Probabilistic Movement Primitives

Alexandros Paraschos
TU Darmstadt
Email: paraschos@ias.tu-darmstadt.de

Gerhard Neumann
TU Darmstadt
Email: neumann@ias.tu-darmstadt.de

Jan Peters
TU Darmstadt
Email: peters@ias.tu-darmstadt.de

***Abstract*—Movement primitives are a promising approach for modular and re-usable movement generation, and suitable for data-driven movement acquisition. Beneficial properties such as simultaneous activation of multiple primitives, optimal movement encoding for stochastic systems, and generalization to new targets, are absent in most common approaches. We propose a probabilistic approach for generating, learning, and re-using movement primitives that overcomes these limitations. We represent a movement primitive as a probability distribution over trajectories. As a consequence, we can activate primitives simultaneously, smoothly blend together, generalize to new target states and encode optimal trajectories in stochastic systems. We compare our approach to the existing state-of-the art and present real robot results for learning from demonstration.**

Movement primitives (MP) are considered to be a state of the art approach for learning robot movement generation. The most commonly used MP representations use time- or phase-dependent policies [1], like the widely used dynamic movement primitive (DMP) approach [1]. DMPs are based on second-order dynamical systems, which are composed of a linear spring-damper system and a learnable non-linear forcing function. Integrating the dynamic system results in the desired trajectory which is followed by feedback control laws.

However, there are properties for a MP representation that are not fulfilled by the DMP approach. Most importantly, multiple DMPs for the same degrees of freedom (DoF) cannot be activated simultaneously without further considerations on prioritized control. Additionally, a DMP cannot encode optimal movement generation in the presence of stochasticity. For stochastic systems, following a single trajectory is always sub-optimal [2] as one cannot adapt the variance of the resulting trajectory distribution. DMPs can be efficiently used to imitate a single trajectory, however, how to generalize the shape of a DMP from multiple demonstrations is an open problem.

We propose a probabilistic approach to movement primitives which we call Probabilistic Movement Primitives (ProMP). ProMPs are represented as a distribution over trajectories $p(\boldsymbol{\tau})$. They do not require a parametric representation of the control policy, and a stochastic feedback controller that exactly follows the given distribution is obtained in closed form. Using trajectory distributions allows us to encode optimal behaviors for stochastic systems as well. Trajectory distributions can also be easily combined by calculating the "intersection" of two distributions as a product of them. This operation allows for a simultaneous activation of several primitives or a smooth switching from one activated primitive to the next, as illustrated in Figure 1. We can also condition $p(\boldsymbol{\tau})$ to reach a desired position or velocity at any point in time as long as the position lies within the distribution. The distribution adapts to the new desired position while simultaneously trying to stay close to the demonstrations. As a result, generalization to new targets or via-points is learned from demonstrations. The generalization to different desired target positions is shown in Figure 1. Our MPs can be easily obtained from imitation, they can be used for both point-to-point and rhythmic movements, and the speed of the movement can be adapted by replacing time by a phase variable. We summarize the basic properties of the ProMPs in comparison to the DMPs in Table I. We evaluated our approach on two real robot tasks with a 7-DoF humanoid arm.

Table I
QUALITATIVE COMPARISON OF PROMPS AND DMPS.

Property	DMPs	ProMPs
Co-Activation	-	product of $p^{[i]}(\boldsymbol{\tau})$
Modulate final positions	heuristics	conditioning
Modulate final velocities	heuristics	conditioning
Modulate via-points	-	conditioning
Optimality	det. systems	stoch. systems
Coupling	mean	mean, covariance
Learning	accelerations	position, velocities
Temporal Scaling	+	+
Rhythmic Movements	+	+

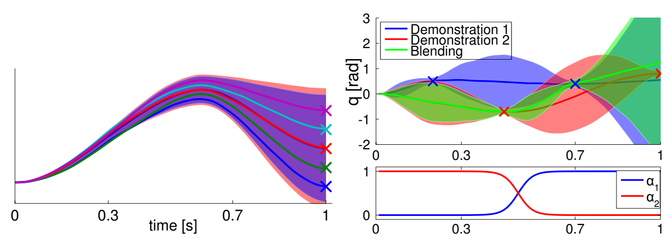

(a) Conditioning (b) Blending

Figure 1. (a) Reproduced generalization to different target states by conditioning. (b) Blending (green) of movement primitives. In the beginning, the red primitive is active. At $t = 0.5s$, we smoothly switch the activation to the blue primitive and the red primitive is ignored.

REFERENCES

[1] Auke Jan Ijspeert and Stefan Schaal. Learning Attractor Landscapes for Learning Motor Primitives. In *Advances in Neural Information Processing Systems 15*, (NIPS). MIT Press, Cambridge, MA, 2003.

[2] E. Todorov and M. Jordan. Optimal feedback control as a theory of motor coordination. *Nature Neuroscience*, 5: 1226–1235, 2002.