

Anticipative Interaction Primitives for Human-Robot Collaboration

Guilherme Maeda¹, Aayush Maloo², Marco Ewerton¹, Rudolf Lioutikov¹, Jan Peters^{1,3}

¹ Intelligent Autonomous Systems, Technische Universitaet Darmstadt, Germany, {maeda, ewerton, lioutikov, peters}@ias.tu-darmstadt.de

² Indian Institute of Technology Madras, Chennai, India, {ed12b001}@smai.iitm.ac.in

³ Max Planck Institute for Intelligent Systems, Tuebingen, Germany, {jan.peters}@tuebingen.mpg.de

Abstract

This paper introduces our initial investigation on the problem of providing a semi-autonomous robot collaborator with anticipative capabilities to predict human actions. Anticipative robot behavior is a desired characteristic of robot collaborators that lead to fluid, proactive interactions. We are particularly interested in improving reactive methods that rely on human action recognition to activate the corresponding robot action. Action recognition invariably causes delay in the robot's response, and the goal of our method is to eliminate this delay by predicting the next human action. Prediction is achieved by using a lookup table containing variations of assembly sequences, previously demonstrated by different users. The method uses the nearest neighbor sequence in the table that matches the actual sequence of human actions. At the movement level, our method uses a probabilistic representation of interaction primitives to generate robot trajectories. The method is demonstrated using a 7 degree-of-freedom lightweight arm equipped with a 5-finger hand on an assembly task consisting of 17 steps.

Introduction

An important feature of a collaborative robot is its ability to correctly reason about the action it should take to satisfy the current (and upcoming) human action(s). At the continuous level, the robot must also move in a manner that matches the predicted human movement, such that coordinated physical collaboration is possible. As illustrated in Figure 1, by observing the movement of the human, a robot partner must not only decide if it should hand over a screw or a plate; but once the human action is evident, the robot must spatially coordinate its trajectory w.r.t the human movement. Chances that the robot will provide the correct part increase if it only reacts after the human request, but ideally the robot should anticipate and deliver such parts beforehand.

Semi-autonomy has been one of the main approaches to address complex decisions in human-robot collaboration. Under the semi-autonomy approach, actions of the robot are made functions of human actions. Thus, the human is put in charge of the supervision and planning of a task. Semi-autonomy has the immediate benefit of making the problem of human-robot collaboration more tractable, in the sense




Figure 1: In this paper a robot collaborator must decide which action it should take in a proactive manner, before the human requests such action. We motivate the method with an experiment of a collaborative toolbox assembly where the robot must anticipate if the human will ask for a plate or for a screw.

that the robot reacts to the high-level decisions of the human rather than trying to predict them. Also, semi-autonomy leads to robots that are fundamentally helpers, as opposed to leaders. This makes semi-autonomous robots easier to be accepted by human coworkers since the human can still impose his/her own preferences.

An interesting challenge for semi-autonomous robots, however, is to address the incorporation of anticipative behavior. In other words, to allow a robot to adapt its own degree of autonomy, for example, as a function of the amount of the experience the robot has with the human. Take the case of the collaborative assembly of multiple identical products. While the human may be willing to dictate which parts he/she wants from the robot during the assembly of the first five products, a natural expectation is that the robot will start to predict the pattern in which the parts are being assembled; and therefore, to anticipate and deliver them to the human before he/she asks for it.

This paper provides our initial investigations and experimental results towards methods that allow semi-autonomous robots to act with less supervision when enough experience is acquired. Here, we opt for a lookup table approach that gathers the demonstration of 26 different users that provide their personal assembly sequence. The closest sequence in the table that matches the actual human actions are used to pre-trigger the actions of the robot. As it will be explained, we also address the case when the sequence of the human ac-

tion changes, and therefore the robot's plan of actions must be adapted accordingly.

