delivery to the goal position, the temporal coherence is used to predict the actions. The action set for a human is $A^h = \{Take_0, Take_1, \dots\}$, where $Take_i$ represents an action of taking object i from its current position to the new position. The action set for the robot arm is defined as $A^r = \{Fetch_0, Fetch_1, \dots\}$.

To evaluate the quality of anticipated trajectory of human motion, modified Hausdorff distance (MHD) [7] between the ground-truth human trajectory and the predicted mean trajectory is used. In our experiments, MHD is measured for an actively moving hand joint over 1 second.

We also measured the smoothness of robot's trajectory with and without human motion prediction results. The smoothness is computed as

$$\frac{1}{T} \int_0^T \sum_{i=1}^n \ddot{\mathbf{q}}_i(t)^2, \quad (12)$$

where the two dots indicate acceleration of joint angles.

Table I highlights the performance of our algorithm in three different variations of this scenario: *arrangements of blocks*, *task order* and *confidence level*. The numbers in the "Task Order" column indicates the identifiers of human actions. Parentheses mean that the human actions in the parentheses

can be performed in any order. Arrows mean that the right actions can be performed only if the left actions are done. For example, $(0, 1) \rightarrow (2, 3)$ means that the possible action orders are $0 \rightarrow 1 \rightarrow 2 \rightarrow 3$, $0 \rightarrow 1 \rightarrow 3 \rightarrow 2$, $1 \rightarrow 0 \rightarrow 2 \rightarrow 3$ and $1 \rightarrow 0 \rightarrow 3 \rightarrow 2$. Table II shows the performance of our algorithm with a real robot. Our algorithm has been implemented on a PC with 8-core i7-4790 CPU. We used OpenMP to parallelize the computation of future human motion prediction and probabilistic collision checking.

A. Benefits of our Prediction Algorithm




Fig. 4: Different block arrangements: *Different arrangements in terms of the positions of the blocks, results in different human motions and actions. Our planner computes their intent for safe trajectory planning. The different arrangements are: (a) 1×4 . (b) 2×2 . (c) 2×4 .*

In the *Different Arrangements* scenarios, the position and layout of the blocks changes. Fig. 4 shows three different arrangements of the blocks: 1×4 , 2×2 and 2×4 . In the two cases 2×2 and 2×4 , where positions are arranged in two rows unlike the 1×4 scenario, the human arm blocks a movement from a front position to the back position. As a result, the robot needs to compute its trajectory accordingly.

Depending on the temporal coherence present in the human tasks, the human intention prediction may or may not improve the performance of our the task planner. It is shown in the *Temporal Coherence* scenarios. In the sequential order coherence, the human intention is predicted accurately with our approach with 100% certainty. In the random order, however, the human intention prediction step is not accurate until the human hand reaches the specific position. The personal order varies for each human, and reduces the possibility of predicting the next human action. When the right arm moves forward a little, $Fetch_0$ is predicted as the human intention with a high probability whereas $Fetch_1$ is predicted with low probability, even though position 1 is closer than position 0.

In the *Confidence Level* scenarios, we analyze the effect of confidence level δ_{CD} on the trajectory computed by the planner, the average task completion time, and the average motion planning time. As the confidence level becomes higher, the robot may not take the smoothest and shortest path so as to compute a collision-free path that is consistent with the confidence level.

In all cases, we observe the prediction results in smoother trajectory, using the smoothness metric defined as Equation (12). This is because the robot changes its path in advance before the human obstacle actually blocks the robot's shortest path if human motion prediction is used.

Scenarios	Arrangement	Task Order	Confidence Level	Average Prediction Time	MHD	Smoothness	
						No Pred.	Pred.
Different Arrangements	1×4	$(0, 1) \rightarrow (2, 3)$	0.95	52.0 ms	6.7 cm	2.96	1.08
	2×2	$(1, 5) \rightarrow (2, 6)$	0.95	72.4 ms	6.2 cm	5.78	1.04
	2×4	$(0, 4) \rightarrow (1, 5) \rightarrow (2, 6) \rightarrow (3, 7)$	0.95	169 ms	10.4 cm	4.82	1.15
Temporal Coherence	1×4	$0 \rightarrow 1 \rightarrow 2 \rightarrow 3$	0.95	52.1 ms	4.3 cm	1.79	0.65
	1×4	Random	0.95	105 ms	8.2 cm	5.49	1.21
	1×4	$(0, 2) \rightarrow (1, 3)$	0.95	51.7 ms	6.8 cm	3.21	1.00
Confidence Level	1×4	$0 \rightarrow (1, 2) \rightarrow 3$	0.90	47.2 ms	7.9 cm	2.90	1.17
	1×4	$0 \rightarrow (1, 2) \rightarrow 3$	0.95	50.7 ms		3.12	1.28
	1×4	$0 \rightarrow (1, 2) \rightarrow 3$	0.99	155 ms		3.76	1.40

TABLE I: Performance of our planner in three different scenarios: Different Arrangements, Temporal Coherence and Confidence Level. We take into account different arrangement of blocks as well as the confidence levels used for probabilistic collision checking. These confidence levels are used for motion prediction. The prediction results in smoother trajectory and we observe up to 4X improvement in our smoothness metric defined in Equation (12). The overall planner runs in realtime.

