
Reinforcement Learning by Reward-weighted Regression for Operational Space Control

Jan Peters

Max-Planck Institute for Biological Cybernetics, 72074 Tuebingen, Germany

JAN.PETERS@TUEBINGEN.MPG.DE

Stefan Schaal

University of Southern California, Los Angeles, CA 90089, USA

SSCHAAL@USC.EDU

Abstract

Many robot control problems of practical importance, including operational space control, can be reformulated as immediate reward reinforcement learning problems. However, few of the known optimization or reinforcement learning algorithms can be used in online learning control for robots, as they are either prohibitively slow, do not scale to interesting domains of complex robots, or require trying out policies generated by random search, which are infeasible for a physical system. Using a generalization of the EM-base reinforcement learning framework suggested by Dayan & Hinton, we reduce the problem of learning with immediate rewards to a reward-weighted regression problem with an adaptive, integrated reward transformation for faster convergence. The resulting algorithm is efficient, learns smoothly without dangerous jumps in solution space, and works well in applications of complex high degree-of-freedom robots.

1. Introduction

In learning control with complex robotic systems like manipulator robots and anthropomorphic robots, the control task is usually defined in external space, e.g., the Cartesian coordinates of our 3D world. In contrast, motor commands need to be generated at the level of the actuators, i.e., the space internal to the robot, as for instance joint angles. In such scenarios, the goal of learning control is to find the motor commands that accomplish the task while simultaneously recruiting all degrees-of-freedom of the robot in an optimal way, e.g.,

Appearing in *Proceedings of the 24th International Conference on Machine Learning*, Corvallis, OR, 2007. Copyright 2007 by the author(s)/owner(s).

by minimizing the amount of movement, the amount of energy consumption, the jerkiness of movement. The most advanced analytical framework to formulate such control problems is called operational space control (Khatib, 1987). Let us assume the robot dynamics can be modeled as a rigid body dynamics system

$$\mathbf{M}(\mathbf{q})\ddot{\mathbf{q}} = \mathbf{F}(\mathbf{q}, \dot{\mathbf{q}}, t) + \mathbf{a}, \quad (1)$$

where $\mathbf{q}, \dot{\mathbf{q}}, \ddot{\mathbf{q}}$ are the generalized positions, velocities, and accelerations of the robot, \mathbf{a} denotes our motor commands, $\mathbf{M}(\mathbf{q})$ the inertia matrix of the robot, and $\mathbf{F}(\mathbf{q}, \dot{\mathbf{q}}, t)$ all the (potentially time t dependent) forces acting on the system (e.g., coriolis forces, centripetal force, gravity, friction). In external space, the task is defined by a desired trajectory $\mathbf{x}_d(t)$, $\dot{\mathbf{x}}_d(t)$, $\ddot{\mathbf{x}}_d(t)$, which is converted into some reference dynamics $\ddot{\mathbf{x}}_{\text{ref}} = \ddot{\mathbf{x}}_d + K_D(\dot{\mathbf{x}}_d - \dot{\mathbf{x}}) + K_P(\mathbf{x}_d - \mathbf{x})$; here $\mathbf{x} = \mathbf{x}(\mathbf{q})$ and $\dot{\mathbf{x}} = \dot{\mathbf{x}}(\mathbf{q}, \dot{\mathbf{q}}) = \mathbf{J}(\mathbf{q})\dot{\mathbf{q}}$ denote the forward kinematics and differential kinematics of the robot endeffector, respectively, and \mathbf{J} the Jacobian of the endeffector kinematics. The resulting control law of such a problem, for instance, typically look like

$$\mathbf{a} = \mathbf{J}^T(\mathbf{J}\mathbf{M}^{-1}\mathbf{J}^T)^{-1}(\ddot{\mathbf{x}}_{\text{ref}} - \dot{\mathbf{J}}\dot{\mathbf{q}} + \mathbf{J}\mathbf{M}^{-1}\mathbf{F}). \quad (2)$$

Operational space control has numerous advantageous features (Khatib, 1987); among the most important ones is that it allows very compliant control of robots, which is one of the key ingredients that will allow robots to become a non-dangerous part of daily human life. However, the drawback of operational space control, which, so far, has prevented its widespread application in robotics, is that the control law is very sensitive to the quality of system identification, i.e., the accurate knowledge of terms like \mathbf{M} , \mathbf{F} , and \mathbf{J} . Small errors in these terms can result in very undesirable behavior and even lead to instability of the entire control system (Nakanishi et al., 2005).