Coordination in Human-Robot Collaboration

One approach to anticipation is to provide a robot with a plan that dictates the sequence in which actions are taken, and that correctly predicts the next human action. In task planning, sequential plans are usually generated in an abstract, high level space where actions are general symbolic descriptions, as opposed to specific motions, and its sequencing can be efficiently optimized. However, the connection between task and motion has its own difficulties. They may require intense planning at the low-level (Plaku and Hager 2010), and due to the combinatorial growth of plans over time, local commitments and short advances in limited horizons must be made for efficiency (Kaelbling and Lozano-Pérez 2011). The high level planning is aggravated in human-robot collaboration due to the multi-agent setting and concurrency in actions (Toussaint et al. 2016).

Many authors approach the problem of joint human-robot actions primarily at the motion level, with less concern on a plan of actions but rather attempt to quickly generate a reactive robot movement. This is a challenging problem because it not only requires a method to quickly recognize the human intention at the early stages, but also to generate the corresponding robot trajectory. For example, Mainprice and Berenson (2013) and Hayne, Luo, and Berenson (2016) use an occupancy grid approach where the human trajectory is predicted and a robot trajectory is optimized such that it interferes the least with the predicted volume swept by the human.

In a similar vein, Tanaka et al. (2012) uses a Markov Model on discretized occupancy grid to predict the duration and trajectories of a human in an assembly line. Gaussian Mixture Models (GMMs) were used to classify the actions of the human. In (Koppula and Saxena 2013), Conditional Random Fields were used to predict the possible actions of a human. In both works, however, the trajectory generation of the robot was not addressed, at least explicitly by their proposed methods. In these cases, the environment must be structured enough such that robot trajectories can be preprogrammed, or an extra step of online trajectory generation must be applied, as in (Hayne, Luo, and Berenson 2016).

To quickly recognize the human action, demonstrated trajectories usually encoded with some probabilistic model can be advantageous. For this purpose, GMMs can be used to explicitly capture the correlation between time and states (Calinon and Billard 2009), and in (Tanaka and Kosuge 2014) have been used to predict interaction between humans. However, a recurrent problem of probabilist representations is that to better capture the spatial correlation, demonstrated trajectories must be temporarily aligned. This problem has motivated the work in (Perez-D'Arpino and Shah 2015), where an online variant of Dynamic Time Warping was applied to the framework of probabilistic flow tubes (Dong and Williams 2012). In (Ben Amor et al. 2014), an iterative version of DTW was presented and used with Interaction Primitives. Apart from (Ben Amor et al. 2014),

in common, the related works do not explicitly address robot trajectories as part of the model. The majority focus on early human action recognition, or treat the design of the robot motion as an independent step that must be executed once the action is recognized (Mainprice and Berenson 2013; Hayne, Luo, and Berenson 2016).

Anticipation can also be achieved by directly modeling motions and the temporal sequence at which they are observed. In this category, Hidden Markov Models (HMMs) have been widely used. In (Kosuge et al. 2003) HMMs were used to provide a dancing robot means to estimate the next step of the human partner. HMMs have been used in conjunction with Gaussian Mixture Regression (GMR) to augment a mixture model with sequential information of motions (Calinon et al. 2010). To address the prediction of whole sequences of motions Inamura, Tanie, and Nakamura (2003) suggested using a HMM to symbolically represent whole sequences of motions. Later, Lee, Ott, and Nakamura (2010) extended this model for physical human-robot interaction.

In teleoperation the problem of predicting the human action is also important as it allows the robot to (partially) take over the motions of the task. The principal difference with our setting is that in teleoperation both human and robot share the same physical interface, while here their interactions with the environment are physically independent, but their actions are functions of each other. Nevertheless, the problem of shared autonomy, which has been investigated in assisted teleoperation (Javdani, Srinivasa, and Bagnell (2015), Dragan and Srinivasa (2013)), provides formalizations that can be used in the context of semi-autonomy, although the bridge between shared and semi autonomy require further clarifications.