Scenarios	Average Prediction Time	MHD	Smoothness	
			No Pred.	Pred.
Waving Arms	20.9 ms	5.0 cm	4.88	0.91
Moving Cans (Table)	51.7 ms	7.3 cm	5.13	1.04

TABLE II: Performance of the planner with a real robot running on the 7-DOF Fetch robot next to dynamic human obstacles. The online motion planner computes safe trajectories for challenging benchmark like "moving cans". We observe almost 5X improvement in the smoothness of the trajectory due to our prediction algorithm.




Fig. 5: **Probabilistic collision checking with different confidence levels:** A collision probability less ($1 - \delta_{CD}$) implies a safe trajectory. The current pose (i.e., blue spheres) and the predicted future pose (i.e. red spheres) are shown. The robot's trajectory avoids these collisions before the human performs its action. The higher the confidence level is, the longer the distance between the human arm and the robot trajectory. (a) $\delta_{CD} = 0.90$. (b) $\delta_{CD} = 0.95$. (c) $\delta_{CD} = 0.99$.

B. Evaluation using 7-DOF Fetch Robot

We integrated our planner with the 7-DOF Fetch robot arm and evaluated in complex 3D workspaces. The robot delivers four soda cans from start locations to target locations on a desk. At the same time, the human sitting in front of the desk picks up and takes away the soda cans delivered to the target positions by the robot, which can cause collisions with the robot arm. In order to evaluate the collision avoidance capability of our approach, the human intentionally starts moving his arm to a soda can at a target location, blocking the robot's initially planned trajectory, when the robot is delivering another can moving fast. Our intention aware planner avoids collisions with the human arm and results in a smooth trajectory compared to motion planner results without human prediction.

Figure 1 shows two sequences of robot's trajectories. In the

first row, the robot arm trajectory is generated an ITOMP [24] re-planning algorithm without human motion prediction. As the human and the robot arm move too fast to re-plan collision-free trajectory. As a result, the robot collides (the second figure) or results in a jerky trajectory (the third figure). In the second row, our human motion prediction approach is incorporated as described in Section V. The robot re-plans the arm trajectory before the human actually blocks its way, resulting in collision-free path.

VIII. CONCLUSIONS, LIMITATIONS, AND FUTURE WORK

We present a novel intention-aware planning algorithm to compute safe robot trajectories in dynamic environments with humans performing different actions. Our approach uses offline learning of human motions and can account for large noise in terms of depth cameras. At runtime, our approach uses the learned human actions to predict and estimate the future motions. We use upper bounds on collision guarantees to compute safe trajectories. We highlight the performance of our planning algorithm in complex benchmarks for human-robot cooperation in both simulated and real world scenarios with 7-DOF robots.

Our approach has some limitations. As the number of human action types increases, the number of states of MDP can increase significantly. Partially Observable MDP (POMDP) for robot motion planning under uncertainty [19] can be a better approach in this case, for the sake of running time. Our probabilistic collision checking formulation is limited to environment uncertainty and does not take into account robot control errors. The performance of motion prediction algorithm depends on the variety and size of the learned data. Currently, we use supervised learning with labeled action types, but it would be useful to explore unsupervised learning based on appropriate action clustering algorithms. In particular, we would to measure the impact of robot actions on human motion, and thereby establish a two-way coupling between robot and human actions.

IX. ACKNOWLEDGEMENT

This research is supported in part by ARO grant W911NF-14-1-0437 and NSF grant 1305286.

