

Anticipatory Action Selection for Human-Robot Table Tennis

Zhikun Wang^{a,b}, Abdeslam Boularias^c, Katharina Mülling^c,
Bernhard Schölkopf^a, Jan Peters^{a,b}

^a*MPI for Intelligent Systems, Spemannstr. 38, 72076 Tübingen, Germany*

^b*TU Darmstadt, Hochschulstr. 10, 64289 Darmstadt, Germany*

^c*Carnegie Mellon University, 5000 Forbes Ave., Pittsburgh, PA 15213*

Abstract

Anticipation can enhance the capability of a robot in its interaction with humans, where the robot predicts the humans' intention for selecting its own action. We present a novel framework of anticipatory action selection for human-robot interaction, which is capable to handle nonlinear and stochastic human behaviors such as table tennis strokes and allows the robot to choose the optimal action based on prediction of the human partner's intention with uncertainty. The presented framework is generic and can be used in many human-robot interaction scenarios, for example, in navigation and human-robot co-manipulation. In this article, we conduct a case study on human-robot table tennis. Due to the limited amount of time for executing hitting movements, a robot usually needs to initiate its hitting movement before the opponent hits the ball, which requires the robot to be anticipatory based on visual observation of the opponent's movement. Previous work on Intention-Driven Dynamics Models (IDDM) allowed the robot to predict the intended target of the opponent. In this article, we address the problem of action selection and optimal timing for initiating a chosen action by formulating the anticipatory action selection as a Partially Observable Markov Decision Process (POMDP), where the transition and observation are modeled by the IDDM framework. We present two approaches to anticipatory action selection based on the POMDP formulation, i.e., a model-free policy learning method based on Least-Squares Policy Iteration (LSPI) that employs the IDDM for belief updates, and a model-based Monte-Carlo Planning (MCP) method, which benefits from the transition and observation model by the IDDM. Experimental results using real data in a simulated environment show the importance of anticipatory action selection, and that POMDPs are suitable to formulate the anticipatory action selection problem by taking into account the uncertainties in prediction. We also show that existing algorithms for POMDPs, such as LSPI and MCP, can be applied to substantially improve the robot's performance in its interaction with humans.

Keywords: Anticipation, Intention-driven dynamics model, Partially observable Markov decision process, Policy iteration, Monte-Carlo planning

1. Introduction

Humans possess the ability to coordinate their actions with others in joint activities, where *anticipation* plays an important role to integrate the prediction of the others' actions in one's own action planning (Sebanz et al., 2006). Such ability of anticipation is also crucial in Human-Robot Interaction (HRI), where an anticipatory robot needs to coordinate for its actions while interacting with humans (Hoffman and Breazeal, 2007), in both competitive and cooperative contexts. An anticipatory robot usually relies on a predictive model of the environment, especially its human partners, which "allows it to change state at an instant in accord with the model's predictions pertaining to a later instant" (Rosen et al., 2012). Predicting the underlying intention and the future actions of the human partners has been extensively studied in robotics (e.g. Kurniawati et al., 2011; Hauser, 2012; Kuderer et al., 2012; Wang et al., 2013; Bandyopadhyay et al., 2013). This prediction capability offers the promise of anticipatory robots that incorporate *anticipatory action selection* in their planning.

In this article, we focus on the anticipatory action selection based on a prediction model. Specifically, the robot chooses an action from a repertoire of motor skills based on the prediction of the human partner's intention. One important problem for the anticipatory action selection is prediction uncertainties that naturally arose due to the complexity and stochasticity of human behavior. The prediction usually becomes more accurate and less uncertain as more input, e.g., observed human movements, is obtained. However, waiting for a confident prediction causes delayed action selection and reduces the available time for the robot to execute its action (Wang et al., 2013). The anticipatory robot is, thus, forced to make decisions given a sequence of uncertain predictions, where a trade-off between prediction accuracy and reaction delay needs to be addressed. To address this issue, we formulate the decision making process as a Partially Observable Markov Decision Process (POMDP), and propose approaches to efficiently choose the robot's optimal action and decide the timing to initiate action. The framework is based on Intention-Driven Dynamics Model (IDDM), a nonparametric Bayesian model for human movements that allows to infer the intention of human partner while he is executing his action (Wang et al., 2013)

The presented framework is generic and can be used in many human-robot interaction scenarios. A typical application would be one where the robot observes the actions of a human, makes predictions about the human's intention with the IDDM, and reacts accordingly by either waiting for more observations or by executing a physical action. For example, a robot that navigates in an indoor environment alongside humans needs to constantly predict their trajectories in order to avoid collisions and to optimize its path (Ziebart et al., 2009). In such environments, humans tend to walk to a target that belongs to a small set of specific predefined targets such as doors, chairs, stairs, etc. Therefore, the human's intention can be discretized and IDDM can be used to compute a belief on where the human is heading based on the observed images

Email address: zhikun.wang@gmail.com (Zhikun Wang)

and on the images used for training. The robot can either adjust its path or slow down and wait for more observations, by using the Monte Carlo POMDP framework to forecast the effect of each decision. It is interesting to note here that humans use seemingly similar mechanisms of stopping, or slowing down, and belief updating in order to avoid collisions with pedestrians.

Another example is the problem of cooperative human-robot co-manipulation that is receiving increasing attention in the domain of intelligent prosthetics. Modern prosthetics contain semi-autonomous robotic joints that react to neurological signals (Velliste et al., 2008). These signals are high-dimensional and inherently noised. One can use IDDM to learn to predict the joint movement intended by the human user and to execute it. However, before executing the joint movement, it is also important to wait until the system is sufficiently certain that the human wants that movement. This is also a clearly optimal stopping movement that can be casted in the proposed framework (Shiryaev, 2007). Maybe a better example of human-robot co-manipulation is the scenario where the human uses a separate, fully autonomous, robotic arm for lifting objects (Jarrassé et al., 2008). In this case, the human simply reaches to an object, the robot perceives the human’s arm trajectory and predicts the intended object and reaches to it. During the lifting action, the robot also needs to predict the direction that the human is trying to follow for moving the object, by using the force and torque feedback, as well as observed trajectories with IDDM. This is also an anticipatory action selection problem that can be solved in the proposed framework.

In this article, we conduct a case study on human-robot table tennis, where anticipation is crucial for giving the robot sufficient time to execute its hitting movements (Wang et al., 2011). This choice is mainly motivated by the need of fast perception and decision-making in table tennis, in addition to the high level of noise encountered in this type of applications.

1.1. Anticipation in Human-Robot Table Tennis

Playing table tennis is a challenging task for robots due to reasons ranging from the robot’s deficiencies in perceiving the environment to the hardware limitations that restrict the actions. Hence, robot table tennis has been used as a benchmark task for high-speed vision (Acosta et al., 2003; Fässler et al., 1990), fast movement generation (Ángel et al., 2005; Mülling et al., 2011), learning (Matsushima et al., 2005; Miyazaki et al., 2005; Mülling et al., 2013), and many other research problems in robotics. Inspired by the fact that human players reconstruct and improve their dynamics during matches (Sève et al., 2002), Mülling et al. (2013) showed that a single type of hitting movement (action), e.g., forehand hitting movement, can be learned and constantly improved while interacting with human players. In practice, the robot often needs a repertoire of actions to cover the potential hitting points in its entire accessible workspace. For example, the biomimetic robot player (Mülling et al., 2011) used in this article employs three actions, namely, forehand, default, and backhand hitting movements, as shown in Figure 6. Each hitting movement is parametrized by duration, amplitude, and final goal position. Although the robot is faster and more precise than a human




Figure 1: The four stages of a typical table tennis ball rally, where the trajectories in red represent the ball trajectories and the trajectories in blue represent the racket’s movements of the robot player. Figure adapted from Mülling et al. (2011).

being, the robot player still suffers from certain hardware limitations, such as torque and acceleration limits, which severely restrict its movement abilities in this high-speed scenario. Even a beginner human player would have the upper hand simply by choosing regions that the robot cannot reach in time if the robot’s action is only based on the extrapolated trajectory of the incoming ball.

The robot’s hitting movement imitates typical human table tennis strokes, as shown in Figure 1, which consist of four stages (Ramanantsoa and Durey, 1994). In the awaiting stage, the ball moves towards the human opponent and is returned back. The robot stays at the *awaiting pose* during this stage. The preparation stage starts when the coming ball passes over the net, and the robot moves to a *preparation pose*. The *hitting stage* begins shortly after the ball bounces of the table and the human has enough information to decide for a hitting state. The racket moves towards this hitting state and hits the ball at the end of the hitting stage. It follows through in the *finishing stage* and recovers to the awaiting pose. The duration of the hitting stage is constant for expert players and lasts approximately 80ms. Even against a slower opponent, which allows less than 300ms for the robot’s hitting movement, this short time often does not suffice for initializing the racket to the preparation position and then moving it to reach the hitting state. Therefore, many desired hitting movements are not feasible due to time limits; the robot needs to initiate its hitting movement during the awaiting stage, that is, before the human opponent returns the ball.

To gain sufficient time for initializing and executing a hitting movement, the robot needs to be anticipatory about its human opponent’s intention, which human players rely heavily on (Alexander and Honish,




Figure 2: Demonstration of the robot’s hitting plane and the human opponent’s target. The hitting point is the intersection x_{hp} of the coming ball’s trajectory and the virtual hitting plane 80 cm behind the table, which is considered the *target* where the human opponent intends to shoot the ball. Our goal is to select an action according to the prediction of the target x_{hp} . Figure adapted from Mülling et al. (2011).

