

Starting from theoretically sound robotic control structures for task representation and execution, we replace analytic modules by more flexible learned ones [13]. To this end, we tackle problems such as accurate but compliant execution, learning of elementary behaviors, hierarchical composition of behaviors, and parsing complex demonstrations into elementary behaviors.

Accurate execution of movements ideally requires only low-gain controls, such that the robot can accomplish tasks without harming humans. Following a trajectory with little feedback requires accurate prediction of needed torques, which often cannot be achieved using classical methods. However, learning such models is hard: the joint-space can never be fully explored and the learning algorithm has to cope with a data stream in real time. We have developed learning methods for tasks represented in joint-space [12] or task-space [5].

Executing a task is important, but often the task itself needs to be learned. We focused on learning elementary tasks or movement primitives, which are parameterized nonlinear differential equations with desired attractor properties. We mimic how children learn new motor tasks, using imitation to initialize these movement primitives, and reinforcement learning to subsequently improve performance. We have learned tasks such as Ball-in-a-Cup or bouncing a ball [1, 2].

For more complex tasks, hierarchical solutions compose behaviors based on a large number of elementary ones. As the ‘drosophila’ of complex behavior, we chose the task of returning table tennis balls over the net. This requires all the methods described in the previous paragraphs, as well as forms of reinforcement learning discussed below. We created a compiler that segments movements of a human teacher into elementary movements [7, 6]. These then train the single primitives discussed above [1, 2]. Novel behaviors, modulated by the opponent’s incoming ball, are composed by mixing motor primitives [11]. The learning system then generalizes between the primitives. This created successful returns of 88% of balls. Further improvement is limited by the robot’s hardware and reaction time.

Human players infer the direction of incoming balls from the opponent’s movement, and so can prepare their stroke before the opponent even hits the ball. Inspired by this, we developed learning methods anticipating the aim of human opponent. They achieved a prediction accuracy of $35cm$, already $320ms$ before the ball is hit

[14], significantly increasing available reaction time.



Figure 0.1: The robot table tennis setup [11].

Recently, we developed a method maintaining multiple conflicting strategies in its motor policy [8]. This generalizes previous work on mixtures of motor primitives [11] and puts it on stronger theoretical footing. It was tested in a game of tetherball (Figure 0.2), and may become useful in table tennis as it allows more ambiguous gameplay.



Figure 0.2: The Tetherball task [8], which won the IROS Best Cognitive Robotics Paper Award while being finalist for both the best paper and the best student paper award.

Motor skill learning can also be applied to grasping and manipulation. We have developed methods [9, 3] learning how to grasp novel objects (a preliminary version [9] won the ICINCO Best Paper Award) and generalize such basic manipulations to different objects [10]. We also studied how to recognize surfaces from tactile frequency patterns obtained by sliding a finger over surfaces [4]. The robot autonomously discovered the most relevant dimensions of the tactile data by training jointly on vision and tactile data from various textured surfaces. Subsequently, when presented with only tactile stimuli, it was able to quickly and reliably recover surface type.

Publications

Journal Articles

- [1] J Kober and J Peters. Imitation and reinforcement learning. *IEEE Robotics and Automation Magazine*, 17(2):55–62, 6 2010. [1](#)
- [2] J Kober and J Peters. Policy search for motor primitives in robotics. *Machine Learning*, 84(1-2):171–203, 7 2011. [1](#)
- [3] O Kroemer, R Detry, J Piater, and J Peters. Combining active learning and reactive control for robot grasping. *Robotics and Autonomous Systems*, 58(9):1105–1116, 9 2010. [1](#)
- [4] O Kroemer, CH Lampert, and J Peters. Learning dynamic tactile sensing with robust vision-based training. *IEEE Transactions on Robotics*, 27(3):545–557, 6 2011. [1](#)
- [5] D Nguyen-Tuong and J Peters. Online kernel-based learning for task-space tracking robot control. *IEEE Transactions on Neural Networks and Learning Systems*, 5 2012. [1](#)

Articles in Conference Proceedings

- [6] MA Alvarez, J Peters, B Schölkopf, and ND Lawrence. Switched latent force models for movement segmentation. In J Lafferty, CKI Williams, J Shawe-Taylor, RS Zemel, and A Culotta, editors, *24th Annual Conference on Neural Information Processing Systems (NIPS)*, pages 55–63, Vancouver, BC, Canada, 2010. Curran. [1](#)
- [7] S Chiappa and J Peters. Movement extraction by detecting dynamics switches and repetitions. In J Lafferty, C KI Williams, J Shawe-Taylor, R S Zemel, and A Culotta, editors, *24th Annual Conference on Neural Information Processing Systems (NIPS)*, pages 388–396, Vancouver, BC, Canada, 2010. Curran. [1](#)
- [8] C Daniel, G Neumann, and J Peters. Learning concurrent motor skills in versatile solution spaces. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2012)*, Vilamoura, Portugal, 2012. [1](#)
- [9] O Kroemer, R Detry, J Piater, and J Peters. Grasping with vision descriptors and motor primitives. In J Filipe, J Andrade-Cetto, and J-L Ferrier, editors, *7th International Conference on Informatics in Control, Automation and Robotics (ICINCO 2010)*, pages 47–54, Funchal, Madeira, Portugal, 6 2010. SciTePress. [1](#)
- [10] O Kroemer, E Ugur, E Oztog, and J Peters. A kernel-based approach to direct action perception. In *IEEE International Conference on Robotics and Automation (ICRA 2012)*, pages 1–6, St. Paul, MN, USA, 5 2012. [1](#)
- [11] K Mülling, J Kober, and J Peters. Learning table tennis with a mixture of motor primitives. In *10th IEEE-RAS International Conference on Humanoid Robots (Humanoids 2010)*, pages 411–416, Nashville, TN, USA, 12 2010. IEEE. [1](#)
- [12] D Nguyen-Tuong and J Peters. Incremental sparsification for real-time online model learning. In YW Teh and M Titterton, editors, *13th International Conference on Artificial Intelligence and Statistics (AISTATS)*, pages 557–564, Chia Laguna Resort, Italy, 5 2010. JMLR. [1](#)
- [13] J Peters, J Kober, K Mülling, D Nguyen-Tuong, and O Kroemer. Robot skill learning. In *20th European Conference on Artificial Intelligence (ECAI 2012)*, pages 1–6, Montpellier, France, 2012. [1](#)
- [14] Z Wang, M Deisenroth, H Ben Amor, D Vogt, B Schölkopf, and J Peters. Probabilistic modeling of human movements for intention inference. In *Robotics: Science and Systems (R:SS 2012)*, Sydney, Australia, 7 2012. [1](#)