REFERENCES

- [1] Tirthankar Bandyopadhyay, Chong Zhuang Jie, David Hsu, Marcelo H Ang Jr, Daniela Rus, and Emilio Frazzoli. Intention-aware pedestrian avoidance. In *Experimental Robotics*, pages 963–977. Springer, 2013.
- [2] Aniket Bera, Sujeong Kim, Tanmay Randhavane, Srihari Pratapa, and Dinesh Manocha. Gmp-realtime pedestrian path prediction using global and local movement patterns. *ICRA*, 2016.
- [3] Lucian Busoni, Robert Babuska, and Bart De Schutter. A comprehensive survey of multiagent reinforcement learning. *Systems, Man, and Cybernetics, Part C: Applications and Reviews, IEEE Transactions on*, 38(2):156–172, 2008.
- [4] Benjamin Choo, Michael Landau, Michael DeVore, and Peter A Beling. Statistical analysis-based error models for the microsoft kinecttm depth sensor. *Sensors*, 14(9):17430–17450, 2014.
- [5] Anca D Dragan and Siddhartha S Srinivasa. A policy-blending formalism for shared control. *The International Journal of Robotics Research*, 32(7):790–805, 2013.
- [6] Anca D Dragan, Shira Bauman, Jodi Forlizzi, and Siddhartha S Srinivasa. Effects of robot motion on human-robot collaboration. In *Proceedings of the Tenth Annual ACM/IEEE International Conference on Human-Robot Interaction*, pages 51–58. ACM, 2015.
- [7] Marie-Pierre Dubuisson and Anil K Jain. A modified hausdorff distance for object matching. In *Pattern Recognition, 1994. Vol. 1-Conference A: Computer Vision & Image Processing., Proceedings of the 12th IAPR International Conference on*, volume 1, pages 566–568. IEEE, 1994.
- [8] Carl Henrik Ek, Philip HS Torr, and Neil D Lawrence. Gaussian process latent variable models for human pose estimation. In *Machine learning for multimodal interaction*, pages 132–143. Springer, 2007.
- [9] Dieter Fox, Wolfram Burgard, and Sebastian Thrun. The dynamic window approach to collision avoidance. *IEEE Robotics & Automation Magazine*, 4(1):23–33, 1997.
- [10] Katerina Fragkiadaki, Sergey Levine, Panna Felsen, and Jitendra Malik. Recurrent network models for human dynamics. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 4346–4354, 2015.
- [11] Chiara Fulgenzi, Anne Spalanzani, and Christian Laugier. Dynamic obstacle avoidance in uncertain environment combining pvos and occupancy grid. In *Robotics and Automation, 2007 IEEE International Conference on*, pages 1610–1616. IEEE, 2007.
- [12] Charles W Groetsch. The theory of tikhonov regularization for fredholm equations of the first kind. 1984.
- [13] Kelsey P Hawkins, Nam Vo, Shray Bansal, and Aaron F Bobick. Probabilistic human action prediction and wait-sensitive planning for responsive human-robot collaboration. In *Humanoid Robots (Humanoids), 2013 13th IEEE-RAS International Conference on*, pages 499–506. IEEE, 2013.
- [14] Michael Hofmann and Darius M Gavrila. Multi-view 3d human pose estimation in complex environment. *International journal of computer vision*, 96(1):103–124, 2012.
- [15] Sujeong Kim, Stephen J Guy, Wenxi Liu, David Wilkie, Rynson WH Lau, Ming C Lin, and Dinesh Manocha. Brvo: Predicting pedestrian trajectories using velocity-space reasoning. *The International Journal of Robotics Research*, 2014.
- [16] Hema S Koppula and Ashutosh Saxena. Anticipating human activities using object affordances for reactive robotic response. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 38(1):14–29, 2016.
- [17] Hema S Koppula, Ashesh Jain, and Ashutosh Saxena. Anticipatory planning for human-robot teams. In *Experimental Robotics*, pages 453–470. Springer, 2016.
- [18] Markus Kuderer, Henrik Kretzschmar, Christoph Sprunk, and Wolfram Burgard. Feature-based prediction of trajectories for socially compliant navigation. In *Robotics: science and systems*. Citeseer, 2012.
- [19] Hanna Kurniawati, Tirthankar Bandyopadhyay, and Nicholas M Patrikalakis. Global motion planning under uncertain motion, sensing, and environment map. *Autonomous Robots*, 33(3):255–272, 2012.
- [20] Jim Mainprice and Dmitry Berenson. Human-robot collaborative manipulation planning using early prediction of human motion. In *Intelligent Robots and Systems (IROS), 2013 IEEE/RSJ International Conference on*, pages 299–306. IEEE, 2013.
- [21] Meinard Müller. Dynamic time warping. *Information retrieval for music and motion*, pages 69–84, 2007.
- [22] Stefanos Nikolaidis, Przemyslaw Lasota, Gregory Rossano, Carlos Martinez, Thomas Fuhlbrigge, and Julie Shah. Human-robot collaboration in manufacturing: Quantitative evaluation of predictable, convergent joint action. In *Robotics (ISR), 2013 44th International Symposium on*, pages 1–6. IEEE, 2013.
- [23] Jia Pan, Sachin Chitta, and Dinesh Manocha. Probabilistic collision detection between noisy point clouds using robust classification. In *International Symposium on Robotics Research (ISRR)*, 2011.
- [24] Chonhyon Park, Jia Pan, and Dinesh Manocha. IT-OMP: Incremental trajectory optimization for real-time replanning in dynamic environments. In *Proceedings of International Conference on Automated Planning and Scheduling*, 2012.
- [25] Chonhyon Park, Jia Pan, and Dinesh Manocha. Real-time optimization-based planning in dynamic environments using GPUs. In *Proceedings of IEEE International Conference on Robotics and Automation*, 2013.
- [26] Chonhyon Park, Jae Sung Park, and Dinesh Manocha. Fast and bounded probabilistic collision detection for high-dof robots in dynamic environments. In *Workshop on Algorithmic Foundations of Robotics*, 2016.
- [27] Jae Sung Park, Chonhyon Park, and Dinesh Manocha. Efficient probabilistic collision detection for non-convex