2009). The prediction of the human partner’s intention can be realized by modeling how the intention directs the dynamics of hitting movement (Wang et al., 2013). In the table tennis scenario, this Intention-Driven Dynamics Model (IDDM) leads to an online algorithm that, given a time series of observations, continually predicts the human player’s intended *target*, i.e., where the human intends to shoot the ball (Wang et al., 2013), as shown in Figure 2.

Anticipatory action selection needs to take into account the prediction uncertainty. The robot is likely to fail to return the ball if it has initiated a forehand (“right” in Figure 2) hitting movement and the ball is shot to its backhand (“left”) region, or vice versa. The prediction of the intended target tends to become increasingly accurate and confident as the human opponent finishes his movement (Wang et al., 2013). On the other hand, the robot requires a certain minimum time to execute its hitting movement. Therefore, the essence of the anticipatory action selection is deciding when and how to initiate the hitting movement based on the increasingly confident predictions.

1.2. Related Work and Contributions

Intention inference has been investigated in different settings, for example, using Hidden Markov Models (HMMs) to model and predict human behavior where different dynamics models were adopted to the corresponding behaviors (Pentland and Liu, 1999). Online learning of intentional motion patterns and prediction of intentions based on HMMs was proposed by Vasquez et al. (2008), which allows efficient inference in real time. The HMM can be learned incrementally to cope with new motion patterns in parallel with prediction (Vasquez et al., 2009).

Anticipation is important in many human-robot interaction scenarios (e.g., Ziebart et al., 2009; Dragan and Srinivasa, 2012; Wang et al., 2012b). Decision making can also be achieved jointly with intention

inference, in scenarios such as autonomous driving (Bandyopadhyay et al., 2013), control (Hauser, 2012), or navigation in human crowds (Kuderer et al., 2012). For example, when the state space is finite, the problem can be formulated as a partially observable Markov decision process (Kurniawati et al., 2011) and solved efficiently (Wang et al., 2012a). The issue of planning a robot’s actions while predicting actions of humans has also been tackled in the literature of human-aware planning. For example, Cirillo et al. (2010) presented a task planner for robots that recognizes activities of humans in a shared living space (e.g. kitchen). In the presented system, an HMM is used for activity recognition, based on partial observations. The planner computes a probability distribution on the actions of the humans, and finds the best robot plan accordingly. The planner also adjusts to detected changes. Compared to our work, the aforementioned work is based on the same concept of combining predictions of human behaviour with planning. Our work is however tailored to predict human motions and trajectories, instead of general plans or activities. Sisbot et al. (2007) also proposed a human-aware motion planner. The proposed system allows robots to take into account human partners by reasoning about their accessibility, vision field and preferences. The proposed system did not include, however, a dynamical model of human movements.

The anticipatory action selection decides when it is time to initiate the hitting movement. This decision is based on a trade-off between prediction accuracy and reaction delay, which is a generalization of optimal stopping problems. The optimal stopping problems have been extensively investigated in sequential analysis using Markov decision process (MDP), where the focus is usually the existence of optimal stopping rules given the transition model or obtaining closed-form solutions in specific problems (e.g., Shiryaev, 2007). We do not have closed-form solutions for the decision making problem in our application due to the complexity of the human dynamics. The optimal stopping problem can be applied in the context of classification and feature selection (Póczos et al., 2009; Gaudel and Sebag, 2010; Dulac-Arnold et al., 2012), using reinforcement learning to obtain an optimal stopping policy. The optimal stopping problem has also been studied under partial observation (Zhou, 2013), and in many applications, such as quality control (Jensen and Hsu, 1993) and finance (Décamps et al., 2005; Rishel and Helmes, 2006).

POMDPs have been used to formulate planning for human-robot interaction, for example, to assist elderly individuals with mild cognitive and physical impairments (Pineau et al., 2003; Boger et al., 2005). Previous work in human-robot interaction focused on POMDPs with representational restrictions, such as finite states (Hansen and Zhou, 2003) or time-aggregated state space (Boger et al., 2005; Broz et al., 2013). Such restrictions make the model difficult to represent complicated and nonlinear human behaviors. This article presents a novel framework that formulates action selection in human-robot interaction with continuous-state POMDPs and handles non-linearity and stochasticity in human movements with Bayesian nonparametric models. This formulation is sufficiently expressive for complicated human movements such as table tennis strokes.

Human-robot table tennis has been used as a benchmark task for high-speed vision (Acosta et al., 2003;

Fässler et al., 1990), fast movement generation (Ángel et al., 2005; Mülling et al., 2011), learning (Matsushima et al., 2005; Miyazaki et al., 2005; Mülling et al., 2013), and many other research problems in robotics. However, previous work in robot table tennis focused on reactive action planning solely based on the ball’s trajectory. This article presents a novel proactive planning framework by modeling and predicting the behavior of human opponent, such that the robot can start planning even before the human has hit the ball. To the best of our knowledge, this is the first work on human-robot table tennis that makes use of perceived human movements.

In addition, this article substantially extends previous work on the Intention-Driven Dynamics Models (IDDM), which was originally a dynamics model for intention inference and for predicting future observations (Wang et al., 2013). We propose an original formulation of the action selection problem in robot table tennis as POMDPs, wherein a Monte-Carlo planning algorithm that makes use of the IDDM was presented to address the trade-off between prediction accuracy and reaction delay. We show that this novel formulation leads to a significant improvement of the robot’s performance.

At a more general level, the most important contribution of this paper to Artificial Intelligence is presenting a well developed study on the combination of anticipation and planning. The anticipation module forecasts future states of a dynamical system based on past observations. Often, the dynamical system corresponds to human actions and intentions that are complex and difficult to model. Therefore, the anticipation module is typically learned from examples and used as a black-box by the planning module. The planning module relates to the robot’s actions. Starting from a given state, the planner inquires the anticipation module about the probabilities of different states in the future, and chooses its actions accordingly. These two modules are complementary because they solve different problems. The two modules can be developed separately. For instance, the anticipation module might have been developed for analyzing table tennis games between human players, while the planner was developed using a simulator or playing against a simple automatic ball launcher. Yet, the two modules can be combined with a minimum effort. We believe that this type of anticipation-planning combinations will play a crucial role in the development of future intelligent robots, as robots are moving out of their traditional controlled environments and increasingly interacting with humans.

The remainder of this article is organized as follows. We formulate the anticipatory action selection using POMDPs, and present approaches to policy learning and decision making in Section 2. In Section 3, we first introduce the setup of the considered robot table tennis player. Subsequently, we evaluate the effectiveness of the proposed approaches, and show that they substantially enhance the capability of the considered robot player. Finally, we conclude and summarize the contributions in Section 4.

2. Anticipatory Action Selection

The essence of anticipatory action selection is decision making given a time series of predictions, where two fundamental issues have to be addressed. First, uncertainties in the prediction and outcome need to be considered. For example in the robot table tennis setup, the uncertainty in prediction is mainly due to the fact that the opponent may still change the target before his racket hits the ball. The prediction of the opponent’s intended target, based on the observed partial movement of his stroke movement, tends to become increasingly accurate as the opponent finishes the hitting movement. Furthermore, the outcome of executing a selected action, e.g., the robot’s success of returning the ball to the opponent’s court with the chosen hitting movement, is not deterministic as the underlying dynamics of the robot arm are often too complicated to be modeled precisely at high speed. The decision making algorithm should be able to deal with the associated uncertainties.

The second fundamental issue is the timing for the robot to initiate the selected action, as the robot often requires sufficient time to execute an action. In the table tennis setup, while the predictions tend to become increasingly accurate, the robot needs sufficient time to move the arm from the awaiting pose to the desired preparation pose. The anticipatory action selection needs to trade off between prediction accuracy and delay in action selection, as both influence the success probability of the selected action. Hence, it is essential to choose the optimal action at the right time.

Partially Observable Markov Decision Processes (POMDPs) are suitable to model the uncertainties in prediction and outcome and formulate anticipatory action selection as a Markov decision process under uncertainty. We present two different approaches to decision making. In the first approach, we transform the POMDP into an equivalent fully observable Markov Decision Process (MDP). The states in the equivalent MDP are the belief states of the POMDP, which are the posterior distribution of the unobserved state given the observation history. We adopt the Intention-Driven Dynamics Model (IDDM) for belief updates and the Least-Squares Policy Iteration (LSPI) for policy learning. However, the model-free LSPI algorithm does not make use of the estimated transition model in the IDDM framework, and can be sample-insufficient in the considered application. Consequently, we present a more sample-efficient approach using the Monte-Carlo planning, where actions are chosen according to the value function estimated using the Monte-Carlo method (Thrun, 2000). The considered optimal stopping problem is a special case of POMDPs; only one action, i.e., waiting, needs to be explored. We applied upper confidence bounds (Auer, 2003) to avoid unnecessary exploration of the waiting action and effectively prune search trees. The resulting algorithm is equivalent to a tailored POMCP (Silver and Veness, 2010) algorithm for optimal stopping.

2.1. Optimal Stopping under Partial Observability

The anticipatory action selection is a generalization of optimal stopping under partial observation (Mazziotto, 1986), and, similarly, can be formulated as a POMDP. A discrete-time POMDP is defined as a tuple

$(\mathcal{S}, \mathcal{A}, \mathcal{Z}, \mathcal{P}, \Omega, \mathcal{R})$, where \mathcal{S} is a state space, $\mathcal{A} = \{0, \dots, n\}$ is a set of actions, and \mathcal{Z} is a observation space. The states evolve following a Markov transition model, governed by \mathcal{P} , where $\mathcal{P}(\mathbf{s}'|\mathbf{s}, a)$ represents the probability of going to state \mathbf{s}' from state \mathbf{s} when taking action a . The observations are generated from the states following the observation model Ω , where $\Omega(\mathbf{z}|\mathbf{s})$ is the probability of observing \mathbf{z} in state \mathbf{s} . The reward function $\mathcal{R}(\mathbf{s}, a)$ represents the expected immediate reward obtained for taking action a in state \mathbf{s} .

