

Stochastic Optimal Control Methods for Investigating the Power of Morphological Computation

Elmar A. Ruckert^{*,**}
Graz University of Technology

Gerhard Neumann[†]
Technische Universität Darmstadt

Abstract One key idea behind morphological computation is that many difficulties of a control problem can be absorbed by the morphology of a robot. The performance of the controlled system naturally depends on the control architecture and on the morphology of the robot. Because of this strong coupling, most of the impressive applications in morphological computation typically apply minimalistic control architectures. Ideally, adapting the morphology of the plant and optimizing the control law interact so that finally, optimal physical properties of the system and optimal control laws emerge. As a first step toward this vision, we apply optimal control methods for investigating the power of morphological computation. We use a probabilistic optimal control method to acquire control laws, given the current morphology. We show that by changing the morphology of our robot, control problems can be simplified, resulting in optimal controllers with reduced complexity and higher performance. This concept is evaluated on a compliant four-link model of a humanoid robot, which has to keep balance in the presence of external pushes.

Keywords

Probabilistic inference, motor planning, stochastic optimal control, morphological computation, compliant robots

A version of this paper with color figures is available online at http://dx.doi.org/10.1162/artl_a_00085. Subscription required.

1 Introduction

The control of compliant robots is inherently difficult, due to their nonlinear stochastic dynamics. Morphological computation poses the vision that parts of the complexity of a control or learning task can be outsourced to the morphology of the robot [11, 12]. This has been illustrated in many impressive biologically inspired robotic applications. For example, Tedrake et al. [14] showed that the task complexity of learning to walk is drastically reduced by exploiting the dynamics of a passive biped walker. Iida and Pfeifer [8] demonstrated that the stability of a simple four-legged robot exhibits surprisingly robust behavior using a mixture of active and passive joints. For swimming [20] or flying [19] robots, investigations of the morphologies yield a rich behavioral diversity using only a single actuator exploiting the dynamics of the environment.

* Contact author.

** Institute for Theoretical Computer Science, Graz University of Technology, Inffeldgasse 16b/1, Graz, 8010, Austria.
E-mail: ruckert@igi.tugraz.at

† Intelligent Autonomous Systems, Technische Universität Darmstadt, Hochschulstraße 10, 64289 Darmstadt, Germany.
E-mail: neumann@ias.tu-darmstadt.de

These approaches typically used minimalistic control architectures, like open loop actuation with sinusoidal drives [8, 20, 19] or a simple linear combination of state-dependent features [14]. The morphology is assumed to absorb much of the computation needed for controlling the robot, and therefore less emphasis is placed on the controller. More sophisticated control architectures were omitted, as these are typically more difficult to fine-tune for a new morphology. Thus, finding a good control law constitutes an individual optimization problem for each morphology, which renders the joint optimization of the morphology and the controller challenging. However, an autonomous system has to encompass both aspects—the morphology and the control architecture—for efficient learning of complex motor skills.

One remarkable exception to the typically used minimalistic control architectures is the theoretical study of the morphological contribution for compliant bodies in Hauser et al. [5]. Those authors have demonstrated in simulations for simple but nonlinear generic models of physical bodies, based on mass-spring systems, that by outsourcing the computation to the physical body, the difficult problem of learning to control a complex body could be reduced to a simple and perspicuous linear learning task. Thus, it cannot get stuck in local minima of an error function. However, learning the optimal morphology and the optimal control law remains an open problem for many motor control tasks using real robots.

As a first step toward this vision, we propose to directly use optimal control methods for eliminating the dependence between the morphology of the robot and the control architecture. Thus, for evaluating the computational power of the morphology, we will always use the optimal control law connected to this morphology. This allows for a fair comparison of different morphologies and simplifies the search for an optimal morphology.

For determining our optimal control law for a given morphology, we can use one of the many tools provided by stochastic optimal control (SOC) methods, such as those of Todorov and Li [15], Kappen, [6], and Toussaint [16]. These methods have been shown to be powerful approaches for movement planning in high-dimensional robotic systems. We will use approximate inference control (AICO) [16], as it is a state-of-the-art planning method for stochastic optimal control tasks. AICO is based on probabilistic inference for motor control. The beauty of this approach is that there is no distinction between sensor and motor, perception and action. We can include a multitude of variables, some of which might represent some features of the state, some of which might represent goals, constraints, or motivations in the future, and some of which might represent future actions or motor signals.