The goal of this paper is to suggest a general learning solution for operational space control, i.e., a method

that does not require prior knowledge of any of the robot’s kinematics and dynamics parameters. For this purpose, first, the next section will discuss a recent insight that many operational space control laws can be derived from a constrained optimal control framework with quadratic cost function on the motor commands. Second, this insight allows us to view learning operational space control as an immediate reward reinforcement learning problem. The EM-based framework by Dayan & Hinton (Dayan & Hinton, 1997) provides a suitable starting point for deriving appropriate learning algorithms. Interestingly, reinforcement learning can be reduced to a reward-weighted nonlinear regression problem in this context, which greatly accelerates the speed of learning and may have a more general applicability in reinforcement learning. Finally, we evaluate our approach on simulated anthropomorphic robot arms and compare the results of learning with analytical solutions.

2. Operational Space Control as an Optimal Control Problem

In order to be able to address operational space control as a learning problem, an associated cost function is required that, when optimized, results in an appropriate operational space controller, e.g., Eq. (2). A key insight of recent work (Peters et al., 2005) was that a large class of operational space control laws can be derived as the solution of a constrained immediate reward optimization problem:

$$\begin{aligned} r(\mathbf{a}) &= -(\mathbf{a} - \mathbf{a}_0)^T \mathbf{N}(\mathbf{a} - \mathbf{a}_0) \\ \text{s.t. } \mathbf{J}\dot{\mathbf{q}} &= \ddot{\mathbf{x}}_{\text{ref}} - \dot{\mathbf{J}}\dot{\mathbf{q}}, \end{aligned} \quad (3)$$

where \mathbf{a}_0 is the default or nominal stabilizing behavior, e.g., a force which pulls the robot towards a static rest posture $\mathbf{a}_0 = -\mathbf{K}_D\dot{\mathbf{q}} - \mathbf{K}_P(\mathbf{q} - \mathbf{q}_{\text{rest}})$, and $\mathbf{a}_1 = \mathbf{a} - \mathbf{a}_0$ corresponds to the control signal that accomplishes the desired task goal characterized by $\ddot{\mathbf{x}}_{\text{ref}}$. \mathbf{N} is a positive definite metric that decides the relative importance of the motor commands in the optimization. If an accurate model is available, the general analytic solution to this optimization problem is given by

$$\begin{aligned} \mathbf{a} &= \mathbf{N}^{-\frac{1}{2}}(\mathbf{J}\mathbf{M}^{-1}\mathbf{N}^{-\frac{1}{2}})^+(\ddot{\mathbf{x}}_{\text{ref}} - \dot{\mathbf{J}}\dot{\mathbf{q}} + \mathbf{J}\mathbf{M}^{-1}\mathbf{F}) \\ &+ \mathbf{N}^{-\frac{1}{2}}(\mathbf{I} - (\mathbf{N}^{-\frac{1}{2}}\mathbf{M}^{-1}\mathbf{J}^T)(\mathbf{J}\mathbf{M}^{-1}\mathbf{N}^{-\frac{1}{2}})^+)\mathbf{N}^{\frac{1}{2}}\mathbf{a}_0, \end{aligned} \quad (4)$$

where the second summand fulfills the nominal control law \mathbf{a}_0 in the null-space of the first term. For example, Eq. (2) is derived for $\mathbf{N} = \mathbf{M}^{-1}$ and $\mathbf{a}_0 = \mathbf{0}$. Of course, this solution is only interesting when the dimensionality of $\ddot{\mathbf{x}}_{\text{ref}}$ is smaller than that of \mathbf{a} , which we assume in all the following.