The set of actions \mathcal{A} consists of a waiting action $a = 0$ and a set of stopping actions $a \in \mathcal{A} \setminus \{0\}$. In the table tennis scenario, taking the waiting action $a = 0$ means to postpone the selection of a hitting action and to wait for the subsequent observation to be available; and each stopping action $a \neq 0$ leads to selecting and initiating a particular type of hitting movement, and, hence, to the termination of an episode. The immediate reward is only nonzero when a stopping action is taken, and corresponds to the outcome of the selected type of hitting movement. We consider a continuous state representation $\mathbf{s} = [\mathbf{x}, g]$ that consists of the state \mathbf{x} of the environment and the intention g of the human. Here, the intention g is assumed to be invariant during an episode; for example, the intended target does not change during the human player’s hitting movement. An observation $\mathbf{z} \in \mathcal{Z}$ includes perceived features of the environment, such as the position and velocity of the ball and the configuration of the opponent’s racket.

Finding the optimal timing to initiate the appropriate action is a POMDP problem. The decision at every time step is made by maximizing the expected future reward that it leads to, e.g., the chance of successfully returning the ball in the table tennis setup. Specifically, we want to maximize the aggregated expected reward during an episode with indefinite horizon (Hansen, 2007), given by

$$J = \mathbb{E}[r(\mathbf{s}_T, a_T)], \quad (1)$$

where the variable T denotes the stopping time, and a reward is obtained only by taking a stopping action a_T . Note that, generally, the stopping time T is unknown and determined by when a stopping action is taken. However, in many scenarios, there exists a maximum length of episode. For example in the case of human-robot table tennis, the robot’s hitting point can usually be precisely predicted 120 ms after the human player hits the ball (Wang et al., 2013), and, thus, an optimal policy would stop waiting and initiate a hitting movement. In practice, we enforce that the robot takes a stopping action when the robot’s hitting point can be reliably predicted by the vision system.

2.2. Belief Update with Intention-Driven Dynamics Model

A key step towards solving the optimal stopping problem is updating the belief on state $\mathbf{s}_t = [\mathbf{x}_t, g]$ according to the history $\mathbf{z}_{1:t}$, given by $p(\mathbf{x}_t, g|\mathbf{z}_{1:t})$.

We apply the online target prediction algorithm using the Intention-Driven Dynamics Model (IDDM). The IDDM is a discrete-time dynamics model for movement modeling and intention inference. In robotics scenarios, we often rely on noisy and high-dimensional sensor data. However, the intrinsic states are typically




Figure 3: The graphical model of the IDDM in an online manner, where we denote the intended target by g , state by \mathbf{x}_t , and observation by \mathbf{z}_t . The proposed model explicitly incorporates the intention as an input to the transition function (Wang et al., 2013). Here, we use gray nodes for the observed variables and white nodes for the latent variables.

not observable, and may have lower dimensions. Therefore, we seek a latent state representation of the relevant information in the data, and then model how the intention governs the dynamics in the latent states \mathbf{x}_t , as shown in Figure 3. The resulting model jointly learns both the latent state representation and the dynamics in the state space.

Designing a parametric dynamics model is difficult due to the complexity of nonlinear and stochastic human movements. Hence, the IDDM uses Gaussian processes to handle both the transition model $p(\mathbf{x}_{t+1}|\mathbf{x}_t, g)$ in the latent state space and the observation model $p(\mathbf{z}_t|\mathbf{x}_t)$ from the latent states to the observations. The IDDM considers the dynamics of latent states \mathbf{x} to follow an unknown function \mathbf{f} , given by

$$\mathbf{x}_{t+1} = \mathbf{f}(\mathbf{x}_t, g) + \mathbf{n}_{x,t}, \quad \mathbf{n}_{x,t} \sim \mathcal{N}(\mathbf{0}, \mathbf{S}_x). \quad (2)$$

The latent state \mathbf{x}_{t+1} at time $t+1$ depends on the latent state \mathbf{x}_t at time t as well as on the intention g , as demonstrated in the graphical model shown in Figure 3. A GP prior $\mathcal{GP}(m_x(\cdot), k_x(\cdot, \cdot))$ is placed on every dimension of \mathbf{f} with shared mean and covariance functions. Subsequently, the predictive distribution of the latent state \mathbf{x}_{t+1} conditioned on the current state \mathbf{x}_t and intention g is a Gaussian distribution given by $\mathbf{x}_{t+1} \sim \mathcal{N}(\mathbf{m}_x([\mathbf{x}_t, g]), \mathbf{\Sigma}_x([\mathbf{x}_t, g]))$ based on training inputs \mathbf{X}_x and outputs \mathbf{Y}_x . Similarly, the measurement mapping function \mathbf{h} from latent state \mathbf{x} to observations \mathbf{z} , given by

$$\mathbf{z}_t = \mathbf{h}(\mathbf{x}_t) + \mathbf{n}_{z,t}, \quad \mathbf{n}_{z,t} \sim \mathcal{N}(\mathbf{0}, \mathbf{S}_z), \quad (3)$$

is modeled by another set of GPs. The predictive probability of the observations \mathbf{z}_t is given by a Gaussian distribution $\mathbf{z}_t \sim \mathcal{N}(\mathbf{m}_z(\mathbf{x}_t), \mathbf{\Sigma}_z(\mathbf{x}_t))$, where the predictive mean and covariance are computed based on training inputs \mathbf{X}_z and outputs \mathbf{Y}_z .

The IDDM is a special case of the Hierarchical Gaussian Process Dynamics Model (Wang, 2013), where the intention is assumed to be time-invariant when learning the model. Nevertheless, this assumption does

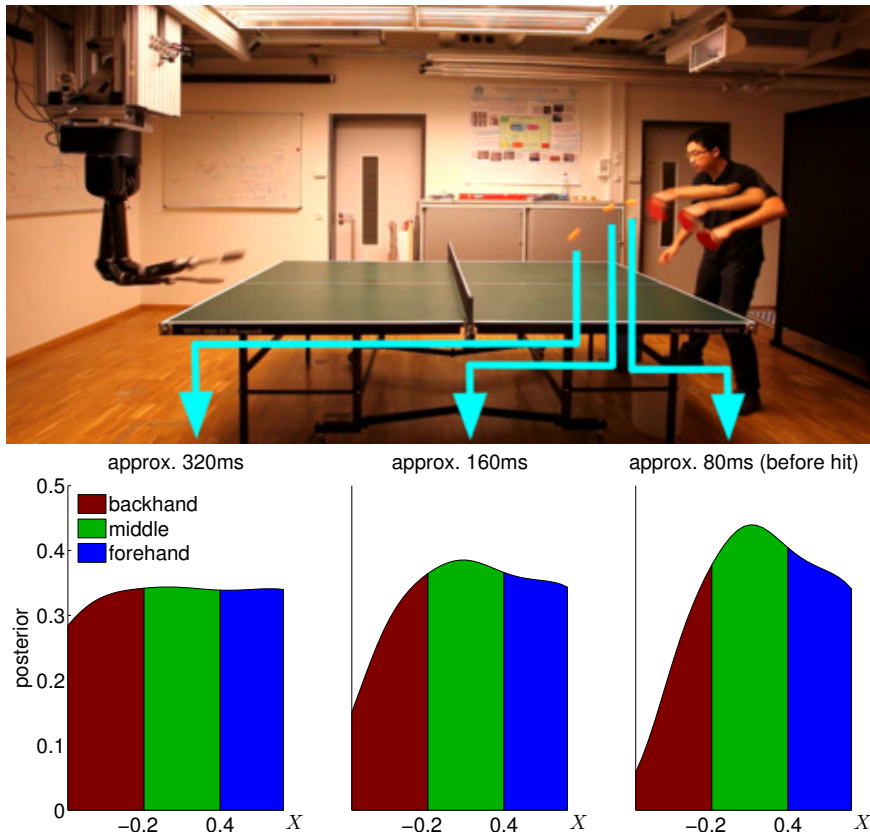


Figure 4: Bar plots show the distribution of the target (X coordinate) at approximately 320ms, 160ms, and 80ms before the player hits the ball. The prediction became increasingly confident as the player finishes the hitting movement, and the robot later chose the default hitting movement accordingly. Figure adapted from Wang et al. (2013).

not limit the model’s capability to infer time-varying intention. In fact, the online inference method (Wang et al., 2013) is able to deal with the change of intention. For more technical details of learning and approximate inference, we refer the reader to Wang et al. (2013) and Wang (2013).

Assuming that the dynamics of the human player’s racket is driven by the intended target g , we can apply the IDDM to predict the target g given a time series of observations $\mathbf{z}_{1:t}$ that are generated from corresponding latent states $\mathbf{x}_{1:t}$. While exact inference of the intention g and states \mathbf{x}_t is not tractable, Wang et al. (2013) presented an efficient online inference algorithm to update the belief $p(g, \mathbf{x}_t | \mathbf{z}_{1:t})$, i.e., the posterior probability of the intention g and the latent states \mathbf{x}_t once a new observation \mathbf{z}_t is obtained. Figure 4 shows that the predictive uncertainties decrease as the human player finishes the hitting movement. To summarize, the IDDM provides an estimate of the transition model \mathcal{T} and of the measurement model Ω in the considered POMDP with Gaussian processes, which are used for updating the belief.