Like all other stochastic optimal control methods, AICO minimizes a cost function, which is in our case given by the quadratic distance to a target state and the energy used in the movement. In order to apply the AICO method to torque-constraint dynamical models, we will briefly explain in Section 2 how to extend the algorithm to systems with control limits.

Quantifying how much computation is done by the morphology of the robot and how much is done by the controller is often a difficult problem. Despite the advances in theoretical models of morphological computation [5] where the computational power of the plant can be assessed in a principled manner, this remains an open problem for many motor control tasks using real robots. Yet, by the use of optimal control laws, the computation done by the morphology can be quantified by investigating the complexity of the controller. Thus, if the controller needs to do a lot of computation, less computation is provided by the morphology in order to fulfill a task.

AICO also provides us with a time-varying linear feedback controller as a policy to generate movement trajectories. We will use the variance of the control gains as a complexity measure of the controller. If the control gains are almost constant in time, the control law is close to linear and does not perform much computation. However, if the control gains are highly varying, the controller needs to do a lot of computation, and therefore, less computation is provided by the morphology.

As different complexity measures are possible, we additionally use the final costs and the total jerk of a movement trajectory for comparison. We illustrate the power of morphological computation combined with optimal control on a dynamical model of a humanoid robot (70 kg, 2 m). The robot is modeled by a four-link pendulum, which has to keep balance in the presence of external pushes. We will show in Section 3 that on changing the morphology of a robot (e.g., the joint friction, the link

lengths, or the spring constants), the resulting optimal controllers have reduced complexity. As we will demonstrate, there are optimal values for the physical properties for a given control task, which can only be found by the use of optimal control laws. With naive control architectures, different, suboptimal morphologies would be chosen.

1.1 Related Work

We propose to use optimal control methods to eliminate the dependence between the control architecture and the morphology of the robot. These SOC methods, such as those of Todorov and Li [15], Kappen [6], and Toussaint [16], have been shown to be powerful methods for controlling high-dimensional robotic systems. For example, the incremental linear quadratic Gaussian (iLQG) [15] algorithm is one of the most commonly used SOC methods. It uses Taylor expansions of the system dynamics and cost function to convert the nonlinear control problem in a linear dynamics, quadratic costs, and Gaussian noise (LQG) system. The algorithm is iterative—the Taylor expansions are recalculated at the newly estimated optimal trajectory for the LQG system.

An SOC problem can be reformulated as an inference problem in a graphical model [7, 16], which has the nice property that there is no distinction between sensor and motor, perception and action. The unknown quantities (i.e., the states and actions) can be efficiently inferred by using, for example, the AICO algorithm [16]. The graphical model is given by a simple dynamic Bayesian network with states \mathbf{x}_t , actions \mathbf{u}_t , and task variables \mathbf{z}_t (representing the costs) as nodes; see Figure 1. In this graphical model, observations are denoted by shaded nodes, in contrast to the unknown random variables, which are indicated by circles. If beliefs in the graphical model are approximated as Gaussian, the resulting algorithm is very similar to iLQG. Gaussian message passing iteratively reapproximates local costs and transitions as LQG around the current mode of the belief within a time slice. One contrast with iLQG is that AICO uses forward messages instead of a forward rollout to determine the point of a local LQG approximation. Thus, AICO allows us to iterate belief reapproximation within a time slice until convergence. This may lead to faster overall convergence [16].

The dynamic Bayesian network shown in Figure 1 is fully specified by the conditional distributions encoded by the cost function and by the state transition model. As with most SOC methods [6, 15, 16], we assume that the cost function and the state transition model are known. However, for model learning, many types of function approximators can be applied [9, 10, 18]. For learning the cost function, inverse reinforcement learning methods such as those of Abbeel and Ng [1] and Boularias et al. [3] could be used. Since we demonstrate in this article how morphological computation and optimal control can benefit from each other in general, an extension to the more complex learning problems remains for future research.

The original formulation of the AICO method [16] does not consider torque limits, which are important for many robotic experiments as well for the dynamic balancing experiments we consider in this article. Therefore we extended the algorithm, as is discussed in the next section.