The conclusion of this brief summary is that operational space control can be viewed as an immediate reward reinforcement learning problem (Kaelbling et al., 1996) with high-dimensional, continuous states $\mathbf{s} = [\mathbf{q}, \dot{\mathbf{q}}, \ddot{\mathbf{x}}_{\text{ref}}, \mathbf{a}_0] \in \mathbb{R}^n$ and actions $\mathbf{a} \in \mathbb{R}^m$. The goal of learning is to obtain an optimal policy

$$\mathbf{a} = \mu(\mathbf{q}, \dot{\mathbf{q}}, \ddot{\mathbf{x}}_{\text{ref}}, \mathbf{a}_0) = \mu(\mathbf{s}) \quad (5)$$

such that the system follows the reference acceleration $\ddot{\mathbf{x}}_{\text{ref}}$ while maximizing the immediate reward $r(\mathbf{a}) = -(\mathbf{a} - \mathbf{a}_0)^T \mathbf{N}(\mathbf{a} - \mathbf{a}_0)$ for any given nominal behavior \mathbf{a}_0 . In order to incorporate exploration during learning, we consider a stochastic control policy $\mathbf{a} = \mu_\theta(\mathbf{q}, \dot{\mathbf{q}}, \ddot{\mathbf{x}}_{\text{ref}}) + \varepsilon$, modeled as a probability distribution $\pi_\theta(\mathbf{a}|\mathbf{s}) = p(\mathbf{a}|\mathbf{s}, \theta)$ with parameter vector θ . The goal of the learning system is thus to find the policy parameters θ that maximize

$$J_r(\theta) = \int p(\mathbf{s}) \int \pi_\theta(\mathbf{a}|\mathbf{s}) r(\mathbf{s}, \mathbf{a}) d\mathbf{a} d\mathbf{s}. \quad (6)$$

$p(\mathbf{s})$ denotes the distribution of states, which is treated as fixed in immediate reward reinforcement learning problems (Kaelbling et al., 1996).

3. Reinforcement Learning by Reward-Weighted Regression

Previous work in the literature suggested a variety of optimizing methods, which can be applied to immediate reward reinforcement learning problems, e.g., gradient based methods (e.g., REINFORCE, Covariant REINFORCE, finite difference gradients, the Kiefer-Wolfowitz procedure, A_{RP} algorithms, CRBP, and many others) and random search algorithms (e.g., simulated annealing or genetic algorithms) (Dayan & Hinton, 1997; Kaelbling et al., 1996; Spall, 2003). However, gradient based methods tend to be too slow for the online learning that we desire in our problem, while randomized search algorithms can create too arbitrary solutions, often not suitable for execution on a robotic system. For learning operational space control, we require a method that is computationally sufficiently efficient to deal with high-dimensional robot systems and large amounts of data, that comes with convergence guarantees, and that is suitable for smooth online improvement. For instance, linear regression techniques and/or methods employing EM-style algorithms are highly desirable.

A good starting point our work is the probabilistic reinforcement learning framework by Dayan & Hinton (Dayan & Hinton, 1997). As we will show in the following, a generalization of this approach allows us to derive an EM-algorithm that essentially reduces the im-

mediate reward learning problem to a reward-weighted regression problem.

3.1. Reward Transformation

In order to maximize the expected return (Eq. 6) using samples, we approximate

$$J_r(\theta) \approx \sum_{i=1}^n \pi_{\theta}(\mathbf{a}_i | \mathbf{s}_i) r_i \quad (7)$$

where $r_i = r(\mathbf{s}_i, \mathbf{a}_i)$. For application of the probabilistic reinforcement learning framework of Dayan & Hinton (Dayan & Hinton, 1997), the reward needs to be strictly positive such that it resembles an (improper) probability distribution. While this can be achieved by linear rescaling for problems for bounded rewards, for unbounded rewards as discussed in this paper, a nonlinear transformation of the reward $u_{\tau}(r)$ is required with the constraint that the optimal solution to the underlying problem remains unchanged. Thus, we require that $u_{\tau}(r)$ is strictly monotonic with respect to r , and additionally that $u_{\tau}(r) \geq 0$ and $\int_0^{\infty} u_{\tau}(r) dr = \text{const}$, resulting in the transformed optimization problem

$$J_u(\theta) = \sum_{i=1}^n \pi_{\theta}(\mathbf{a}_i | \mathbf{s}_i) u_{\tau}(r_i). \quad (8)$$

The reward transformation plays a more important role than initially meets the eye: as pointed out in (Dayan & Hinton, 1997), convergence speed can be greatly affected by this transformation. Making $u_{\tau}(r)$ an adaptive part of the learning algorithm by means of some internal parameters τ can greatly accelerate the learning speed and help avoiding local minima during learning. Figure 1 demonstrates this issue with a 1D continuous state and 1D continuous action example, where the goal is to learn an optimal linear policy. Using the algorithm that we will introduce below, an adaptive reward transformation accelerated the convergence by a factor of 4, and actually significantly help avoiding local minima during learning.