Interaction Primitives

Imitation learning (e.g. see (Schaal 1999)) is a paradigm that has been proposed to alleviate the expert programming problem. Movement primitives are representations to learn generalizable robot movements by means of demonstration, and have been a key component in many applications in imitation learning.

The use of movement primitives for learning human-robot interaction movements with Dynamical Movement Primitives was introduced in (Ben Amor et al. 2014) under the name Interaction Primitives (IPs), and later used for action recognition and multiple interactions (Maeda et al. 2016). However, although suitable for low-level policy representation, such methods do not address any high-level mechanism that incorporate their sequences of actions, that is, the sequence in which such Interaction Primitives are executed. Thus, such robots act based on human action recognition and semi-autonomy.

In regards to Interaction Primitives, the contribution of this paper is to provide IPs with a layer of prediction. That is, we are interesting in providing the robot the ability to predict the sequences in which IPs should be executed to accomplish a task. Although motivated by IPs, the method is not associated with a particular representation of movements. Other representations that demand ac-

tion recognition—such as GMMs and HMMs—can potentially profit from the method. In this work we opt for IPs because of its simplicity. The single joint probability that models the interaction can be used to both recognize the action and generate the robot trajectory. While other methods such as (Mainprice and Berenson 2013) and (Tanaka et al. 2012) demand an independent step for robot trajectory generation.

Anticipative Interaction Primitives

The proposed method uses IPs where robot movement primitives are activated based on the recognition of the current human action. We propose adding a mechanism that pre-triggers the initial parts of the robot’s movement primitive, based on training data that reveal the most likely sequence of future actions.

The workflow of the method is shown in Figure 2. The training phase consists in learning multiple Interaction Primitives, for each of the interaction patterns involved in the task $\{S_1, \dots, S_N\}$. In the case of our experiment, one pattern would be a plate handover, for example. At the planning level, we also learn the most likely sequence S^* at which such interaction patterns will occur. During execution, the plan is used to predict the next interaction $S_n \in S^*$, and to pre-trigger the robot action by executing the first part of its trajectory. Interaction Primitives are then executed to recognize the human action. If the recognized action matches the prediction, the remainder of the trajectory is executed. Otherwise, a returning trajectory is used to reset the robot state and execute the correct trajectory for the recognized action. The following Sections will detail each component of the method.

Lookup Table of Actions

For this initial investigation, we adopt a straightforward lookup table of interaction sequences. We assume a task can be segmented by the sequential execution of N interaction patterns or actions. We assume the existence of a number of D demonstrations of different users indicating their preferences in executing such sequences. These sequences are used to fill the lookup table $L \in \mathbb{R}^{D \times N}$ with elements S_{ij} .

Given the current, partial history of interactions $S' = \{S'_1, \dots, S'_K\}$ with $K < N$, the closest sequence in the lookup table is retrieved. The optimal predictive sequence is taken as the i -th row of L with the minimum distance ($S^* = L_{i^*j}$, $j = \{1..N\}$), with

$$i^* = \underset{i}{\operatorname{argmin}} \sum_{j=1}^K |S_{ij} - S'_j|, \forall i \in [1..D]. \quad (1)$$

The predictive sequence can be used to predict the sequence of actions from steps $K + 1$ to N .

For the first iteration, as S' is empty, two natural initializations are possible. One is to not predict, and wait for the human to take the actions of the first iterations, such that (1) can be subsequently used. A second approach is to use the sequence whose incidence is the highest in the table, and assume this sequence as the most likely sequence that the user will follow.

Note that the algorithm must account for the case where the next predicted human action $S^*[K + 1]$ may differ from the actual action. In this case, (1) must be called to recompute a corrected sequence. Since the lookup table is expected to be of small size ([number of demonstrations \times number of parts]), the computation of the predictive sequence is computationally inexpensive.

