

Learning Throwing and Catching Skills*

Jens Kober, Katharina Muelling, and Jan Peters

Abstract—In this video, we present approaches for learning throwing and catching skills. We first show how a hitting skill (i.e., table tennis) can be learned using a combination of imitation and reinforcement learning. This hitting skill is subsequently generalized to a catching skill. Secondly, we show how a robot can adapt a throwing skill to new targets. Finally, we demonstrate that a BioRob and a Barrett WAM can play catch together using the previously acquired skills.

I. INTRODUCTION

We show how a motor skill can be learned on multiple levels in this video. As a first step, we need a representation of the skill, which we introduce in Sect. II. We take hitting skills as an example and show how such skills can be learned using a combination of imitation and reinforcement learning. In Sect. III, we show how such skills can be refined and generalized using reinforcement learning on different levels. As hitting skills are highly related to catching skills, we also demonstrate in Sect. IV how to generalize hitting to catching skills. In a similar fashion, throwing skills can be learned as outlined in Sect. V. Combining the acquired catching and throwing skills enables a robotic tandem to play catch (see Sect. VI). Finally, we provide some details on the ball tracking system in Sect. VII.


Catching and throwing in robotics has previously been studied by [1–5] and many others. This video focuses on learning, generalizing, and combining skills.

II. REPRESENTING MOTOR SKILLS

Learning new skills can frequently be helped significantly by choosing a movement template representation that facilitates the process of acquiring and refining the desired behavior. The dynamical systems-based motor primitives (DMPs) [6] are a popular choice for representing motor skills in robotics as they are invariant under rescaling of both duration, movement amplitude, and changes of the final position. Furthermore, DMPs guarantee the stability of the movement generation. In [7], we have introduced a modification of the DMPs that allows to generalize them to arbitrary velocities and positions at the end of the movement. Moreover, DMPs are linear in parameters and, hence, it is straightforward to learn a movement by imitation learning

*The project receives funding from the European Community’s Seventh Framework Programme under grant agreement no. ICT-248273 GeRT. The project receives funding from the European Community’s Seventh Framework Programme under grant agreement no. ICT-270327 CompLACS.

All authors are with the Max Planck Institute for Intelligent Systems, Department Empirical Inference & Department Autonomous Motion, Spemannstr. 38, 72076 Tübingen, Germany and the Technische Universität Darmstadt, Intelligent Autonomous Systems Group, Hochschulstr. 10, 64289 Darmstadt, Germany {firstname.lastname}@tuebingen.mpg.de



(a) Kinesthetic teach-in of the hitting skill for imitation learning.

(b) The final table tennis skill.

Fig. 1: Learning hitting skills.

[6] and to refine it by reinforcement learning [8]. The video and Fig. 1a show a demonstration of table tennis strokes by kinesthetic teach-in (i.e., by taking the racket by the hand) that is used to learn the hitting skill by imitation learning.

III. LEARNING HITTING SKILLS

Reinforcement learning can be employed to refine the hitting skill acquired by imitation learning according to a cost function (e.g., to minimize the required torques), see [8]. A skill learned using these techniques is only applicable if the ball is always played exactly to the same location. In order to be able to deal with different ball trajectories, we now need to learn how to take advantage of the generalization properties of the movement primitives, i.e., we need to learn a mapping from the ball trajectory to the required hitting point and velocity. In [9], we proposed a non-parametric approach to learn how to adapt motor primitives to such new situations. Not every hitting movement is equally well-suited for all locations. Hitting movements can be generalized with a Mixture of Motor Primitives [10], i.e., by selecting, weighting and combining appropriate motor primitives according to the desired hitting point and velocity. Fig. 1b illustrates the resulting table tennis skill.

IV. GENERALIZING TO CATCHING SKILLS

A catching skill can be seen as a sub-set of hitting skills, if we ignore the mechanism for retaining the ball. Both skills require the end-effector to intersect the trajectory of the ball, however for the catching skill does not require a pre-determined velocity at the contact point. For the catching skill we replaced the table tennis racket by a net of similar size. In contrast to the hitting movement a catching skill does not need to take into account where the ball should be returned on the opponent’s side of the table. The ball is

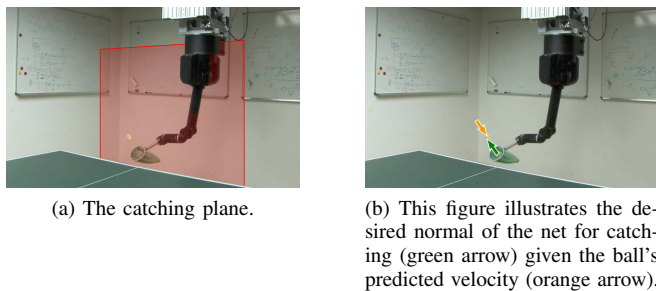


Fig. 2: Generalizing hitting skills to a catching skill.