Figure 1. The graphical model for probabilistic planning, where we consider finite-horizon tasks with T time steps. The state variable \mathbf{x}_t , for example, is composed of the joint angles and joint velocities of a robot. Controls are labeled by \mathbf{u}_t . The beauty of probabilistic inference for motor control is that we can include a multitude of variables, some of which might represent some features of the state, some of which might represent goals, constraints, or motivations in the future, and some of which might represent future actions or motor signals. These constraints are expressed in the model by the task variables \mathbf{z}_t .

2 Probabilistic Inference for Motor Planning

We use the probabilistic planning method AICO [16] as the optimal control method. Most applications of AICO are in the kinematic planning domain. Here, we want to apply AICO to fully dynamic, torque-controlled robot simulations. Therefore, we had to extend the AICO framework with control or torque limits, as is explained in the next sections.

2.1 Approximate Inference Control

We will briefly clarify the notation for our discussion. Let \mathbf{x}_t denote the state, and \mathbf{u}_t the control vector at time step t . A trajectory τ is defined as a sequence of state-control pairs, $\tau = \langle \mathbf{x}_{1:T}, \mathbf{u}_{1:T-1} \rangle$, where T is the length of the trajectory. Each trajectory has associated costs

$$L(\tau) := \sum_{t=1}^T c_t(\mathbf{x}_t, \mathbf{u}_t), \quad (1)$$

where $c_t(\mathbf{x}_t, \mathbf{u}_t)$ represents the cost function for a single time step, which is in our case given by the quadratic distance to a target state and the energy used for the movement. Only the cost function $c_{1:T}(\cdot)$ and the initial state \mathbf{x}_1 are known to the optimal control algorithm. The unknown trajectory τ is the result of the inference process.

AICO uses message passing in graphical models to infer the trajectory τ . In order to transform the minimization of $L(\tau)$ into an inference problem, for each time step, an individual binary random variable z_t is introduced. This random variable indicates a reward event. Its probability is given by

$$P(z_t = 1 | \mathbf{x}_t, \mathbf{u}_t, t) \propto \exp(-c_t(\mathbf{x}_t, \mathbf{u}_t)).$$

AICO now assumes that a reward event $z_t = 1$ is observed at every time step; see Figure 1. Given that evidence, AICO calculates the posterior distribution $P(\mathbf{x}_{1:T}, \mathbf{u}_{1:T} | z_{1:T} = 1)$ over trajectories.

We will use the simplest version of AICO, where an extended Kalman smoothing approach is used to estimate the posterior $P(\mathbf{x}_{1:T}, \mathbf{u}_{1:T} | z_{1:T} = 1)$. The extended Kalman smoothing approach uses Taylor expansions to linearize the system and subsequently uses Gaussian messages for belief propagation in a graphical model. Gaussian message passing iteratively reapproximates local costs and transitions as an LQG system around the current mode of the belief within a time slice. Concise derivations of the messages for AICO are given in Appendix 1, as they are used to further extend the algorithm to implement control constraints.

AICO provides us with a linear feedback controller for each time slice of the form

$$\mathbf{u}_t = \mathbf{O}_t \mathbf{x}_t + \mathbf{o}_t,$$

where \mathbf{O}_t is the inferred feedback control gain matrix and \mathbf{o}_t is the linear feedback controller term. For our evaluations we also use a complexity measure that is proportional to the variance of this feedback control law:

$$\text{Var}(\mathbf{O}_{1:T-1}, \mathbf{o}_{1:T-1}) = \sum_{j=1}^{D^x} \sum_{j=1}^{D^u} \text{var}(\mathbf{O}_{j,1:T-1}) + \sum_{j=1}^{D^u} \text{var}(\mathbf{o}_{j,1:T-1}), \quad (2)$$

where the dimension of the states is denoted by D^x and the dimension of the controls is denoted by D^u .

AICO is only a local optimization method, and we have to provide an initial solution that is used for the first linearization. We will simply use the first state as the initialization: $\mathbf{x}_{2:T} = \mathbf{x}_1$.

If we use AICO with a constant cost and dynamic model for each time step, the algorithm reduces to calculating a linear-quadratic regulator (LQR), which is often used in optimal control. An LQR is the optimal linear feedback controller for a given linear system. In contrast, AICO uses a time-varying linear feedback controller, which may be different for each time step. In our experiments we will compare the two approaches on a dynamic nonlinear balancing task. Thus we evaluate the benefit of using AICO (time-varying linear feedback control) and an LQR (constant linear feedback control).