Interaction Primitives

Each element in the sequence S^* is used to pre-trigger the execution of the corresponding Interaction Primitive. We use the method of Interaction Probabilistic Movement Primitives (Maeda et al. 2016) since it addresses human action recognition. The method assumes that, for each interaction pattern, distributions of robot and human trajectories are normally correlated. This distribution is obtained from pairs of human-robot demonstrated trajectories. Assume a demonstration for a plate handover, where both the robot and human move simultaneously—one agent to handover, the other agent to receive—the plate¹. The trajectories are usually encoded as weight vectors using radial-basis-functions and linear regression (Bishop 2006). Thus, for the d -th demonstration define the weight vector

$$w_d = [(w_1^H)^T, \dots, (w_P^H)^T, (w_1^R)^T, \dots, (w_Q^R)^T]^T, \quad (2)$$

containing the concatenation of the parameterized trajectories of the P degrees-of-freedom (DoFs) of the human, with the Q DoFs of the robot.

Multiple demonstrations are then stacked in a matrix of weights $[w_1, \dots, w_D]^T$. Under the normal assumption $w \sim \mathcal{N}(\mu_w, \Sigma_w)$, the trajectories of both agents $y_{1:T}$ can be predicted with

$$p(y_{1:T}) = \int p(y_{1:T}|w)p(w)dw, \quad (3)$$

where $y_{1:T} \in \mathbb{R}^{(P+Q) \times T}$ is a concatenation of trajectories of all involved DoFs.

Action Recognition

Given the N possible types of interaction, the current interaction can be recognized from the human observation with Baye’s theorem

$$p(n|y_{t:t'}) \propto p(y_{t:t'}^*|n)p(n), \quad (4)$$

where $p(n)$ is a prior distribution of the n -th interaction, and $y_{t:t'}^* \in \mathbb{R}^{P \times T^*}$ are the observed trajectories of the human, observed in the interval $[t \ t']$. The likelihood in (4) can be computed with

$$\begin{aligned} p(y_{t:t'}^*|n) &= \int p(y_{t:t'}^*|\mathbf{H}_{t:t'}^T w, \Sigma_y^*)p(w)dw \\ &= \mathcal{N}(y_{t:t'}^*|\mathbf{H}_{t:t'}^T \mu_w, \mathbf{H}_{t:t'}^T \Sigma_w \mathbf{H}_{t:t'} + \Sigma_y^*), \end{aligned} \quad (5)$$

where $\mathbf{H}_{t:t'}^T$ is the matrix with the basis functions corresponding to the observed time steps, and Σ_y^* is the observation noise. The interaction pattern is then recognized from the human action with

$$\hat{S} = \underset{n}{\operatorname{argmax}} p(n|y_{t:t'}^*). \quad (6)$$

¹The robot being moved by kinesthetic demonstrations, for example.