2.3. Policy Learning with Belief MDP

A POMDP can be transformed to an equivalent fully observable MDP by using the belief state $\boldsymbol{\theta}_t \in \Theta$ as the observable state, which is the belief of the unobserved state $\mathbf{s}_t = [\mathbf{x}_t, g]$ given the observation history $\mathbf{z}_{1:t}$. Subsequently, we can apply reinforcement learning algorithms to learn and improve a policy π , mapping a belief state $\boldsymbol{\theta}_t$ to the action that maximizes the expected reward. Estimating the transition model $\mathcal{P}(\boldsymbol{\theta}'|\boldsymbol{\theta}, a)$ is difficult mainly due to the complicated behavior of the nonparametric dynamics model employed. Hence, we need model-free reinforcement learning algorithms that do not rely on an estimate of the transition model for belief states. Specifically, we consider the Q-learning algorithm (Sutton and Barto, 1998), and learn a state-action value function $Q^\pi(\boldsymbol{\theta}, a)$ of a policy π . The value function $Q^\pi(\boldsymbol{\theta}, a)$ measures the expected future reward when taking action a according to the current belief $\boldsymbol{\theta}$ and following the policy π thereafter. We can write the value function according to the Bellman equation, given by

$$Q^\pi(\boldsymbol{\theta}, a = \text{waiting}) = \int \mathcal{P}(\boldsymbol{\theta}'|\boldsymbol{\theta}, a = \text{waiting}) \int \pi(a'|\boldsymbol{\theta}') Q^\pi(\boldsymbol{\theta}', a') da' d\boldsymbol{\theta}', \quad (4)$$

for the waiting action, and

$$\forall a \neq 0, \quad Q^\pi(\boldsymbol{\theta}, a) = \mathcal{R}(\boldsymbol{\theta}, a), \quad (5)$$

for any stopping action.

Since the belief states are continuous, the value function Q^π cannot be written in a tabular form. A common way to deal with large or infinite state space is value function approximation by a linear combination of basis functions $\phi(\boldsymbol{\theta}, a)$, given by

$$Q^\pi(\boldsymbol{\theta}, a) \approx \hat{Q}^\pi(\boldsymbol{\theta}, a; \mathbf{w}) = \phi(\boldsymbol{\theta}, a)^T \mathbf{w}, \quad (6)$$

with parameters \mathbf{w} . Taking into account these factors, we choose the Least-Squares Policy Iteration (LSPI) algorithm (Lagoudakis and Parr, 2003), a model-free reinforcement learning algorithm with the function approximation. The LSPI has been used successfully to solve several large scale problems.

The LSPI algorithm consists of a policy evaluation step, which estimates the value function \hat{Q}^π for the current policy π , and a policy improvement step, which improves the policy π by fixing the obtained value function \hat{Q}^π . Using the linear function approximation, the policy evaluation step boils down to estimating the parameters \mathbf{w} . The integration in Eq. (4) is intractable due to the unknown transition model \mathcal{P} , and is replaced by sampling in the Q-learning framework. Derived from MDP with finite state space, the update rule of the approximation parameters \mathbf{w} can be straightforwardly applied for the considered continuous state space. The approximation parameters \mathbf{w} can be obtained from a finite set of samples $\mathcal{D} = \{(\boldsymbol{\theta}_i, a_i, r_i, \boldsymbol{\theta}'_i) | i = 1, \dots, L\}$, where $\boldsymbol{\theta}_i$ is the output of the online inference algorithm (Wang et al.,

Algorithm 1: The LSPI algorithm, which iteratively updates the approximated value function and the optimal policy.

Input : Obtained samples \mathcal{D}

Output: Approximation parameters \mathbf{w}

```

1  $\mathbf{w}' \leftarrow 0$ ;
2 repeat
3    $\mathbf{w} \leftarrow \mathbf{w}'$ ;
4   foreach  $(\boldsymbol{\theta}_i, a_i, r_i, \boldsymbol{\theta}'_i) \in \mathcal{D}$  do
5      $a' \leftarrow \pi(\boldsymbol{\theta}'_i) = \arg \max_a \phi(\boldsymbol{\theta}'_i, a)^T \mathbf{w}$ ;
6      $\hat{\mathbf{A}} \leftarrow \hat{\mathbf{A}} + \frac{1}{L} \phi(\boldsymbol{\theta}_i, a_i)(\phi(\boldsymbol{\theta}_i, a_i) - \phi(\boldsymbol{\theta}'_i, a'))^T$ ;
7      $\hat{\mathbf{b}} \leftarrow \hat{\mathbf{b}} + \frac{1}{L} \phi(\boldsymbol{\theta}_i, a_i)r_i$ ;
8    $\mathbf{w}' \leftarrow (\hat{\mathbf{A}} + \delta \mathbf{I})^{-1} \hat{\mathbf{b}}$ ;
9 until convergence;
```

2013). Given a policy π , the estimates

$$\hat{\mathbf{A}} = \frac{1}{L} \sum_{l=1}^L \phi(\boldsymbol{\theta}_l, a_l)(\phi(\boldsymbol{\theta}_l, a_l) - \phi(\boldsymbol{\theta}'_l, \pi(\boldsymbol{\theta}'_l)))^T, \quad (7)$$

$$\hat{\mathbf{b}} = \frac{1}{L} \sum_{l=1}^L \phi(\boldsymbol{\theta}_l, a_l)r_l, \quad (8)$$

are used to update the approximation parameters

$$\mathbf{w} = (\hat{\mathbf{A}} + \delta^2 \mathbf{I})^{-1} \hat{\mathbf{b}}, \quad (9)$$

where the sufficient small δ^2 is used to avoid numerical error in the inversion of $\hat{\mathbf{A}}$ (Lagoudakis and Parr, 2003).

In the policy improvement step, we improve the policy π by a new policy π' that maximizes the expected reward according to the estimated value function \hat{Q}^π . The optimal policy π' greedily chooses the action that maximizes the corresponding value function \hat{Q}^π . Therefore, we obtain an improved policy

$$\pi'(\boldsymbol{\theta}) = \arg \max_a \hat{Q}^\pi(\boldsymbol{\theta}, a). \quad (10)$$

We can obtain the optimal policy by iteratively executing the policy evaluation and improvement steps, as summarized in Algorithm 1.

2.4. Monte-Carlo Planning with POMDP

The LSPI algorithm as described above employs a model-free approach to policy learning, using the Intention-Driven Dynamics Model as a black box for updating belief $\boldsymbol{\theta}_t$ given a history of observations

Algorithm 2: The particle projection algorithm that estimates the value function of postponing the decision for one time step.

Input : Current belief $\boldsymbol{\theta}_t$
Input : Number of samples I
Output: Estimate of value function $Q(\boldsymbol{\theta}_t, a_t = \text{waiting})$

- 1 Collection of sampled rewards $\Phi = \emptyset$;
- 2 **for** $i \leftarrow 1, \dots, I$ **do**
- 3 Sample current state and intention $\mathbf{s}_t = [\mathbf{x}_t, g]$ according to belief $\boldsymbol{\theta}_t$;
- 4 Sample subsequent state $\mathbf{x}_{t+1} \sim P(\mathbf{x}_{t+1}|\mathbf{x}_t, g)$ using transition model of IDDM ;
- 5 Sample subsequent observation $\mathbf{z}_{t+1} \sim P(\mathbf{z}_{t+1}|\mathbf{x}_{t+1})$ using measurement model of IDDM ;
- 6 Update belief $\boldsymbol{\theta}_{t+1}$ provided observation \mathbf{z}_{t+1} using IDDM ;
- 7 Compute maximal expected reward for stopping $r^i = \max_{a \neq 0} \hat{Q}(\boldsymbol{\theta}_{t+1}, a; \mathbf{w})$;
- 8 Update collection of sampled rewards $\Phi = \Phi \cup \{r^i\}$;
- 9 Return $Q(\boldsymbol{\theta}_t, a_t = \text{waiting}) \approx \frac{1}{I} \sum_{r^i \in \Phi} r^i$;

$\mathbf{z}_{1:t}$. Besides the capability of belief update, the IDDM in fact provides a transition model in state space, estimated from its training data, which has not been exploited by the LSPI algorithm. Here, we present Monte-Carlo Planning (MCP), a model-based approach to action selection as an alternative to the LSPI algorithm.

Rather than estimating the value function Q^π given a policy π , we directly consider the value function Q for an optimal policy. The considered optimal stopping problem is a special case of POMDPs; there is only one single waiting action that needs to be explored. As a stopping action terminates the decision process immediately, the value function for the stopping actions $\forall a \in \mathcal{A} \setminus \{0\} : Q^\pi(\boldsymbol{\theta}, a) = Q(\boldsymbol{\theta}, a)$ holds for any belief state $\boldsymbol{\theta}$ and any policy π . It is sensible to estimate the value function for stopping actions offline, as running simulation repeatedly, especially with real-time simulator (e.g., simulated robot table tennis), during planning is expensive. We, hence, reuse the estimated value function for stopping actions by LSPI, and focus on the value function for the waiting action $Q(\boldsymbol{\theta}_t, a_t = \text{waiting})$, given by the expected future reward

$$Q(\boldsymbol{\theta}_t, a_t = \text{waiting}) = \mathbb{E} \left[\max_{a_{t+1}} Q(\boldsymbol{\theta}_{t+1}, a_{t+1}) \right] \quad (11)$$

with respect to the subsequent belief state $\boldsymbol{\theta}_{t+1}$.