3.2. EM-Reinforcement Learning with Adaptive Reward Transformation

To derive our learning algorithm, similar as in (Dayan & Hinton, 1997), we start by establishing the lower

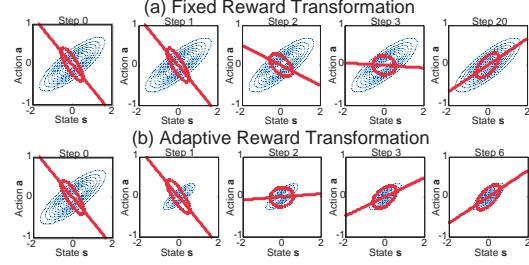


Figure 1. A comparison of fixed and adaptive reward transformation for learning a linear policy $\pi(a|s) = N(a|\theta_1 s + \theta_2, \sigma^2)$ under the transformed reward $u(r(s, a)) = \exp(-\tau(q_1 a^2 + q_2 a s + s q_3^2))$. The transformed reward is indicated by the dotted blue ellipses, the variance of the action distribution is indicated by the red thick ellipse, and the mean of the linear policy is shown by the red thick line. With τ being adaptive, significantly faster learning of the optimal policy is achieved. Step 0 shows the initial policy and initial transformed reward, Step 1 shows the initial policy with adapted transformed reward.

bound

$$\begin{aligned} \log J_u(\theta) &= \log \sum_{i=1}^n q(i) \frac{\pi_{\theta}(\mathbf{a}_i | \mathbf{s}_i) u_{\tau}(r_i)}{q(i)} \quad (9) \\ &\geq \sum_{i=1}^n q(i) \log \frac{\pi_{\theta}(\mathbf{a}_i | \mathbf{s}_i) u_{\tau}(r_i)}{q(i)} \\ &= \sum_{i=1}^n q(i) [\log \pi_{\theta}(\mathbf{a}_i | \mathbf{s}_i) + \log u_{\tau}(r_i) - \log q(i)] \\ &= \mathcal{F}(q, \theta, \tau), \quad (10) \end{aligned}$$

due to Jensen's inequality. The re-weighting distribution $q(i)$ obeys the constraint $\sum_{i=1}^n q(i) - 1 = 0$. The ensuing EM algorithm is:

Algorithm 1 An EM algorithm for optimizing both the expected reward as well as the reward-transformation is given by

E-Step:

$$q_{k+1}(j) = \frac{\pi_{\theta_k}(\mathbf{a}_j | \mathbf{s}_j) u_{\tau_k}(r_j)}{\sum_{i=1}^n \pi_{\theta}(\mathbf{a}_i | \mathbf{s}_i) u_{\tau_k}(r_i)}, \quad (11)$$

M-Step Policy:

$$\theta_{k+1} = \arg \max_{\theta} \sum_{i=1}^n q_{k+1}(i) \log \pi_{\theta}(\mathbf{a}_i | \mathbf{s}_i), \quad (12)$$

M-Step Reward Transformation:

$$\tau_{k+1} = \arg \max_{\tau} \sum_{i=1}^n q_{k+1}(i) \log u_{\tau}(r_i). \quad (13)$$

Derivation. The E-Step is given by

$$q = \arg \max_q \mathcal{F}(q, \theta, \tau)$$

while fulfilling the constraint $0 = \sum_{i=1}^n q(i) - 1$. Thus, we obtain a constraint optimization problem with La-

grange multiplier λ :

$$L(\lambda, q) = \sum_{i=1}^n q(i) [\log \pi_\theta(\mathbf{a}_i | \mathbf{s}_i) + \log u_\tau(r_i) - \log q(i) + \lambda] \quad (14)$$