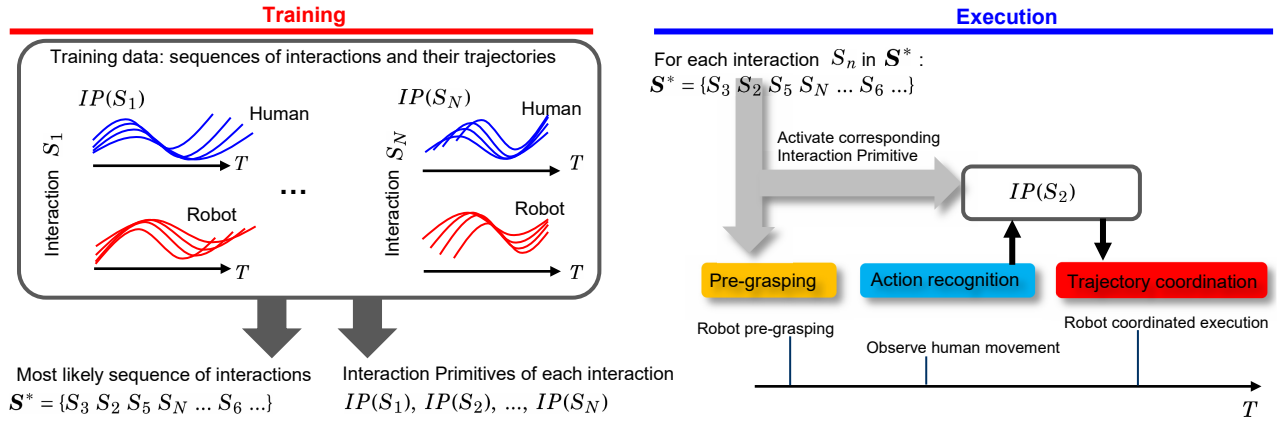


Figure 2: The training phase and execution phase of the proposed method. Multiple demonstrations are provided in the form of sequences of interactions, and in the form of pairs of human-robot trajectories. The training sequences are used to construct a lookup table while the pairs of trajectories are used to build Interaction Primitives (IPs) for each interaction. During execution, the sequences are used to pre-trigger the corresponding IP.

Coordination of Trajectories

Once the human action is recognized, the corresponding Interaction Primitive for the action \hat{S} can be conditioned on the same human observations $y_{t:t'}^*$ to produce a posterior distribution of both agent's trajectories. These trajectories predict the whole human and robot's movement. The predicted mean of the robot trajectories can then be used as a reference for the feedback controller of the robot.

Conditioning is achieved in closed-form with

$$\begin{aligned} \mu_w^+ &= \mu_w + \mathbf{K}(y_{t:t'}^* - \mathbf{H}_{t:t'}^T \mu_w), \\ \Sigma_w^+ &= \Sigma_w - \mathbf{K}(\mathbf{H}_{t:t'}^T \Sigma_w), \end{aligned} \quad (7)$$

where $\mathbf{K} = \Sigma_w \mathbf{H}_{t:t'}^T (\Sigma_w^* + \mathbf{H}_{t:t'}^T \Sigma_w \mathbf{H}_{t:t'}^T)^{-1}$. Equation (3) is used for the trajectory prediction, with the difference that the parameters of the distribution are updated $w \sim \mathcal{N}(\mu_w^+, \Sigma_w^+)$.

Experiments

As a proof-of-concept, we propose an experiment where the robot plays the role of a coworker that helps a human assembling a toolbox. Interaction Primitives are constructed by pairing joint trajectories of the robot with the Cartesian trajectories of the human. In the experiment, the human wrist provides the trajectories via tracking of optical markers. The assembly consists of a total of 17 steps where 10 actions represent screw handovers and 7 actions are given by plate handovers. The goal of the robot is to predict if the human needs a screw or a plate, and according to this prediction execute the initial pre-grasping of parts. To finalize the grasping of the object and hand it over to the human, the robot must first infer the human action \hat{S} with the procedure described in the previous Section.

The sequence in which the different plates and screws are needed are provided by the lookup table. A picture of the toolbox used in the experiments is shown in Fig. 3, with the labels of each part. The table was constructed by asking a




Figure 3: The experiment consists in assembling a toolbox with 7 plates and 10 screws. Each part was labeled and a total of 26 users were asked to show the sequences that they prefer to receive the parts.

total of 26 users to indicate which order they would assemble the toolbox. For example, the following three sequences

```

FP SP SC SC FP SC SC BP MP SC SC BP SP ...
FP SP SC SC FP SC SC BP MP SC SC BP SP ...
FP SC SC SP SC SC FP BP MP SC SC BP SP ...

```

show that the first two users prefer to have the front plate (FP) and side plate (SP) first, and then ask for the two screws (SC) that attach them together. The third user, on the other hand, prefers to ask for one plate, position the two screws on the hole, and then ask for the second plate.