The value function $Q(\boldsymbol{\theta}_t, a_t = \text{waiting})$ measures the expected reward if the agent chooses to wait for more observation. While computing the exact expectation is intractable, the value function can be estimated using Monte-Carlo approximation (Thrun, 2000), where we replace the expectation operator by an empirical average over sampled belief states. Each time we draw a sample of the current state \mathbf{s}_t , the subsequent state

\mathbf{s}_{t+1} , and the subsequent observation \mathbf{z}_{t+1} , compute the subsequent belief state $\boldsymbol{\theta}_{t+1}$ based on the IDDM, and estimate $\max_{a_{t+1}} Q(\boldsymbol{\theta}_{t+1}, a_{t+1})$ recursively. One can see that estimating the value function for waiting at the next time step $Q(\boldsymbol{\theta}_{t+1}, a_{t+1} = \text{waiting})$ again relies on the Monte-Carlo approximation, and, hence, that the number of sampled decision trees grows exponentially with the horizon. Although the horizon is often finite for the anticipatory action selection, e.g., the optimal hitting action can be chosen straightforwardly once the human player has hit the ball, we need to limit the depth of sampled decision trees in consideration of restrictive time constraints in robotics. Here, we only plan for one step ahead; namely, we estimate the value function of postponing the decision for one time step, given by

$$Q(\boldsymbol{\theta}_t, a_t = \text{waiting}) = \mathbb{E} \left[\max_{a_{t+1} \neq 0} \hat{Q}(\boldsymbol{\theta}_{t+1}, a_{t+1}; \mathbf{w}) \right]. \quad (12)$$

This estimation can be achieved by using particle projection routine (Thrun, 2000), as described in Algorithm 2.

The particle projection may still be too time-consuming to be applicable to online planning, as the Monte-Carlo approximation requires a certain amount of samples to achieve a reliable estimate. Nevertheless, consider the online planning that finds the optimal action

$$a_t = \operatorname{argmax}_{a \in \mathcal{A}} Q(\boldsymbol{\theta}_t, a), \quad (13)$$

where one only needs to know if the waiting action leads to higher expected reward than all the stopping actions, rather than the actual expected reward of waiting. We can terminate the particle projection routine if the expected reward of waiting is very likely to be higher or lower than that of the optimal stopping action $\max_{a \neq 0} Q(\boldsymbol{\theta}_t, a)$. Inspired by upper confidence bound algorithms (Auer, 2003), we use the confidence interval estimate of the expected reward for waiting to terminate the particle projection routine before the Monte-Carlo sampling completes. To obtain a confidence interval estimator, we assume that the future reward for waiting at time step t is Gaussian distributed. Given a set of sampled rewards $\Phi = \{r^1, \dots, r^n\}$, the confidence interval given a confidence level α is

$$\left[\bar{r} - \frac{cs}{\sqrt{n}}, \bar{r} + \frac{cs}{\sqrt{n}} \right], \quad (14)$$

where $n = |\Phi|$ is the number of samples, \bar{r} the sample mean, s the sample standard deviation, and c is the α percentile of a Student's t-distribution with $n - 1$ degrees of freedom. Algorithm 3 describes the resulting Monte-Carlo planning algorithm. Note that this pruning algorithm can be seen as a Partially Observable Monte-Carlo Planning (POMCP) algorithm (Silver and Veness, 2010, Algorithm 1) tailored for optimal stopping problems, where the search procedure can be even more effectively pruned before reaching timeout. The differences between the presented algorithm and POMCP are twofold. First, simulation is often computationally demanding in many robotics applications; it is sensible to estimate the value function of stopping actions offline, as stopping actions terminate the POMDP immediately and no exploration is

Algorithm 3: The MCP algorithm with early termination according to the estimate of confidence interval.

Input : Current belief θ_t

Input : Number of samples I

Input : Function approximation parameters \mathbf{w}

Input : Confidence level α

Output: Action a_t

- 1 Collection of sampled rewards $\Phi = \emptyset$;
- 2 **for** $i \leftarrow 1, \dots, I$ **do**
- 3 Sample current state and intention $\mathbf{s}_t = [\mathbf{x}_t, g]$;
- 4 Sample subsequent state $\mathbf{x}_{t+1} \sim P(\mathbf{x}_{t+1}|\mathbf{x}_t, g)$;
- 5 Sample subsequent observation $\mathbf{z}_{t+1} \sim P(\mathbf{z}_{t+1}|\mathbf{x}_{t+1})$;
- 6 Update belief θ_{t+1} provided observation \mathbf{z}_{t+1} ;
- 7 Compute expected reward $r^i = \max_{a \neq 0} \hat{Q}(\theta_{t+1}, a; \mathbf{w})$;
- 8 Update collection of sampled rewards $\Phi = \Phi \cup \{r^i\}$;
- 9 **if** *Number of samples $|\Phi|$ sufficiently large* **then**
- 10 Compute upper confidence bound U given sample Φ ;
- 11 **if** $U < \max_{a \neq 0} Q(\theta_t, a)$ **then**
- 12 \quad Return the optimal stopping action $a_t = \operatorname{argmax}_{a \neq 0} Q(\theta_t, a)$;
- 13 Compute lower confidence bound L given sample Φ ;
- 14 **if** $L > \max_{a \neq 0} Q(\theta_t, a)$ **then**
- 15 \quad Return the waiting action $a_t = 0$;
- 16 Expected reward for waiting $Q(\theta_t, a_t = \text{waiting})$ is the mean of sampled rewards in Φ ;
- 17 Return the optimal action $a_t = \operatorname{argmax}_a Q(\theta_t, a)$;

needed. Therefore, the presented algorithm estimates the value of stopping actions using all samples obtained from simulation, whereas the vanilla POMCP algorithm explores the stopping action during planning, which is not realistic for real-time simulators such as simulated robot table tennis (Mülling et al., 2011). Second, only the waiting action needs to be explored in each step, and the search procedure in POMCP can stop before timeout when either (1) the upper confidence bound of expected reward for waiting is lower than expected reward for any stopping action or (2) the lower confidence bound is higher than expected reward for all stopping actions.

In comparison to the model-free LSPI method, the MCP algorithm exploits the estimated transition model in the IDDM, and is expected to be more sample-efficient.

2.5. Basis Functions

The presented LSPI and MCP methods both employ function approximation, relying on a set of basis functions of the belief state θ . The belief state obtained by the IDDM is represented by a vector that consists of the mean and covariance of the belief on the latent state \mathbf{x} and a discretized histogram over the intended target g . The discretized histogram is illustrated in Figure 4.

We consider a set of radial basis functions for approximating the value function. We collected all the encountered belief states on the training data, and chose K centers $\bar{\theta}_1, \dots, \bar{\theta}_K$ using K-means clustering. The basis functions are given by

$$\phi(\theta, a) = [\delta_{a,0}\phi'(\theta)^T, \delta_{a,1}\phi'(\theta)^T, \dots, \delta_{a,|\mathcal{A}|}\phi'(\theta)^T]^T, \quad (15)$$

where δ is the Kronecker delta and we consider the radial basis functions

$$\phi'(\theta) = [\exp\{-\eta\|\theta - \bar{\theta}_1\|^2\}, \dots, \exp\{-\eta\|\theta - \bar{\theta}_K\|^2\}]^T \quad (16)$$

for the belief states.

3. Application in Human-Robot Table Tennis

We evaluated the LSPI and the MCP algorithms for anticipatory action selection in human-robot table tennis setup. The presented action selection methods were evaluated on the biomimetic robot table tennis player (Mülling et al., 2011), as this setup allowed exhibiting how much of an advantage such an anticipation system may offer. We expect that the system will help similarly when deployed within the skill learning framework (Mülling et al., 2013) as well as many of the recent table tennis learning systems (Huang et al., 2013; Yang et al., 2010; Matsushima et al., 2005).

3.1. Robot Player

To quantitatively evaluate the performance of the proposed methods in terms of success rates, we used the SL framework (Schaal, 2009), which consisted of a real-world setup and a sufficiently realistic simulation. The setup included a Barrett WAMTM arm with seven degrees of freedom, which is capable of high speed motion. A racket was attached to the end-effector. Table, racket and ball were compliant with the international rules of table tennis. We used a vision system of seven cameras to collect real data during a table tennis game between two human players (Wang et al., 2011), recording the players' movement and ball's trajectory. The collected data was used by the SL simulation for the following experiments, as demonstrated in Figure 5. Previous work showed that the SL framework is sufficiently realistic to simulate the robot table tennis setting¹, and that similar performance can be expected when the real robot is used (Mülling et al., 2011; Mülling et al., 2013; Wang et al., 2013).

¹See <http://robot-learning.de/Research/ProbabilisticMovementModeling> for a demonstration of the real-robot setup.




Figure 5: The SL simulated environment with the state of the robot arm and the information obtained from the vision system, including states of the opponent’s racket the ball.

In the considered setting, the robot always hits the ball on a virtual hitting plane 80 cm behind the table, as shown in Figure 2. We defined the human’s intended target as the intersection of the returned ball’s trajectory with the robot’s virtual hitting plane. As the x -coordinate (see Figure 2) was most important for choosing the type of hitting movements (Wang et al., 2013), the intention g considered here was the x -coordinate of the hitting point. Physical limitations of the robot restricted the x -coordinate to the range of ± 1.2 m from the robot’s base (the table is 1.52 m wide).

The robot player can execute three types of hitting movement (actions) that were refined and optimized for hitting points in a specific region, as shown in Figure 6. Hence, the action set \mathcal{A} consisted of one waiting action and three stopping actions, each stopping action corresponding to a type hitting movement. Note that the action selection was only used to choose a hitting type, i.e., default, forehand, or backhand. Fine-tuning of the robot’s movement can be done when the robot has initiated an action and once the returned ball can be reliably predicted from the ball’s trajectory alone. However, returning the ball outside the corresponding hitting region is difficult once the robot has initiated the chosen action (Wang et al., 2013).