- shapes. In *Proceedings of IEEE International Conference on Robotics and Automation*. IEEE, 2017.
- [28] Claudia Pérez-D’Arpino and Julie A Shah. Fast target prediction of human reaching motion for cooperative human-robot manipulation tasks using time series classification. In *2015 IEEE International Conference on Robotics and Automation (ICRA)*, pages 6175–6182. IEEE, 2015.
- [29] S. Petti and T. Fraichard. Safe motion planning in dynamic environments. In *Proceedings of IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 2210–2215, 2005.
- [30] *Prime Sensor NITE 1.3 Algorithms notes*. PrimeSense Inc., 2010. URL <http://www.primesense.com>. Last viewed 19-01-2011 15:34.
- [31] Morgan Quigley, Ken Conley, Brian Gerkey, Josh Faust, Tully Foote, Jeremy Leibs, Rob Wheeler, and Andrew Y Ng. Ros: an open-source robot operating system. In *ICRA workshop on open source software*, volume 3, page 5. Kobe, Japan, 2009.
- [32] Jamie Shotton, Toby Sharp, Alex Kipman, Andrew Fitzgibbon, Mark Finocchio, Andrew Blake, Mat Cook, and Richard Moore. Real-time human pose recognition in parts from single depth images. *Communications of the ACM*, 56(1):116–124, 2013.
- [33] Emrah Akin Sisbot, Luis F Marin-Urias, Rachid Alami, and Thierry Simeon. A human aware mobile robot motion planner. *Robotics, IEEE Transactions on*, 23(5): 874–883, 2007.
- [34] Edward Snelson and Zoubin Ghahramani. Sparse gaussian processes using pseudo-inputs. In *Advances in neural information processing systems*, pages 1257–1264, 2005.
- [35] Richard S Sutton and Andrew G Barto. *Reinforcement learning: An introduction*. MIT press, 1998.
- [36] Pavan Turaga, Rama Chellappa, Venkatramana S Subrahmanian, and Octavian Udrea. Machine recognition of human activities: A survey. *Circuits and Systems for Video Technology, IEEE Transactions on*, 18(11):1473–1488, 2008.
- [37] Vaibhav V Unhelkar, Claudia Pérez-D’Arpino, Leia Stirling, and Julie A Shah. Human-robot co-navigation using anticipatory indicators of human walking motion. In *Robotics and Automation (ICRA), 2015 IEEE International Conference on*, pages 6183–6190. IEEE, 2015.
- [38] Raquel Urtasun, David J Fleet, and Pascal Fua. 3d people tracking with gaussian process dynamical models. In *Computer Vision and Pattern Recognition, 2006 IEEE Computer Society Conference on*, volume 1, pages 238–245. IEEE, 2006.
- [39] Jur van den Berg, Sachin Patil, Jason Sewall, Dinesh Manocha, and Ming Lin. Interactive navigation of multiple agents in crowded environments. In *Proceedings of the 2008 symposium on Interactive 3D graphics and games*, pages 139–147. ACM, 2008.
- [40] Dizan Vasquez, Thierry Fraichard, and Christian Laugier. Growing hidden markov models: An incremental tool for learning and predicting human and vehicle motion. *The International Journal of Robotics Research*, 2009.
- [41] Ji Zhu and Trevor Hastie. Kernel logistic regression and the import vector machine. *Journal of Computational and Graphical Statistics*, 2012.
- [42] Brian D Ziebart, Andrew L Maas, J Andrew Bagnell, and Anind K Dey. Maximum entropy inverse reinforcement learning. In *AAAI*, pages 1433–1438, 2008.