Fig. 4 shows snapshots of the experiment for the handover of plate (upper row) and the handover of the screw (bottom row). Note that during the pre-grasping phase the actions of the robot are independent of the human. While the human is busy working on his own, the robot queries the next action in the sequence S^* to find out which object it should prepare to grasp. The robot then waits at the pre-grasping




Figure 4: Anticipative human-robot collaboration on the assembly task. The snapshots show the handover of plate (upper row) and the handover of the screw (bottom row). Note that during the pre-grasping phase the human and robot act independent from each other, thus increasing the overall efficiency of the task. The full experiment can be watched at <https://vimeo.com/180552453>.

position. Action recognition is then executed, providing the robot a confirmation that the pre-grasping of the chosen part is actually correct. The pictures at the last column show the spatially coordinated handover of the plate and screw. Coordination means that the position at which the robot delivers the part is conditioned on the position of the human marker. Coordination is achieved using the procedure (7) by conditioning the confirmed Interaction Primitive for the task on the current position of the human hand.

In the case the recognized human action differs from the pre-triggered action, the robot must give priority to the human decision. This means that the robot must then back off and switch to the recognized human action. This case represents a possible drawback w.r.t. robots that only operate in semi-autonomous mode as the latter do not have the ability to be pro-active, and therefore, to make mistakes in its decisions. For the experiments, we hand-coded a contingency trajectory that allows the robot to return from the erroneous pre-grasp and switch to the correct grasping. Snapshots of this procedure can be seen in Fig. 5. Note that the optimal sequence must be revised using (1) whenever the recognized action does not match the action predicted in the sequence. The experiment can be watched at <https://vimeo.com/180552453>.

Table 1 summarizes the reduction of waiting time due to the pre-grasping of the parts. The duration of the trajectory that moves the robot hand from the rest posture to the pre-grasping of a plate takes 10 seconds. The duration of the trajectory that grasps and delivers the plate to the human hand takes 17 seconds (the computation time required for action recognition and conditioning is negligible given the time scale, and usually less than 0.5 seconds). If preemption is not used, the human has then to wait for the whole pre-grasping plus handover. Thus, the reduction of the waiting time in comparison of a purely semi-autonomous task is of $10/(10+17) = 37\%$. Similar for the screw, the waiting time

reduction achieved was of 35%. Considering that there are 10 screws and 7 plates involved, the pre-triggering saves the user from waiting approximately 3 minutes when considering the whole toolbox assembly.

Table 1: Reduction of waiting time due to pre-grasping

	Pre-grasping	Grasp + handover	Reduction
Plate	10 sec	17 sec	37 %
Screw	11 sec	20 sec	35 %

Conclusions

We presented a method for a semi-autonomous robot to anticipate and initiate the movements that most likely address the next human action. The method uses a lookup table to encode demonstrated sequences of actions, and nearest neighbor to retrieve the, presumably, most likely sequence of human actions. Pre-triggering actions offers the potential to greatly decrease the amount of waiting time.

A limitation of nearest neighbor in our problem is that all possible sequences that the human may ever execute, must be present on the table. If the human executes a sequence that is not in the lookup table, the robot will opt for the closest but not exact one. Potentially, such a sequence may have many disagreements with the way the user wants to assemble the toolbox. As a consequence, the switching and execution to the correct action—which takes longer than simply not predicting the human—may lead to frustrating interactions. We are currently addressing more suitable methods for dealing with this problem.

Acknowledgment

The research leading to these results has received funding from the European Communitys Seventh Framework Programme (FP7-ICT-2013-10) under grant agreement 610878 (3rdHand).




Figure 5: In this example, the robot’s predicted action was a plate and the robot prepared the pre-grasping of the plate. The human, however, grasped the screw driver, indicating that the correct action is the screw handover. The robot then switches the action. Note that the optimal sequence must be recomputed.

