Movement generation by learning from demonstration and generalizing to new targets

Peter Pastor, Heiko Hoffmann, and Stefan Schaal
University of Southern California, Los Angeles, USA

We provide a general approach for learning robotic movements from human demonstration. To represent a recorded movement, a non-linear differential equation is adapted such that it reproduces this movement. Based on this representation, we build a library of movements by labeling each recorded movement according to task and context (e.g., grasping, placing, and releasing). Our differential equation is formed such that generalization can be achieved simply by adapting a start and a goal parameter in the equation to the desired position values of a movement. The feasibility of our approach is demonstrated with the Sarcos slave robot arm; the robot pours water into several cups after we demonstrated the movement for one cup.

Humanoid robots assisting humans can become widespread only if they are easy to program. Easy programming might be achieved through learning from demonstration [1]. A human movement is recorded and later reproduced by a robot. Three challenges need to be mastered for this imitation: the correspondence problem, generalization, and robustness against perturbation.

The correspondence problem means that links and joints between human and robot may not match. Generalization is required because we cannot demonstrate every single movement that the robot is supposed to make. Learning by demonstration is feasible only if a demonstrated movement can be generalized to other contexts, like different goal positions. Finally, we need robustness against perturbation. Replaying exactly an observed movement is unrealistic in a dynamic environment, in which obstacles may appear suddenly.

To address these issues, we present a model that is based on the dynamic movement primitive (DMP) framework [2, 3]. In this framework, any recorded movement can be represented with a set of differential equations. Representing a movement with a differential equation has the advantage that a perturbance can be automatically corrected for by the dynamics of the system; this behavior addresses the above mentioned robustness. Furthermore, the equations are formulated in a way that adaptation to a new goal is achieved by simply changing a goal parameter. This characteristic allows generalization. Here, we will present a new version of the dynamic equations with improved adaptation to goal changes.

In the present work, we use the dynamic movement primitives to represent a movement trajectory in end-effector space; thus, we address the above-mentioned correspondence problem. In our robot demonstration, we use standard inverse kinematics to map the end-effector position and gripper orientation onto the appropriate joint angles.

To deal with complex motion, the above framework can be used to build a library of movement primitives out of which the complex motion can be composed by sequencing. For example, the library may contain a grasping, placing, and releasing motion. Each of these movements is recorded from a human demonstrator, represented by a differential equation, and labeled accordingly. For example, to move an object on a table, a grasping-placing-releasing sequence is required, and the corresponding primitives are recalled from the library. Due to the generalization ability of each dynamic movement primitive, an object may be placed between two arbitrary positions on the table based solely on the three demonstrated movements.

Dynamic movement primitives

Dynamic movement primitives can be used to generate discrete and rhythmic movements [2, 3]. Here, we focus on discrete movements and present a new variant of the equations. A movement is generated by integrating the following set of differential equations (which we will refer to as 'transformation system'):

$$\begin{aligned} \tau \dot{v} &= K(g-x) - Dv - K(g-x_0)\theta + Kf(\theta) \quad (1) \\ \tau \dot{x} &= v \ , \quad (2) \end{aligned}$$

where x and v are position and velocity of the system; x_0 and g are the start and goal position; τ is a temporal scaling factor; K and D are constants; D is chosen such that the system is critically damped, and f is a non-linear function that can be adapted to allow the generation of arbitrary complex movements. Equation (1) is motivated from human behavioral data and leg force fields observed in frog after stimulating the spinal cord [4].

The non-linear function is defined as

$$f(\theta) = \frac{\sum_i w_i \psi_i(\theta)}{\sum_i \psi_i(\theta)} \theta \ , \quad (3)$$

where ψ_i are Gaussian basis functions, $\psi_i(\theta) = \exp(-h_i(\theta - c_i)^2)$ with center c_i and width h_i, and w_i are adjustable weights. The function f does not directly depend on time; instead, it depends on a phase variable θ, which goes from 1 towards 0 during a movement and is obtained by the equation

$$\tau \dot{\theta} = -\alpha \theta \ . \quad (4)$$

where α is a pre-defined constant.

To learn a movement from demonstration, first, a movement $x(t)$ is recorded and its derivatives $v(t)$ and $\dot{v}(t)$ are computed for each time step t. Second, $f(t)$ is computed based on (1). Third, (4) is integrated and $\theta(t)$ evaluated. Using these arrays, we find the weights w_i in

(3) by linear regression, which can be solved efficiently. For many-dimensional movements, we have a transformation system for each dimension and, for each, learn the weights separately. A learned movement can be generalized to new targets by changing the variable g. This goal may change also online during a movement.

Different from previous formulations of DMP is the transformation system (1). The new formulation fixes a problem with the original DMP: if start and goal position, x_0 and g, of a movement were the same, then the system remained at x_0. Furthermore, if $g - x_0$ were close to zero, a small change in g would lead to huge accelerations, which can break the limits of a robot (Fig. 1). The modified form solves these problems.

