2016 IEEE-RAS 16th International Conference on
Humanoid Robots (Humanoids)
Cancun, Mexico, Nov 15-17, 2016

Learning Human-Robot Interactions from Human-Human Demonstrations (with Applications in Lego Rocket Assembly)

David Vogt[1], Simon Stepputtis[2], Richard Weinhold[2], Bernhard Jung[1], Heni Ben Amor[3]

***Abstract*— This video demonstrates a novel imitation learning approach for learning human-robot interactions from human-human demonstrations. During training, the movements of two human interaction partners are recorded via motion capture. From this, an interaction model is learned that inherently captures important spatial relationships as well as temporal synchrony of body movements between the two interacting partners. The interaction model is based on interaction meshes that were first adopted by the computer graphics community for the offline animation of interacting virtual characters. We developed a variant of interaction meshes that is suitable for real-time human-robot interaction scenarios. During human-robot collaboration, the learned interaction model allows for adequate spatio-temporal adaptation of the robots behavior to the movements of the human cooperation partner. Thus, the presented approach is well suited for collaborative tasks requiring continuous body movement coordination of a human and a robot. The feasibility of the approach is demonstrated with the example of a cooperative Lego rocket assembly task.**

I. Introduction

In this video, we propose to extract the interaction dynamics during a collaborative task using *learning by demonstration*. Our objective is to develop learning methods, that allow robots to gradually increase their repertoire of interaction skills without additional effort for a human programmer. For this, example *human-human* demonstrations of the collaborative task are used to infer the robot's role in each situation. The recorded demonstrations are analyzed and an interaction model is extracted. The model is composed of an Hidden Markov Model (HMM) and a set of Interaction Meshes (IM). The HMM allows a robot to identify different *contexts*. Once a specific context is identified, a corresponding IM is selected and used to generate the robot responses. An IM holds information about the spatial relationship between the pose of the robot and the pose of the user. Using an optimization scheme, the IM is deformed to best fit the current situation. The approach builds upon previous results [1], [2], [3] and extends them to complex physical interaction scenarios, e.g., collaborative assembly (see Fig. 1). In particular, complex spatio-temporal generalization can be achieved through the combination of motion recognition in low-dimensional posture spaces and a variant of IMs for real-time robot response generation.

[1]David Vogt and Bernhard Jung are with Faculty of Mathematics and Informatics, Technical University Bergakademie Freiberg, 09599, Freiberg, Germany surname.name@informatik.tu-freiberg.de

[2]Simon Stepputtis and Richard Weinhold are with Faculty of Mathematics and Informatics, Technical University Bergakademie Freiberg, 09599, Freiberg, Germany stepputt/weinhol3@student.tu-freiberg.de

[3]Heni Ben Amor is with the School of Computing, Informatics, Decision Systems Engineering, Arizona State University, 699 S Mill Ave, Tempe, AZ 85281 hbenamor@asu.edu

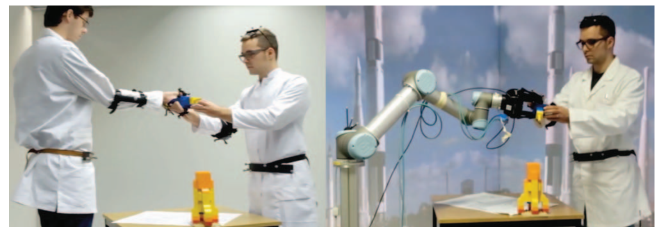

Fig. 1. From motion captured human-human demonstrations an interaction model of a cooperative task is learned. Using the learned model in a human-robot collaboration, the robot's behavior is continuously synchronized, both spatially and temporally, with the actions of the human.

II. Learning an Interaction Model

The interaction model serves to identify a robot's response to the movement of the human user in cooperative tasks. Structurally, the interaction model consists of a large database of IMs as well as several data structures for identifying the best matching IM during an ongoing human-robot interaction. Each IM represents, for one time-step, a pair of postures of the human-human demonstration. They capture the spatial relationships of the interaction and can be adapted at runtime to the specifics of the human-robot interaction such as spatial distances between body parts. We developed a correlation-based variant of IMs that include only the most relevant joints of the two human demonstrators in order to make the adaptation real-time capable.

The selection of a suitable IM during runtime is a multi-stage process involving several data structures that are learned from the interaction demonstrations. A low-dimensional global posture space, in conjunction with a Gaussian mixture model (GMM) and an HMM are used to identify the human-human demonstration that fits the current interaction best. A low-dimensional local posture space, defined for each human-human demonstration, and further segmented into smaller parts of the interaction, in combination with dynamic time warping (DTW) serves to identify the best matching IM (also adapting the robot's response to the timing of the human user).