References

- Ben Amor, H.; Neumann, G.; Kamthe, S.; Kroemer, O.; and Peters, J. 2014. Interaction primitives for human-robot co-operation tasks. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*.
- Bishop, C. 2006. *Pattern recognition and machine learning*, volume 4. Springer New York.
- Calinon, S., and Billard, A. 2009. Statistical learning by imitation of competing constraints in joint space and task space. *Advanced Robotics* 23(15):2059–2076.
- Calinon, S.; D’halluin, F.; Sauser, E. L.; Caldwell, D. G.; and Billard, A. G. 2010. Learning and reproduction of gestures by imitation. *Robotics & Automation Magazine, IEEE* 17(2):44–54.
- Dong, S., and Williams, B. 2012. Learning and recognition of hybrid manipulation motions in variable environments using probabilistic flow tubes. *International Journal of Social Robotics* 4(4):357–368.
- Dragan, A., and Srinivasa, S. 2013. Formalizing assistive teleoperation. In *Proceedings of Robotics: Science and Systems*.
- Hayne, R.; Luo, R.; and Berenson, D. 2016. Considering avoidance and consistency in motion planning for human-robot manipulation in a shared workspace. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*.
- Inamura, T.; Tanie, H.; and Nakamura, Y. 2003. From stochastic motion generation and recognition to geometric symbol development and manipulation. In *International Conference on Humanoid Robots*.
- Javdani, S.; Srinivasa, S.; and Bagnell, A. 2015. Shared autonomy via hindsight optimization. In *Proceedings of Robotics: Science and Systems*.
- Kaelbling, L. P., and Lozano-Pérez, T. 2011. Hierarchical task and motion planning in the now. In *Robotics and Automation (ICRA), 2011 IEEE International Conference on*, 1470–1477. IEEE.
- Koppula, H. S., and Saxena, A. 2013. Anticipating human activities using object affordances for reactive robotic response. In *Robotics: Science and Systems*.
- Kosuge, K.; Hayashi, T.; Hirata, Y.; and Tobiya, R. 2003. Dance partner robot-ms dancer. In *Intelligent Robots and Systems, 2003.(IROS 2003). Proceedings. 2003 IEEE/RSJ International Conference on*, volume 4, 3459–3464. IEEE.
- Lee, D.; Ott, C.; and Nakamura, Y. 2010. Mimetic communication model with compliant physical contact in human-humanoid interaction. *The International Journal of Robotics Research* 29(13):1684–1704.
- Maeda, G.; Neumann, G.; Ewerton, M.; Lioutikov, R.; Kroemer, O.; and Peters, J. 2016. Probabilistic movement primitives for the coordination of multiple human-robot collaborative tasks. In press: *Autonomous Robots. Special Issue on Assistive and Rehabilitation Robotics* 1–20.
- Mainprice, J., and Berenson, D. 2013. Human-robot collaborative manipulation planning using early prediction of human motion. In *Intelligent Robots and Systems (IROS), 2013 IEEE/RSJ International Conference on*, 299–306. IEEE.
- Perez-D’Arpino, C., and Shah, J. A. 2015. Fast target prediction of human reaching motion for cooperative human-robot manipulation tasks using time series classification. In *Robotics and Automation (ICRA), 2015 IEEE International Conference on*, 6175–6182. IEEE.
- Plaku, E., and Hager, G. D. 2010. Sampling-based motion and symbolic action planning with geometric and differential constraints. In *Robotics and Automation (ICRA), 2010 IEEE International Conference on*, 5002–5008. IEEE.
- Schaal, S. 1999. Is imitation learning the route to humanoid robots? *Trends in cognitive sciences* 3(6):233–242.
- Tanaka, Y., and Kosuge, K. 2014. Dynamic attack motion prediction for kendo agent. In *Intelligent Robots and Systems (IROS 2014), 2014 IEEE/RSJ International Conference on*, 2187–2193. IEEE.
- Tanaka, Y.; Kinugawa, J.; Sugahara, Y.; and Kosuge, K. 2012. Motion planning with worker’s trajectory prediction for assembly task partner robot. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 1525–1532. IEEE.
- Toussaint, M.; Munzer, T.; Mollard, Y.; Wu, L. Y.; and Lopes, M. 2016. Temporal segmentation of pair-wise interaction phases in sequential manipulation demonstrations. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*.