Optimizing $L(\lambda, q)$ with respect to q and λ results in Eq. (11). Optimizing $\mathcal{F}(q_{k+1}, \theta, \tau)$ with respect to θ and τ yields Eqs.(12, 13). ■

3.3. Reinforcement Learning by Reward-Weighted Regression

Let us assume the specific class of normally distributed policies

$$\pi_\theta(\mathbf{a} | \mathbf{s}) = \mathcal{N}(\mathbf{a} | \mu_\theta(\mathbf{s}), \sigma^2 \mathbf{I}) \quad (15)$$

with $\mu_\theta(\mathbf{s}) = \phi(\mathbf{s})^T \theta$ where $\phi(\mathbf{s})$ denotes some fixed preprocessing of the state by basis functions and $\sigma^2 \mathbf{I}$ determines the exploration¹. Furthermore, we choose the reward transformation $u_\tau(r) = \tau \exp(-\tau r)$ which fulfills all our requirements on a reward transformation (cf. Sec.3.1) for $r > 0$. Algorithm 1 thus becomes:

Algorithm 2 *The update equations for the policy $\pi_\theta(\mathbf{a} | \mathbf{s}) = \mathcal{N}(\mathbf{a} | \mu_\theta(\mathbf{s}), \sigma^2 \mathbf{I})$ are:*

$$\theta_{k+1} = (\Phi^T \mathbf{W} \Phi)^{-1} \Phi^T \mathbf{W} \mathbf{Y}, \quad (16)$$

$$\sigma_{k+1}^2 = \|\mathbf{Y} - \theta_{k+1}^T \Phi\|_{\mathbf{W}}^2, \quad (17)$$

where

$$\mathbf{W} = \text{diag}(u_\tau(r_1), u_\tau(r_2), \dots, u_\tau(r_n)) / \sum_{i=1}^n u_\tau(r_i)$$

denotes a diagonal matrix with transformed rewards, and

$$\Phi = [\phi(\mathbf{s}_1), \phi(\mathbf{s}_2), \dots, \phi(\mathbf{s}_n)]^T, \\ \mathbf{Y} = [\mathbf{a}_1, \mathbf{a}_2, \dots, \mathbf{a}_n]^T$$

denote basis functions and actions, respectively. The update of the reward transformation $u_\tau(r) = \tau \exp(-\tau r)$ is

$$\tau_{k+1} = \frac{\sum_{i=1}^n u_\tau(r_i)}{\sum_{i=1}^n u_\tau(r_i) r_i}. \quad (18)$$

Derivation. When computing $q_{k+1}(j)$ from samples in Eq. (11), we have $q_{k+1}(j) = u_\tau(r_j) / \sum_{i=1}^n u_\tau(r_i)$ as the probabilities are replaced by relative frequencies. We insert the policy

$$\pi_\theta(\mathbf{a} | \mathbf{s}) = (2\pi\sigma^2)^{-d/2} \exp\left(-\frac{1}{2\sigma^2} \|\mathbf{a} - \phi(\mathbf{s})^T \theta\|^2\right)$$

¹Note that $\sigma^2 \mathbf{I}$ could be replaced by a full variance matrix with little changes in the algorithm. However, this would result in a quadratic growth of parameters with the dimensionality of the state and is therefore less desirable.

into Eq. (12). By differentiating with respect to θ and equating the result to zero, we obtain

$$\theta = \left(\sum_{i=1}^n q_{k+1}(i) \phi_i \phi_i^T \right)^{-1} \left(\sum_{i=1}^n q_{k+1}(i) \phi_i \mathbf{a}_i \right)$$

with $\phi_i = \phi(\mathbf{s}_i)$. In matrix vector form, this corresponds to Eq. (16). Analogously, the reward transformation is obtained from differentiation w.r.t. τ as

$$\sum_{i=1}^n q_{k+1}(i) \frac{\partial}{\partial \tau} \log u_\tau(r_i) \\ = \sum_{i=1}^n q_{k+1}(i) (\tau^{-1} - r_i) = \mathbf{0}, \quad (19)$$

which results in Eq. (18). ■

The presented algorithm is focussed on real-time improvement of the policy and, thus, highly adapted to this setting. However, for offline learning of an initial policy, an improved algorithm can be used. In this case, we make use of Bayesian importance sampling (Andrieu et al., 2003) in the expectation step in order to improve efficiency and thus obtain a faster, computationally more efficient convergence. This algorithm offers an improvement in off-policy setting as at each step we are now allocating samples based on how well the maximizations step is doing. However, as a re-iteration of the data is not possible on-policy in real-time settings, we can only use it for pre-training the initial policy.