When the robot’s hitting point can be precisely predicted, an optimal policy would take a stopping action to maximize the time for executing planned movement. Therefore, we enforce that the robot takes a stopping action when the robot’s hitting point can be reliably predicted by the vision system, which typically happens within 120 ms after the human player returns the ball (Wang et al., 2013).

We used a data set with recorded 270 trials. Each trial started when the ball passed over the net while flying towards the human player, and ended when the ball returned by the human player reached the robot’s hitting plane (see Figure 2), including the trajectories of the ball and the player’s racket. We excluded the trials where the ball was shot outside the overall hitting region of the robot, as shown in Figure 6,




Figure 6: Three types of hitting movement (actions) of the robot table tennis player. Each action was optimized for hitting points in a specific region as shown by the shaded rectangle.

and evaluated the performance of the policy learning algorithm on the remaining 207 trials. For the basis functions used for function approximation in Equation (16), we chose 100 centers in the experiments. Note that the IDDM model had been trained with 100 additional trials (Wang et al., 2013), and this set of data was not used in the experiments in this article.

3.2. Experimental Results

We evaluated the algorithms using ten-fold cross-validation. We carried out each round of the cross-validation for ten times to reduce the randomness in the robot player. In the first experiment, for each round of the cross-validation, we obtained sampled episodes $\mathcal{D} = \{(\theta_i, a_i, r_i, \theta'_i) | i = 1, \dots, L\}$ on the nine subsamples for training, where all valid combinations of action and time to initiate were enumerated on each trial. For each episode, the reward, i.e., success of the robot's return, was obtained from the SL simulation, where a reward of 1 was given for successful return of the ball and -1 for failing to return. For the LSPI algorithm, the policy was learned from the sample \mathcal{D} using Algorithm 1, and evaluated on the one subsample for test. The anticipatory action selection following the learned policy by LSPI led to successful returns for the robots in 153.1 ± 0.8 times, with an average success rate of 74%. We subsequently evaluated the MCP algorithm in the same manner. The action selection following Algorithm 3 led to successful returns for the robots in 153.3 ± 0.6 times, an average success rate of 74%.

To demonstrate the importance of action selection, we evaluated a baseline that exclusively used a single action, i.e., the robot always use the same type of hitting movement. Every type of hitting movement yielded a relatively high success rate in its designated regions. However, the overall rate on the entire data set was considerably reduced due to its poor performance in the other regions. Figure 7 showed the number of successful returns for using each action initiated at 240ms before the opponent hits the ball, such that the robot had sufficient time to complete the movement. The robot player without anticipation would achieve an average success rate of 64% for using only the default hitting movement, and 53% and 59% for using only the forehand and backhand movements, respectively. Therefore, both the LSPI algorithm and the MCP

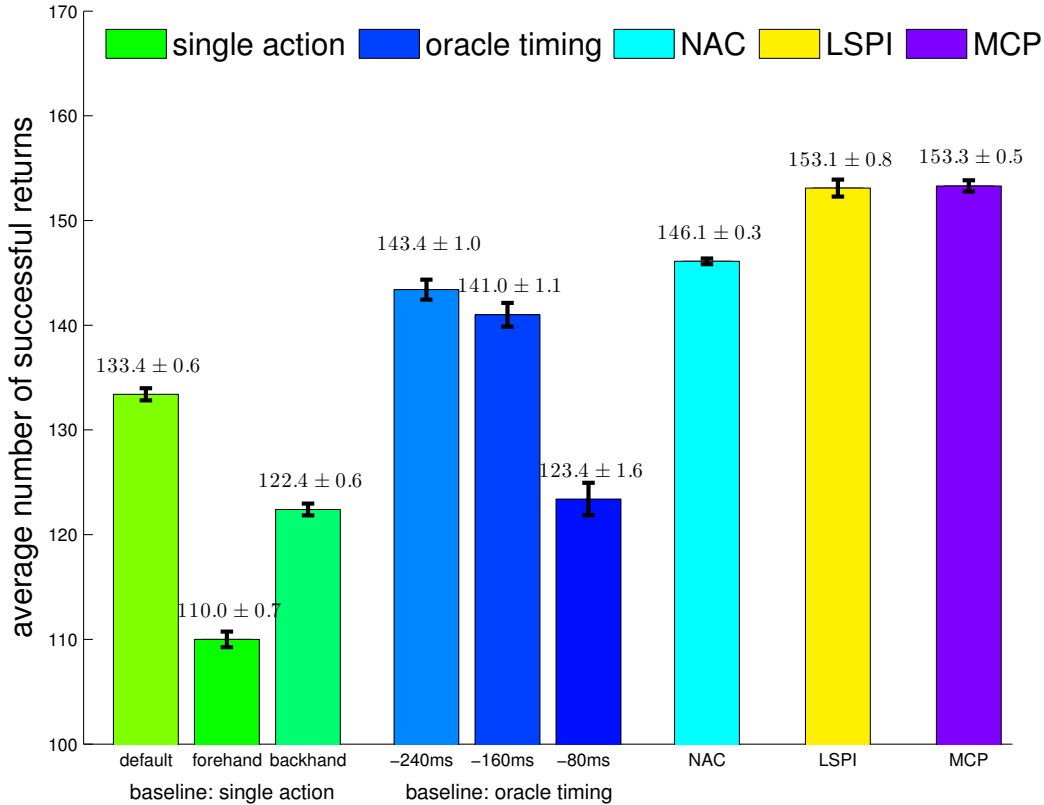


Figure 7: Performance of different methods evaluated in ten repetitions on the test data with 207 valid trials, based on enumerated training episodes. For each method, we evaluated the performance in terms of averaged number of successful returns, using cross-validation in each repetition. The first group of bars showed the baseline performance using only a single action (default, forehand, or backhand hitting movement) for all the test trials. The second group of bars showed the performance with oracle timing, where the action was always selected at a specific time (240ms, 160ms, and 80ms) before the opponent hit the ball. The last three bars showed the performance of the NAC algorithm, the LSPI algorithm, and the MCP algorithm. The error bars correspond to standard error of the mean estimated from ten repetitions.

algorithm significantly improved the performance of the robot player (p-value 0.03, using chi-squared test) by learning which action should be executed.

We considered another baseline with oracle timing, i.e., the action was always chosen at a specific time before the opponent has hit the ball. This hypothetical setting allows to demonstrate the importance of choosing an optimal stopping time. Without the waiting action, the LSPI algorithm and the MCP algorithm were equivalent, and both achieved the success return 143.4 ± 1.0 times at 240ms before the opponent hits the ball, 141.0 ± 1.1 times at 160ms, and 123.4 ± 1.6 times at 80ms, where the performance tended to decline as the reaction delay was increased. One can see that the waiting action played an important role in trading off the prediction accuracy and the reaction delay. Note that these time steps were only used for this baseline method with oracle timing, and that the other evaluated algorithms were not aware of the hitting time of

the opponent in advance. Nevertheless, both the LSPI algorithm and the MCP algorithm outperformed this hypothetical baseline by learning when a chosen action should be initiated.

In addition, we consider a naive baseline where the action is chosen only based on the most likely prediction without considering predictive uncertainty, i.e., the belief state θ is represented only by the mean of latent state and the most likely intention. We first consider the setting with oracle timing. The robot successfully returned the ball 131.0 ± 4.9 , 128.8 ± 5.5 , and $111.7.4 \pm 5.8$ times when the hitting movement was chosen at 240ms, 160ms, and 80ms, respectively, before the opponent hits the ball. The performance is not better than simply using a single action, as it is crucial to take into account the uncertainty in prediction for action selection. We then consider the setting where the timing for action selection is decided by the learned policy using LSPI. The robot successfully returned the ball 127.2 ± 4.2 times. The significantly degenerated performance (p-value 0.006) indicates that the uncertainty in belief state is also crucial for optimal timing. The presented formulation with POMDPs in this article is suitable to represent such optimal stopping problems and propagate the uncertainty in prediction. We do not consider Monte-Carlo planning algorithm in this baseline setting, as the uncertainty in belief state are needed for sampling future states.

Furthermore, we also compared to policy gradient methods (Sutton et al., 1999). Specifically, we adopted episodic Natural Actor-Critic (NAC) algorithm (Peters and Schaal, 2008) and formulated the policy by softmax function (Sutton and Barto, 1998). The NAC is a policy search method that directly estimates the gradient of the value function, which can be easier to do than estimating the value function itself. Moreover, one can use available domain knowledge to choose the type of policies that are searched. However, as an on-policy approach, the NAC is not very efficient in the use of sampled data. In addition, as in all gradient methods, the NAC can suffer from a premature convergence to a local optimum depending on the chosen step-size. In this application, the LSPI and MCP algorithms both outperformed the NAC algorithm, as shown in Figure 7.

In the second experiment, we compared LSPI and MCP in terms of sample efficiency. We obtained samples \mathcal{D} from the training data by following a uniform policy, which chose the waiting action with a probability of $\frac{1}{2}$ and each of the three stopping actions with a probability of $\frac{1}{6}$. The first group of bars in Figure 8 was obtained by sampling one episode for each trial in the training data. This amount of data was insufficient for LSPI to acquire a reliable estimate of the value function, especially for the waiting action. While the LSPI method performed even worse than the baseline that only chose the default hitting movement, the MCP method already achieved a substantial improvement by taking advantage of the estimated transition model in the IDDM. As we increased the number of sampled episode for each trial, the performance was improved for both methods. The MCP method always outperformed the LSPI method, although the advantages became smaller as more samples were available. The experimental results shown in Figure 8 verified that the MCP can be more sample-efficient than the LSPI. Due to the fact that sampling is expensive in many real-robot applications, we advocate the use of the MCP method for anticipatory action selection.