2.2 Cost Function for Constrained Systems

In order to apply the AICO algorithm to torque-controlled robots, we have to extend the framework to incorporate control limits, as the available motor torque is typically limited. This is done by adding a control-dependent punishment term $c_t''(\mathbf{u}_t)$ to the cost function in Equation 1. Thus for a trajectory τ we specify the costs

$$L(\tau) := \sum_{t=1}^T c_t(\mathbf{x}_t, \mathbf{u}_t) + c_t''(\mathbf{u}_t).$$

We use this additional term to punish controls \mathbf{u}_t that exceed a given bound \mathbf{u}_B , at time t :

$$c_t''(\mathbf{u}_t) = (\mathbf{u}_t - \mathbf{u}_{B_t})^T \mathbf{H}_{B_t} (\mathbf{u}_t - \mathbf{u}_{B_t}).$$

As a consequence, the resulting Gaussian distributions that are used to represent the costs c_t change. These distributions typically have zero mean in the control space, due to the typically used quadratic control costs. In the case of control limits, the mean of the distribution is nonzero. Consequently, the message-passing update equations are also used for the AICO algorithm change. The exact message-passing equations of AICO with control limits are presented in the Appendix.

3 Experiments

We investigate morphological computation combined with optimal control on dynamic nonlinear balancing tasks [2], where a robot gets pushed with a specific force and has to move so that it maintains balance. The optimal strategy is a nonlinear control law that returns the robot to the upright position. For our evaluations, different initial postures and multiple forces are used, as sketched in Figure 2. We study the morphological computation by changing different physical properties of a four-link model of a humanoid robot. In particular, we investigate the computation done by different friction coefficients and link lengths. Inspired by the work of Hauser et al. [5], we also incorporate springs into our model for analyzing different spring constants.

3.1 Setting

We use a four-link robot as a simplistic model of a humanoid (70 kg, 2 m) [2]. The eight-dimensional state \mathbf{x}_t is composed of the ankle, the knee, the hip, and the arm positions and their velocities. Table 1 in




Figure 2. Different initial postures of a four-link pendulum modeling a humanoid robot (70 kg, 2 m). The robot gets pushed with specific forces, F_1 and F_2 , and has to move so that it maintains balance. At the end of the movement the robot should stabilize at the upright position.

Appendix 2 shows the initial velocities (resulting from the force F , which always acts at the shoulder of the robot) and the valid joint angle range for the task. In addition to the joint limits, the controls are limited to the intervals $[\pm 70, \pm 500, \pm 500, \pm 250]$ Nm (ankle, knee, hip, and arm). For more detail we refer to Atkeson and Stephens [2].

We use a quadratic cost function given by

$$c_t(\mathbf{x}_t, \mathbf{u}_t) = (\mathbf{x}_t - \mathbf{x}^*)^T \hat{\mathbf{R}}_t (\mathbf{x}_t - \mathbf{x}^*) + \mathbf{u}_t^T \hat{\mathbf{H}}_t \mathbf{u}_t,$$

where the final state is denoted by $\mathbf{x}^* = \mathbf{0}$. The precision matrix $\hat{\mathbf{R}}_t$ determines how costly it is not to reach \mathbf{x}^* . The diagonal elements of $\hat{\mathbf{R}}_{1:T}$ are set to 2×10^3 for joint angles and to 0.2 for joint velocities. The controls are punished by $\hat{\mathbf{H}}_{1:T-1} = 0.05\mathbf{I}$.

The movement trajectories are simulated with a time step of 0.5 ms. For the AICO algorithm, we use a planning time step of $\Delta t = 5$ ms and a horizon of $T = 500$, which results in a movement duration of 2.5 s. We use a torque-dependent noise model—the controls are multiplied by Gaussian noise ϵ with mean 1 and standard deviation 0.25, that is, we use 25% torque-dependent noise. The noise affects the system dynamics $\dot{\mathbf{x}} = f(\mathbf{x}_t, \mathbf{u}_t + \epsilon)$ while simulating a trajectory, where $\epsilon \sim \mathcal{N}(\epsilon | 0, 0.25^2 \cdot \text{abs}(\mathbf{u}_t))$. The experiments are performed for multiple forces sampled from $F \sim \mathcal{U}[-15, +15]$ Ns, as shown in Figure 2.