4. Application to Learning Operational Space Control

We are now in the position to return to our original objective of learning operational space control. As stated in Eq. (5), learning operational space control is equivalent to obtaining a mapping $\mathbf{s} = (\mathbf{q}, \dot{\mathbf{q}}, \ddot{\mathbf{x}}_{\text{ref}}, \mathbf{a}_0) \rightarrow \mathbf{a}$ from sampled data using a function approximator. As the dimensionality of the task-space reference behavior $\ddot{\mathbf{x}}_{\text{ref}}$ is lower than that of the motor commands \mathbf{a} , infinitely many solutions \mathbf{a} exist for a given \mathbf{s} . The solution space of motor commands \mathbf{a} achieving the same reference acceleration $\ddot{\mathbf{x}}_{\text{ref}}$ does usually not form a convex set in most robots, a problem first described in the context of learning inverse kinematics (Jordan & Rumelhart, 1992). Thus, when learning $\mathbf{s} \rightarrow \mathbf{u}$ as a function approximation problem, the learning algorithm can create physically invalid solutions.

Nevertheless, the non-convexity issue can be resolved by employing a spatially localized supervised learning system, an approach that was first introduced in the context of inverse kinematics learning (Bullock et al., 1993). The feasibility of this idea can be demonstrated by re-writing Eq. (2) in its proper functional form, i.e., not as an inverse function:

$$\ddot{\mathbf{x}} = \mathbf{J}(\mathbf{q})\mathbf{M}^{-1}(\mathbf{q}) (\mathbf{a} + \mathbf{F}(\mathbf{q}, \dot{\mathbf{q}}) + \dot{\mathbf{J}}(\mathbf{q}, \dot{\mathbf{q}})\dot{\mathbf{q}})$$

If we partition the state space of the robot, spanned by $\mathbf{q}, \dot{\mathbf{q}}$, into regions where $\mathbf{q}, \dot{\mathbf{q}}$ are approximately constant, the average over all solutions resulting in a particular $\ddot{\mathbf{x}}_{\text{ref}}$ can be written as:

$$\begin{aligned} \ddot{\mathbf{x}}_{\text{ref}} &= \langle \ddot{\mathbf{x}}_{\text{ref}} \rangle = \langle \mathbf{J}\mathbf{M}^{-1}(\mathbf{a} + \mathbf{F}) + \dot{\mathbf{J}}\dot{\mathbf{q}} \rangle \\ &= \mathbf{J}\mathbf{M}^{-1}(\langle \mathbf{a} \rangle + \mathbf{F}) + \dot{\mathbf{J}}\dot{\mathbf{q}}. \end{aligned} \quad (20)$$

Thus, in the vicinity of same $\mathbf{q}, \dot{\mathbf{q}}$ all possible \mathbf{a} that achieve the same $\ddot{\mathbf{x}}_{\text{ref}}$ form a convex solution set, since any average over different solutions $\mathbf{a}_1, \dots, \mathbf{a}_n$ will be guaranteed to still achieve the given $\ddot{\mathbf{x}}_{\text{ref}}$ ². Consequently, our approach to learning operational space control will partition the control law in form of locally linear controllers

$$\mathbf{a}^i = [\ddot{\mathbf{x}}_{\text{ref}}^T, \dot{\mathbf{q}}^T, 1]^T \beta^i, \quad (21)$$

which are active only in a region around a particular $\mathbf{q}^i, \dot{\mathbf{q}}^i$. Note that we added constant input in Eq. (21) to account for the intercept of a linear function. From a control engineering point of view, this argument corresponds to the insight that nonlinear control can often be accomplished through local linearizations at the point of interest, and that, in general, linear systems do not have the problem of non-convexity of the solution space when learning an inverse function.