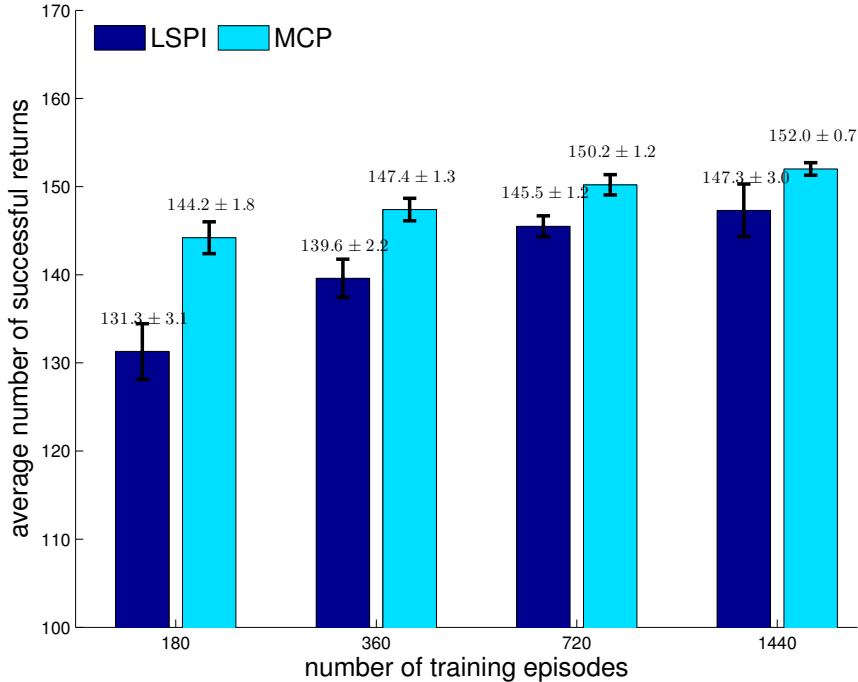


Figure 8: Performance of the LSPI and MCP methods evaluated in ten repetitions on the test data with 207 valid trials, based on sampled episodes on the training data. The numbers on the X-axis showed the number of sampled episodes on the training data.

3.3. Discussion

We conducted a case study with a simplified human-robot table tennis scenario. We assume that the recruited non-professional players do not deliberately mislead the opponent by changing their intended target during the execution of stroke. More generally, IDDM assumes that the intention is time-invariant and the main driving factor when learning the dynamics model. For example, the speed and spin of the table tennis ball, which also affect the dynamics of strokes, are not explicitly considered in the model². This assumption does not necessarily limit the capability of IDDM for intention inference (Wang et al., 2013) and planning. Moreover, this assumption can be further relaxed by taking into account other driving factors as exogenous variables in the dynamics model, leading to Hierarchical Gaussian Process Dynamics Models (H-GPDMs), discussed by (Wang, 2013). Note that, however, the method is not applicable for adversarial planning, as the assumption of time-invariant intention is violated in adversarial scenarios. Incorporation of game-theoretic perspectives in the framework of H-GPDMs is a future direction of this work.

Despite the simplification, the case study clearly showed the importance of combining anticipation and planning and the feasibility of presented methods. We discuss implications of the experimental results in

²The speed and spin of the ball are implicitly modeled by the data-driven, non-parametric dynamics model.

the general context of human-robot interaction as follows. **Action selection is important.** Robot players with action selection (baseline: oracle timing, NAC, LSPI, and MCP) outperformed that using a single action (baseline: single action), as shown in Figure 7. **Anticipatory action selection substantially improves the robot’s performance.** The baseline (oracle timing) in Figure 7 shows that it is important for the robot to be anticipatory and to choose optimal action even before the human’s intention can be accurately predicted. **Optimal stopping is essential in anticipatory action selection.** Comparison between methods for optimal stopping (NAC, LSPI, and MCP) and baseline (oracle timing) shows that it is important to trade off between confidence in prediction and time for executing action. POMDPs are suitable formulation of such optimal stopping problems, and existing algorithms for POMDPs can be applied to solve the optimal stopping problem. **Uncertainties in prediction needs to be taken into account for optimal stopping.** The baseline that only used predictive mean of latent state and most likely intention has substantially degenerated performance, indicating that uncertainties in prediction are important and informative for optimal stopping. **Model-based planning can be more sample-efficient than model-free policy learning.** Comparison between LSPI (model-free policy learning) and MCP (model-based planning) in Figure 8 shows that model-based planning tends to be more sample-efficient than model-free policy learning.

Besides the considered scenario, the method is generic and can be used in other human-robot interaction scenarios. Human-robot co-manipulation is another type of problems that can benefit from the proposed approach. In this problem, a robot assists a human in lifting and moving objects. Co-manipulation is needed in industrial settings where the objects are often too heavy, it is also increasingly used in robotic prosthetics where a robotic arm replaces a lost human arm. In these scenarios, the robot should predict which object among several ones the human is trying to reach and whether the human will need help. Upon becoming certain about the intended target, the robotic arm’s trajectory should be planned and the robotic hand should start moving toward the intended object. Adapting our approach to this problem seems straightforward. We only need to re-define the intention \mathbf{g} as the coordinates of the intended object in a three-dimensional space. Observations $\mathbf{z}_{1..t}$ can remain the observed trajectory of the human’s arm. The robot decides to initiate a movement toward the intended object, to wait until it becomes more certain, or to do nothing if the object can be lifted by the human without any help. The presented method is applicable when the intention is multi-dimensional variables, in which case MCP is a preferred tool for solving POMDPs; it is non-trivial for LSPI to model and parametrize the distribution of multi-dimensional and potentially multi-modal intention.

Note that, however, the proposed method does not straightforwardly work for structured intention prediction such as predicting trajectories of cyclist. There are two potential approaches based on the proposed method for dealing with structured intention. Take the avoidance of a cyclist as an example. Rather than predicting the trajectory, we can formulate the intention as the direction where the cyclist wants to go in

a short period of time. To take into account the intention change over time, one can model the dynamics of the cyclist with a H-GPDM (Wang, 2013), where the intention g_t follows a Markov chain. The other potential approach is to design a new covariance function $k_x(\cdot, \cdot)$ in the transition function in Eq. (2) for structured inputs (Gärtner, 2003), for example measuring the similarity of cyclist trajectories.

4. Conclusions

In this article, we introduced novel formulation and methods for anticipatory action selection for human-robot interaction, combining anticipation and planning. We formulated the anticipatory action selection as optimal stopping in a partially observable Markov decision processes. The presented formulation, using latent state variables and Bayesian nonparametric model, is expressive for nonlinear and stochastic human behaviors, such as table tennis strokes. We first presented a policy learning approach using the Least-Squares Policy Iteration algorithm. However, the LSPI can be sample-inefficient, as it does not exploit the transition model in the Intention-Driven Dynamics Models. Consequently, we presented the Monte-Carlo Planning approach, which benefits from the transition model estimated by the IDDM framework. The presented framework is generic and can be used in many other human-robot interaction scenarios, for example, in navigation and human-robot co-manipulation. A typical application would be one where the robot observes the actions of a human, makes predictions about the human’s intention with the IDDM, and reacts accordingly by either waiting for more observations or by executing a physical action, where the uncertainties in prediction needs to be considered in deciding when and how to plan the robot’s own actions.

In this article, we motivated and evaluated the proposed framework with human-robot table tennis games, where a bottleneck is the limited amount of time for the robot to execute a hitting movement. Movement initiation requires an early decision on the type of action, such as a forehand or backhand movement, before the opponent has hit the ball. Experimental results using real data and a simulated environment showed that the anticipatory action selection can be used for a robot table tennis player to enhance its performance against human players, where the robot decided the timing to initiate a selected hitting movement according to the prediction of the human opponent. We concluded that anticipatory action selection substantially improves the robot’s performance, and that uncertainties in prediction needs to be taken into account for optimal stopping. POMDPs had been shown to be suitable to formulate the anticipatory action selection problem by taking into account the uncertainties in prediction. We also showed that existing algorithms for POMDPs, such as LSPI and MCP, can be applied to substantially improve the robot’s performance in its interaction with humans.

Acosta, L., Rodrigo, J., Mendez, J., Marichal, G., Sigut, M., 2003. Ping-pong player prototype. *IEEE Robotics and Automation Magazine* 10 (4), 44–52.

Alexander, M., Honish, A., 2009. Table tennis: a brief overview of biomechanical aspects of the game for coaches and players.