3.2 Complexity Measures

We evaluate the average values of the final costs $L(\tau)$ in Equation 1, the variance of the time-varying controller gains returned by the AICO approach in Equation 2, and the total jerk $J(\tau) = \Delta t \sum_{t=0}^T \dot{\mathbf{u}}_t^T \dot{\mathbf{u}}_t$ of the trajectory (proportional to the squared derivative of the torque). Different complexity measures would be possible. However, the chosen ones are plausible, since the complexity of controlling the robot is reflected by how well the costs $L(\tau)$ are optimized. As the jerk of a movement tracks the derivatives of the torques, this also seems to be a reliable complexity measure. Finally, the variance of the time-varying controller gains quantifies the complexity of the control law in comparison with linear controllers (no variance).

3.3 Friction Induces Morphological Computation

In this experiment, we evaluate the influence of the friction coefficient γ on the quality of the optimal controller for the four-link model. The friction coefficient directly modulates the acceleration of the joints, that is, $\ddot{\phi}_\gamma = -\gamma \dot{\phi}$.

Figure 3a shows the resulting final costs for different friction coefficients. In order to illustrate the need for sophisticated control architectures, we compare the optimal control method AICO with a LQR controller and a simple constant linear feedback controller. While AICO calculates the optimal




Figure 3. The influence of the friction coefficient on controlling the four-link model using the optimal control methods AICO and LQR. Illustrated are the mean and the standard error over 50 trajectories. As complexity measures of the morphology we used (a) the final costs, (b) the total jerk, and (c) the controller variance. For the final costs we additionally compare with a simple linear feedback controller, denoted by LC. For (c) we compare with an LQR controller with time-varying costs, in contrast to the standard LQR method, which uses constant costs for all time steps.

control law for the nonlinear system, the LQR controller linearizes the system at the upright position and subsequently calculates the closed form solution of the optimal controller for the resulting LQG system. The constant linear feedback controller does not adapt to the morphology of the robot. As we can see from Figure 3a, we cannot determine the optimal morphology with the simple linear controller, while we can recognize a clear minimum with the AICO and the LQR controller.

AICO could find the simplest control law according to the jerk criterion in Figure 3b and was also able to find control laws with considerably decreased costs. In Figure 3c, we illustrate the variance of the time-varying controller gains. The complexity of the resulting controllers is reduced by changing the friction coefficient. The minimum lies in the same range as the minimum of the cost function. Thus, in this case, a simpler controller resulted in better performance because the morphology (the friction) of the robot has simplified the control task. Still, we needed a more complex optimal control algorithm to discover this simpler control law than that given by the LQR.

In Figure 4 we depict the joint and torque trajectories for the hip and the arm joint, where we applied the friction coefficients $\gamma = 0$, $\gamma = 12$, and $\gamma = 30$. With $\gamma = 0$, the trajectories overshoot as shown in Figure 4a, whereas in Figure 4c more energy is needed for stabilizing the robot with $\gamma = 30$. The trajectories for the optimal friction coefficient $\gamma = 12$ (according to the jerk complexity measure) are shown in Figure 4b.

3.4 The Influence of the Link Lengths

To show the effects of morphological computation on a more complex variation of morphologies, we consider in this experiment the variation of the shank (i.e., the calf), the thigh, the trunk, and the arm length, $\mathbf{l} = [l_1, l_2, l_3, l_4]$. Initially the link lengths are $\mathbf{l}_0 = [0.5, 0.5, 1, 1]$ m and the weights of these links are specified by $\mathbf{m}_0 = [17.5, 17.5, 27.5, 7.5]$ kg.

We consider variations of \mathbf{l} of up to 50% from their initial configuration. Different link lengths of the four-link robot result in different initial velocities [2], which influences the complexity of the motor control problem. For this reason, we consider constant inertias for a fair comparison. Thus, for each link length \mathbf{l} the masses of the links are given by $\mathbf{m} = \mathbf{m}_0 \mathbf{l}_0^2 / \mathbf{l}^2$. To match the physical properties of humans we assume equal lengths for the shank and the thigh link: $l_1 = l_2$.