Fig. 3: The throwing setup. The robot learns to hit the targets.

always caught on a vertical plane, called the catching plane¹, illustrated in Fig. 2a. We predict the intersection of the ball trajectory and the hitting plane, as well as its timing. The net is moved in such a way that it reaches the intersection point at the predicted catching time. The orientation of the opening of the net is chosen such that it is maximal with respect to the ball path. This can be achieved by having the normal of the net parallel to the velocity of the ball at the predicted catching point (see Fig. 2b for an illustration).

V. LEARNING THROWING SKILLS

In [11], we show how a throwing skill can be learned jointly with a higher level goal. The robot learns to throw at targets while simultaneously playing a blackjack inspired game. The actual learning of the throwing skill is also based on the idea of adapting a previously acquired skill to new targets [9]. See Fig. 3 for an illustration of the throwing setup.

VI. COMBINING CATCHING AND THROWING SKILLS

It is straightforward to combine the catching and throwing skills to have the BioRob and the Barrett WAM playing catch. The throwing robot learns to throw into the reachable workspace of the catching robot, i.e., within a diameter of approximately 1m. The catching robot did not need to adapt its catching. The robot tandem managed to reliably throw and catch ten consecutive balls until we had to stop the experiment in order to empty the net.

VII. VISION SYSTEM

The ball is tracked by two overlapping high speed stereo vision setups with 200Hz cameras. The position of the ball on

¹Instead of employing the concept of a catching plane, more advanced criteria like “minimal movement of the end-effector” could be employed to determine the catching point [4].




Fig. 4: The BioRob and the Barrett WAM playing catch.

the frames is determined by a GPU-based blob tracker [12]. In order to obtain better estimates of the current position and to calculate the velocities, the raw 3D positions are filtered by an extended Kalman filter [13] that takes contacts of the ball with the table and the racket into account.

VIII. SUMMARY

In this video we have shown how catching skills can be learned as a generalization of hitting skills using a combination of imitation learning and reinforcement learning employed to refine the movement, generalize it to new situations, and to select appropriate primitives. Throwing at different target can also be learned by generalizing a throwing skill with a fixed target. Combining the catching and throwing skills enables a BioRob to throw a ball that is caught by a Barrett WAM.

REFERENCES

- [1] E. W. Aboaf, C. G. Atkeson, and D. J. Reinkensmeyer, “Task-level robot learning: Ball throwing,” MIT, Tech. Rep., 1987.
- [2] B. Hove and J.-J. E. Slotine, “Experiments in robotic catching,” in *American Control Conference*, 1991.
- [3] M. Riley and C. G. Atkeson, “Robot catching: Towards engaging human-humanoid interaction,” *Auton. Robots*, vol. 12, no. 1, pp. 119–128, 2002.
- [4] B. Bäuml, T. Wimböck, and G. Hirzinger, “Kinematically optimal catching a flying ball with a hand-arm-system,” in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2010, pp. 2592–2599.
- [5] A. Ude, A. Gams, T. Asfour, and J. Morimoto, “Task-specific generalization of discrete and periodic dynamic movement primitives,” *IEEE Transactions on Robotics*, vol. 26, no. 5, pp. 800–815, 2010.
- [6] A. J. Ijspeert, J. Nakanishi, and S. Schaal, “Learning attractor landscapes for learning motor primitives,” in *Advances in Neural Information Processing Systems 15 (NIPS)*, 2002, pp. 1523–1530.
- [7] J. Kober, K. Muelling, O. Kroemer, C. H. Lampert, B. Schölkopf, and J. Peters, “Movement templates for learning of hitting and batting,” in *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 2010, pp. 853–858.
- [8] J. Kober and J. Peters, “Policy search for motor primitives in robotics,” *Machine Learning*, vol. 84, no. 1-2, pp. 171–203, 2011.
- [9] J. Kober, E. Oztop, and J. Peters, “Reinforcement learning to adjust robot movements to new situations,” in *Proceedings of Robotics: Science and Systems Conference (R:SS)*, 2010.
- [10] K. Muelling, J. Kober, and J. Peters, “Learning table tennis with a mixture of motor primitives,” in *Proceedings of the 10th IEEE-RAS International Conference on Humanoid Robots (Humanoids)*, 2010, pp. 411–416.
- [11] J. Kober and J. Peters, “Learning elementary movements jointly with a higher level task,” in *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2011, pp. 338–343.
- [12] C. H. Lampert and J. Peters, “Real-time detection of colored objects in multiple camera streams with off-the-shelf hardware components,” *Journal of Real-Time Image Processing*, vol. 7, no. 1, pp. 31–41, 2012.
- [13] R. Kalman, “A new approach to linear filtering and prediction problems,” *Transactions of the ASME—Journal of Basic Engineering*, vol. 82, no. Series D, pp. 35–45, 1960.