- Tech. rep., Faculty of Kinesiology and Recreation Management, University of Manitoba.
 URL http://umanitoba.ca/faculties/kinrec/research/media/table_tennis.pdf
- Ángel, L., Sebastián, J., Salterén, R., Aracil, R., Gutiérrez, R., 2005. RoboTennis: design, dynamic modeling and preliminary control. In: Proceedings of IEEE/ASME International Conference on Advanced Intelligent Mechatronics. pp. 747–752.
- Auer, P., 2003. Using confidence bounds for exploitation-exploration trade-offs. *Journal of Machine Learning Research* 3, 397–422.
- Bandyopadhyay, T., Won, K. S., Frazzoli, E., Hsu, D., Lee, W. S., Rus, D., 2013. Intention-aware motion planning. In: *Algorithmic Foundations of Robotics X*. Vol. 86. Springer Berlin Heidelberg, pp. 475–491.
- Boger, J., Poupart, P., Hoey, J., Boutillier, C., Fernie, G., Mihailidis, A., 2005. A decision-theoretic approach to task assistance for persons with dementia. In: Proceedings of the 19th international joint conference on Artificial intelligence. pp. 1293–1299.
- Broz, F., Nourbakhsh, I., Simmons, R., 2013. Planning for human–robot interaction in socially situated tasks. *International Journal of Social Robotics* 5 (2), 193–214.
- Cirillo, M., Karlsson, L., Saffiotti, A., 2010. Human-aware task planning: an application to mobile robots. *ACM Transactions on Intelligent Systems and Technology* 1 (2), 15.
- Décamps, J.-P., Mariotti, T., Villeneuve, S., 2005. Investment timing under incomplete information. *Mathematics of Operations Research* 30 (2), 472–500.
- Dragan, A. D., Srinivasa, S. S., 2012. Formalizing assistive teleoperation. In: *Proceedings of Robotics: Science and Systems*.
- Dulac-Arnold, G., Denoyer, L., Preux, P., Gallinari, P., 2012. Sequential approaches for learning datum-wise sparse representations. *Machine Learning* 89 (1-2), 87–122.
- Fässler, H., Beyer, H., Wen, J., 1990. A robot ping pong player: optimized mechanics, high performance 3d vision, and intelligent sensor control. *Robotersysteme* 6 (3), 161–170.
- Gärtner, T., 2003. A survey of kernels for structured data. *ACM SIGKDD Explorations Newsletter* 5 (1), 49–58.
- Gaudel, R., Sebag, M., 2010. Feature selection as a one-player game. In: Proceedings of the 27th International Conference on Machine Learning. pp. 359–366.
- Hansen, E. A., 2007. Indefinite-horizon pomdps with action-based termination. In: *Proceedings of the National Conference on Artificial Intelligence*. Vol. 22. Menlo Park, CA; Cambridge, MA; London; AAAI Press; MIT Press; 1999, p. 1237.
- Hansen, E. A., Zhou, R., 2003. Synthesis of hierarchical finite-state controllers for pomdps. In: *Proceedings of the Thirteenth International Conference on Automated Planning and Scheduling*. pp. 113–122.
- Hauser, K., 2012. Recognition, prediction, and planning for assisted teleoperation of freeform tasks. In: *Proceedings of Robotics: Science and Systems*.
- Hoffman, G., Breazeal, C., 2007. Cost-based anticipatory action selection for human–robot fluency. *IEEE Transactions on Robotics* 23 (5), 952–961.
- Huang, Y., Xu, D., Tan, M., Su, H., 2013. Adding active learning to LWR for ping-pong playing robot. *IEEE Transactions on Control Systems Technology* 21 (4), 1489 – 1494.
- Jarrassé, N., Paik, J., Pasqui, V., Morel, G., 2008. How can human motion prediction increase transparency? In: *IEEE International Conference on Robotics and Automation*. IEEE, pp. 2134–2139.
- Jensen, U., Hsu, G.-H., 1993. Optimal stopping by means of point process observations with applications in reliability. *Mathematics of Operations Research* 18 (3), 645–657.
- Kuderer, M., Kretschmar, H., Sprunk, C., Burgard, W., 2012. Feature-based prediction of trajectories for socially compliant navigation. In: *Proceedings of Robotics: Science and Systems*.
- Kurniawati, H., Du, Y., Hsu, D., Lee, W., 2011. Motion planning under uncertainty for robotic tasks with long time horizons. *International Journal of Robotics Research* 30 (3), 308–323.
- Lagoudakis, M., Parr, R., 2003. Least-squares policy iteration. *The Journal of Machine Learning Research* 4, 1107–1149.

- Matsushima, M., Hashimoto, T., Takeuchi, M., Miyazaki, F., 2005. A learning approach to robotic table tennis. *IEEE Transactions on Robotics* 21 (4), 767–771.
- Mazziotto, G., 1986. Approximations of the optimal stopping problem in partial observation. *Journal of Applied Probability* 23, 341–354.
- Miyazaki, F., Matsushima, M., Takeuchi, M., 2005. Learning to dynamically manipulate: A table tennis robot controls a ball and rallies with a human being. *Advances in Robot Control*, 3137–341.
- Mülling, K., Kober, J., Kroemer, O., Peters, J., 2013. Learning to select and generalize striking movements in robot table tennis. *International Journal of Robotics Research* 32 (3), 263–279.
- Mülling, K., Kober, J., Peters, J., 2011. A biomimetic approach to robot table tennis. *Adaptive Behavior* 19 (5), 359–376.
- Pentland, A., Liu, A., 1999. Modeling and prediction of human behavior. *Neural Computation* 11 (1), 229–242.
- Peters, J., Schaal, S., 2008. Natural actor-critic. *Neurocomputing* 71 (7), 1180–1190.
- Pineau, J., Montemerlo, M., Pollack, M., Roy, N., Thrun, S., 2003. Towards robotic assistants in nursing homes: Challenges and results. *Robotics and Autonomous Systems* 42 (3), 271–281.
- Póczos, B., Abbasi-Yadkori, Y., Szepesvári, C., Greiner, R., Sturtevant, N., 2009. Learning when to stop thinking and do something! In: *Proceedings of the 26th Annual International Conference on Machine Learning*. ACM, pp. 825–832.
- Ramanantsoa, M., Durey, A., 1994. Towards a stroke construction model. *International Journal of Table Tennis Science* 2, 97–114.
- Rishel, R., Helmes, K., 2006. A variational inequality sufficient condition for optimal stopping with application to an optimal stock selling problem. *SIAM Journal on Control and Optimization* 45 (2), 580–598.
- Rosen, R., Rosen, J., Kineman, J., Nadin, M., 2012. *Anticipatory Systems: Philosophical, Mathematical, and Methodological Foundations*. IFSR International Series on Systems Science and Engineering. Springer.
- Schaal, S., 2009. The SL simulation and real-time control software package. Tech. rep., University of Southern California.
- Sebanz, N., Bekkering, H., Knoblich, G., 2006. Joint action: bodies and minds moving together. *Trends in Cognitive Sciences* 10 (2), 70–76.
- Sève, C., Saury, J., Theureau, J., Durand, M., 2002. Activity organization and knowledge construction during competitive interaction in table tennis. *Cognitive Systems Research* 3 (3), 501–522.
- Shiryayev, A. N., 2007. *Optimal stopping rules*. Vol. 8. Springer.
- Silver, D., Veness, J., 2010. Monte-carlo planning in large pomdps. In: *Advances in Neural Information Processing Systems*. Vol. 23. pp. 2164–2172.
- Sisbot, E. A., Marin-Urias, L. F., Alami, R., Simeon, T., 2007. A human aware mobile robot motion planner. *IEEE Transactions on Robotics* 23 (5), 874–883.
- Sutton, R., Barto, A., 1998. *Reinforcement learning: An introduction*. The MIT press.
- Sutton, R. S., McAllester, D. A., Singh, S. P., Mansour, Y., 1999. Policy gradient methods for reinforcement learning with function approximation. In: *Advances in Neural Information Processing Systems* 12. pp. 1057–1063.
- Thrun, S., 2000. Monte Carlo POMDPs. In: *Advances in Neural Information Processing Systems* 12. MIT Press, pp. 1064–1070.
- Vasquez, D., Fraichard, T., Aycard, O., Laugier, C., 2008. Intentional motion on-line learning and prediction. *Machine Vision and Applications* 19 (5), 411–425.
- Vasquez, D., Fraichard, T., Laugier, C., 2009. Growing hidden Markov models: An incremental tool for learning and predicting human and vehicle motion. *International Journal of Robotics Research* 28 (11-12), 1486–1506.
- Velliste, M., Perel, S., Spalding, M. C., Whitford, A. S., Schwartz, A. B., 2008. Cortical control of a prosthetic arm for self-feeding. *Nature* 453 (7198), 1098–1101.
- Wang, Y., Won, K., Hsu, D., Lee, W., 2012a. Monte Carlo Bayesian reinforcement learning. In: *Proceedings of the 29th International Conference on Machine Learning*. pp. 1135–1142.

- Wang, Z., 2013. Intention inference and decision making with hierarchical gaussian process dynamics models. Ph.D. thesis, TU Darmstadt.
URL <http://tuprints.ulb.tu-darmstadt.de/3617/>
- Wang, Z., Deisenroth, M. P., Amor, H. B., Vogt, D., Schölkopf, B., Peters, J., 2012b. Probabilistic Modeling of Human Movements for Intention Inference. In: Proceedings of Robotics: Science and Systems (R:SS).
- Wang, Z., Lampert, C. H., Mulling, K., Scholkopf, B., Peters, J., 2011. Learning anticipation policies for robot table tennis. In: Proceedings of IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS). pp. 332–337.
- Wang, Z., Mülling, K., Deisenroth, M. P., Amor, H. B., Vogt, D., Schölkopf, B., Peters, J., 2013. Probabilistic Movement Modeling for Intention Inference in Human-Robot Interaction. *International Journal of Robotics Research* 32 (7), 841–858.
- Yang, P., Xu, D., Wang, H., Zhang, Z., 2010. Control system design for a 5-DOF table tennis robot. In: Proceedings of International Conference on Control Automation Robotics and Vision. pp. 1731–1735.
- Zhou, E., 2013. Optimal stopping under partial observation: Near-value iteration. *IEEE Transactions on Automatic Control* 58 (2), 500–506.
- Ziebart, B. D., Ratliff, N., Gallagher, G., Mertz, C., Peterson, K., Bagnell, J. A., Hebert, M., Dey, A. K., Srinivasa, S., 2009. Planning-based prediction for pedestrians. In: Proceedings of IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS). pp. 3931–3936.