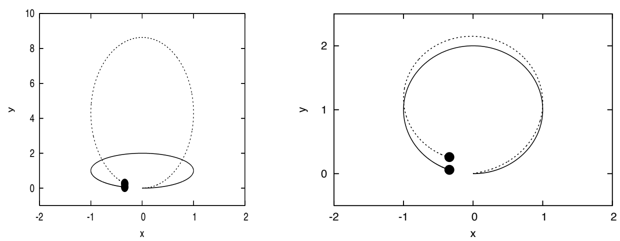

Fig. 1: Comparison of goal adaptation between old (Left) and new (Right) DMP formulation in operational space (x, y). The same original movement (solid line) and goals are used for both formulations. The dashed lines show the results of changing the goal g for the entire movement.

Robot experiment

We demonstrate the utility of our framework in a robot demonstration of serving water (Fig. 2). The single parts of the movement (grasping, pouring, retreating bottle, and releasing) are pre-recorded from a human demonstrator. Afterwards, each of these parts is represented by our differential equation. The equation describes the end-effector position and orientation of the gripper. By appropriately sequencing these movement parts, the robot served water into three cups. It could generalize to different cup positions simply through changing the goal position and orientation for the pouring-movement primitive.

Conclusions

This article extends the approach of dynamic movement primitives to sequential movements, task-space control, and improved generalization to new goals. Semantic information was added to the movement primitives, such that they encode object-oriented action. We demonstrated the feasibility of our approach in an imitation learning setting, where a robot learned to serve water and could generalize this task to novel situations.

References

[1] C. G. Atkeson and S. Schaal, "Robot learning from demonstration," in *International Conference on Machine Learning*, 1997.

[2] A. J. Ijspeert, J. Nakanishi, and S. Schaal, "Learning attractor landscapes for learning motor primitives." in *Advances in Neural Information Processing Systems*, S. Becker, S. Thrun, and K. Obermayer, Eds., vol. 15. MIT Press, Cambridge, MA, 2003, pp. 1523–1530.

[3] S. Schaal, "Dynamic movement primitives: A framework for motor control in humans and humanoid robotics," in *2nd International Symposium on Adaptive Motion of Animals and Machines (AMAM)*, 2003.

[4] H. Hoffmann and S. Schaal, "Human movement generation based on convergent flow fields: a computational model and a behavioral experiment," in *Advances in Computational Motor Control VI*, R. Shadmehr and E. Todorov, Eds., San Diego, CA, 2007.

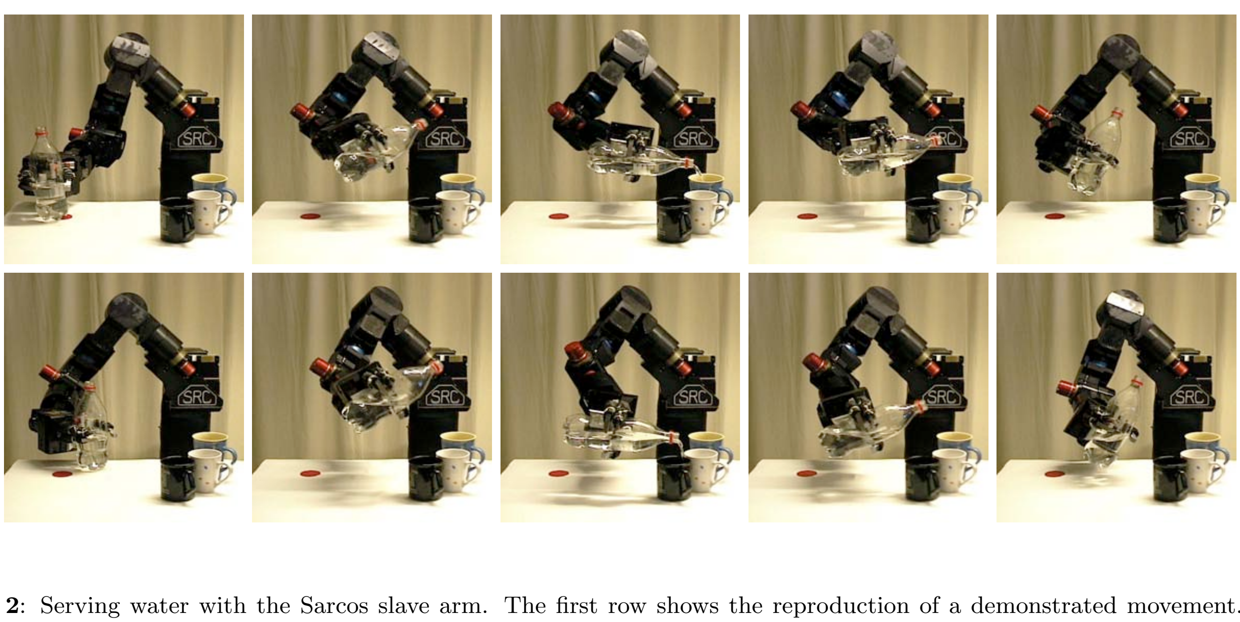

Fig. 2: Serving water with the Sarcos slave arm. The first row shows the reproduction of a demonstrated movement. The second row shows the generalization to a new cup position.