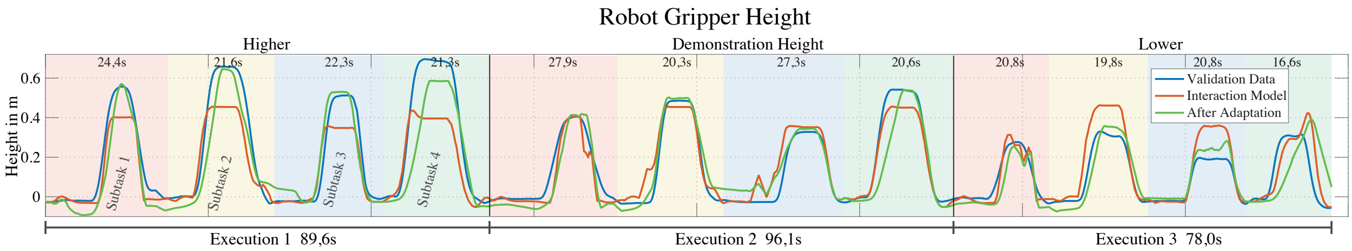

Fig. 2. An important feature of our approach is the generalization of learned behaviors to new situations. To test the generalization capabilities of our system we recorded several repetitions of the rocket assembly task. The figure shows 3 variations with differing hand-over heights. On the left, input user motions significantly higher than the original recording are shown. In the middle hand-over heights similar to the training data and on the right lower hand-over heights are illustrated. The figure depicts how a frame from the interaction model (red trajectory) is adapted to the new situation (blue trajectory). The green trajectory shows gripper heights after our adaptation.

III. Spatiotemporal Generalization

An important feature of our system is spatiotemporal generalization of trained behaviors. We evaluated our approach using one motion captured demonstration of a cooperative Lego rocket assembly task as well as additional 13 motion capture recordings of human-human interactions as validation data. Between the 13 task executions, the handover positions of the manipulated object were varied both in height. In a simulation environment, the recorded motions of the observed agent were applied to a simulated human while the responses of a simulated robot were computed using our interaction model. The simulated robot's motions were then compared to the motions of the human assistant in the validation cases. Fig. 2 depicts the robot's response in three executions of the assembly tasks including all substasks. In the figure, blue trajectories depict the height of the human hand during validation while red trajectories show hand hight in the demonstration used to train the interaction model. The robot's adapted gripper height is shown in green.

Temporal generalization is achieved using a two stage process. First, the user's motion is matched against interaction demonstrations using the HMM and, then, aligned locally in the corresponding local posture space using DTW. Fig. 2 shows three repetitions of the Lego assembly task with varying execution speeds, with a time difference between the slowest and fastest task completion of $\sim 20\,\text{s}$. The actual execution time in the validation cases is indicated by the blue trajectory. As can be seen in the Figure, the shapes of the red trajectory (selection of an appropriate IM) and the green trajectory (adaption of the selected IM to the current situation) closely match the shape of the blue trajectory. This shows that our method is able to maintain a close temporal synchrony between the movements of the human and the robot even if the task execution time is quite different from the training example.

IV. Conclusion

In the video we demonstrate our learning from demonstration system that enables robots to seamlessly interact with humans in a collaborative Lego rocket assembly tasks. Whereas current imitation learning methods almost exclusively focus on a single agent, our method is based on parallel behavior demonstrations by two interaction partners. In addition to being based on two-person behavior demonstrations, the learned interaction models also inherently capture the important spatial relationships and synchrony of body movements between two interacting partners. At runtime, the robot's reaction to the human's motion can be efficiently extracted from the IM in real-time. Therefore the presented approach is well suited for collaborative tasks requiring continuous body movement coordination of human and robot.

Our interaction models are composed of an Hidden Markov Model and a set of IMs. At runtime, the HMM is able to reliably select the appropriate interaction example from the set of training data. In particular, the HMM exhibits enough hysteresis in order to avoid an undesirable "jumping" between different interaction examples in consecutive frames that we observed in alternative action recognition methods we experimented with.

References

[1] E. S. L. Ho, T. Komura, and C.-L. Tai, "Spatial relationship preserving character motion adaptation," *ACM Trans. Graph.*, vol. 29, no. 4, jul 2010.

[2] H. Ben Amor, D. Vogt, M. Ewerton, E. Berger, B. Jung, and J. Peters, "Learning responsive robot behavior by imitation," in *2013 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, nov 2013, pp. 3257–3264.

[3] D. Vogt, B. Lorenz, S. Grehl, and B. Jung, "Behavior generation for interactive virtual humans using context-dependent interaction meshes and automated constraint extraction," *Computer Animation and Virtual Worlds*, vol. 26, no. 3-4, pp. 227–235, may 2015.

Appendix

A high-resolution video demonstrating the proposed method on a tube assembly task can be found at `https://youtu.be/_2qcU4FcGyE`.

A video showcasing our methodology in a handover task using multiple *Kinect V2* sensors can be found at `https://youtu.be/KhcvUUO-ZE0`.