Next we need to address how to find an appropriate piecewise linearization for the locally linear controllers. For this purpose, we learn a locally linear forward or predictor model

$$\ddot{\mathbf{x}}^i = [\dot{\mathbf{q}}^T, \mathbf{a}^T, 1]^T \hat{\beta}^i. \quad (22)$$

Learning this forward model is a standard supervised learning problem, as the mapping is guaranteed to be a proper function. A method of learning such a forward model is Locally Weighted Projection Regression (LWPR) (Schaal et al., 2002). This fast online learning method scales into high-dimensions, has been used for inverse dynamics control of humanoid robots, and can automatically determine the number of local models that are needed to represent the function. The membership to a local model is computed using a weight generated from a Gaussian kernel:

$$w^i(\mathbf{q}, \dot{\mathbf{q}}) = \exp\left(-\frac{1}{2}\left(\begin{bmatrix} \mathbf{q} \\ \dot{\mathbf{q}} \end{bmatrix} - \mathbf{c}^i\right)^T \mathbf{D}^i \left(\begin{bmatrix} \mathbf{q} \\ \dot{\mathbf{q}} \end{bmatrix} - \mathbf{c}^i\right)\right)$$

centered at a fixed \mathbf{c}_i in $(\mathbf{q}, \dot{\mathbf{q}})$ -space, and shaped by a diagonal distance metric \mathbf{D}_i . For a closer description

²Note that the localization in velocity $\dot{\mathbf{q}}$ can be dropped for a pure rigid body formulation as it is linear in the $\dot{q}_i \dot{q}_j$ for all degrees of freedom i, j ; this, however, is not necessarily desirable as it will add new inputs to the local regression problem whose number grows quadratically with the number of degrees of freedom.

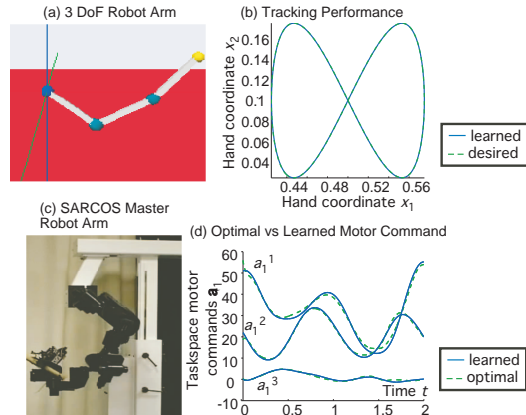


Figure 2. Systems and results of evaluations: (a) screen shot of the 3 DoF arm simulator, (c) Sarcos robot arm, used as simulated system and for actual robot evaluations in progress. (b) Tracking performance for a planar figure-8 pattern for the 3 DoF arm, and (d) comparison between the analytically obtained optimal control commands in comparison to the learned ones for one figure-8 cycle of the 3DoF arm.

of this statistical learning algorithm see (Schaal et al., 2002).

As learning the forward model provides a suitable partitioning of the robot’s state space into regions of local linearity, we re-use this partitioning to learn a local controller in each partition, an approach that resembles several previous forward-inverse model learning approaches (Haruno et al., 2001). The local controllers are learned using the reward-weighted regression approach described in Section 3, with metric $\mathbf{N} = \mathbf{I}$ in Eq. (3) – the reinforcement learning ensured that all local controllers learned a globally consistent solution to the operational space control task. Each local controller maintained its own adaptive reward transformation and associated parameter τ_i . The nominal control law \mathbf{a}_0 is learned separately; it consists of a gravity compensation component which is obtained directly with LWPR by supervised learning and a joint-space force pushing the robot towards a rest posture as explained in the introduction.

We evaluated our approach on two different simulated, physically realistic robots: (i) a three degree-of-freedom (DoF) planar robot arm and (ii) a seven DoF simulated SARCOS master robot arm (Figure 2 (a)). An implementation on the real, physical SARCOS master robot arm (Figure 2(c)) is currently in progress. Experiments were conducted as follows: first, learning the forward models and an initial control policy in each local model was obtained from random point-to-point movements in joint space using a simple PD control

law. This “motor babbling” exploration was necessary in order bootstrap learning with some initial data to prevent the rather slow learning as observed in similar direct-inverse learning approaches (Jordan & Rumelhart, 1992). The measured end-effector accelerations served as desired acceleration in Eq. (21), and all other variables for learning the local controllers were measurable as well. Subsequently, the learning controller was used on-policy with the normally distributed actuator noise serving as exploration. Both robots learned to track desired trajectories with high accuracy, as shown in Figure 2 (b). For the three DoF arm, we verified the quality of the learned control commands in comparison to the analytical solution, given in Eq. (4): Figure 2(d) demonstrates that the motor commands of the learned and analytically optimal case are almost identical. Learning results of the simulated seven DoF SARCOS robot achieved almost the same endeffector tracking quality as shown in Figure 2(b) (and is therefore not shown separately here). However, the joint commands were not quite as close to the optimal ones as for the 3 DoF arm - the rather high dimensional learning space of the 7 DoF arm most likely requires more extensive training and more careful tuning of the LWPR learning algorithm to achieve local linearizations with very high accuracy and with enough data to find the optimal solution. The 3 DoF required about 2 hours of real-time training, while the 7 DoF arm was run for about 6 hours of real-time training.

5. Conclusion

This paper contributes in two different ways to advancing the state-of-the-art of learning control. First, we introduced a framework for learning operational space control, a type of controller that has found little practical realizations due to problems with system identification in actual complex robots. Our learning methods avoid system identification entirely. Second, we introduced the idea of reinforcement learning by reward-weighted regression. While we realized this method here for immediate reward problems, i.e., finding optimal solutions in resolving redundancy in operational space control, we believe that there are much broader applications also in the realm of temporally delayed rewards, in particular for learning from trajectories or roll-outs. Reinforcement learning by reward-weighted regression has some of the flavor that was envisioned for modern approaches to reinforcement learning, i.e., the transformation of the reinforcement learning problem into an efficient supervised learning problem. We demonstrated the success of our approach on implementations on complex robot simulations, which will be followed by actual robot implementations in the

near future.

Acknowledgements

We thank the anonymous reviewers for their insightful comments and suggesting the improvements for our initial policy training.

References

- Andrieu, C., de Freitas, N., Doucet, A., & Jordan, M. (2003). An introduction to mcmc for machine learning. *Machine Learning*, 50, 5–43.
- Bullock, D., Grossberg, S., & Guenther, F. H. (1993). A self-organizing neural model of motor equivalent reaching and tool use by a multijoint arm. *Journal of Cognitive Neuroscience*, 5, 408–435.
- Dayan, P., & Hinton, G. E. (1997). Using expectation-maximization for reinforcement learning. *Neural Computation*, 9, 271–278.
- Featherstone, R. (1987). *Robot dynamics algorithms*. Kluwer Academic Publishers.
- Haruno, M., Wolpert, D. M., & Kawato, M. (2001). Mosaic model for sensorimotor learning and control. *Neural Comput*, 13, 2201–20.
- Jordan, I. M., & Rumelhart (1992). Supervised learning with a distal teacher. *Cognitive Science* (pp. 307–354).
- Kaelbling, L. P., Littman, M. L., & Moore, A. W. (1996). Reinforcement learning: A survey. *Journal of Artificial Intelligence Research*, 4, 237–285.
- Khatib, O. (1987). A unified approach for motion and force control of robot manipulators: The operational space formulation. *IEEE Journal of Robotics and Automation*, 3, 43–53.
- Nakanishi, J., Cory, R., Mistry, M., Peters, J., & Schaal, S. (2005). Comparative experiments on task space control with redundancy resolution. *IROS 2005*.
- Peters, J., Mistry, M., Udwadia, F., & Schaal, S. (2005). A unifying methodology for the control of robotic systems. *IROS 2005*.
- Schaal, S., Atkeson, C. G., & Vijayakumar, S. (2002). Scalable techniques from nonparametric statistics for real-time robot learning. *Applied Intelligence*, 17, 49–60.
- Spall, J. C. (2003). *Introduction to stochastic search and optimization: Estimation, simulation, and control*. Hoboken, NJ: Wiley.