In addition to the link lengths of the thigh, the trunk, and the arm, we still optimize the friction coefficient, resulting in four parameters (l_2, l_3, l_4 , and γ) of our morphology. A naive search for optimal parameters as in the previous experiment is intractable in the multidimensional case. Therefore, we apply the nonlinear stochastic optimizer of Hansen et al. [4] as a local search method.

The stochastic optimizer locally explores the parameter space (l_2, l_3 , and l_4) until convergence. For example, the resulting final cost values for different link lengths are illustrated in Figure 5 when using AICO as the optimal control law. For this experiment, we used a friction coefficient of $\gamma = 12$, and we averaged over multiple initial states and multiple forces. The optimizer converged to link lengths of $l_2 = 0.39$, $l_3 = 1.18$, and $l_4 = 0.5$ for the thigh, the torso, and the arm.

For multiple friction coefficients, an evaluation of the presented complexity measures is shown in Figure 6. According to the final costs in Figure 6a, the best controller was found with a friction




Figure 4. The mean trajectories of the hip (ϕ_3) and the arm (ϕ_4) joints over 100 runs for the first 0.5 s, where we used the AICO approach for the friction coefficients $\gamma = 0$, $\gamma = 12$, and $\gamma = 30$.




Figure 5. Explored link lengths for the thigh, the torso, and the arm (l_2 , l_3 , and l_4) using AICO with the friction coefficient $\gamma = 12$. For this illustration, we discretized the parameter space to visualize the final cost values, which are denoted by the color of the dots. In the electronic version, darker dots correspond to lower cost values as specified by the color bar on the right. The optimal values of the link lengths are denoted by the large crosses ($l_2 = 0.39$, $l_3 = 1.18$, and $l_4 = 0.5$).

coefficient of $\gamma = 9$. For the jerk criteria in Figure 6b and the variance as complexity measure shown in Figure 6c, the minimum of the complexity measures lies in the same range.

The morphologies found by the optimizer for the AICO and the LQR controller are sketched in Figure 7 for the friction coefficients $\gamma = 0$, $\gamma = 9$, and $\gamma = 30$. The initial configuration of the robot is shown in Figure 7a and in Figure 7e. Interestingly, with increasing values of the friction coefficients, the link lengths change to compensate for the larger motor controls. This is shown in Figure 7b–d for the AICO controller and in Figure 7f–h for the LQR controller. AICO could again find significantly different morphologies from the LQR controller. As is illustrated in the labels of Figure 7, the morphologies found by AICO could produce lower costs than the morphologies found by LQR.

3.5 Robot Model with Linear Springs

Inspired by the theoretical analysis of [5], we also evaluate the computational power of linear springs in our model. The springs are mounted on the four-link model as illustrated in Figure 8. These springs support the control of the robot as they push it back into its upright position. Thus, the torques applied to the joints are given by the sum of the controls provided by the feedback controller and the spring forces $\mathbf{u}^* = \mathbf{u}_f + \text{diag}(\mathbf{k})\Delta\mathbf{l}_s$. The vector $\Delta\mathbf{l}_s$ contains the displacements of the four springs, and \mathbf{k} is a vector containing the four spring constants. Note that due to the calculation of the displacements, the springs act nonlinearly on the joint torques.

The adjustable parameters of the morphology include the friction coefficient and the four spring constants $\mathbf{k} = [k_1, k_2, k_3, k_4]$ for the ankle, the knee, the hip, and the arm joint. Initially the spring constants are set to $\mathbf{k}_0 = [1, 10^3, 10^3, 1]$. For the optimization, a lower bound of $0.1\mathbf{k}_0$ and an upper bound of $100\mathbf{k}_0$ are used.

As in the previous experiment, for different friction coefficients γ , we optimize the spring constants using the stochastic search method of Hansen et al. [4]. The results for the final costs, the total jerk, and the variance of the feedback controller as complexity measures are shown in Figure 9. The benefit of the supporting springs is best illustrated by the final costs (AICO: 109.5; LQR: 122.2) in Figure 9a,




Figure 6. The results for optimizing the link lengths and the friction coefficient of the four-link model. As before, we used (a) the final costs, (b) the total jerk, and (c) the variance of the time-varying feedback controller as complexity measures.




Figure 7. The learned link lengths for multiple friction coefficients $\gamma = 0, \gamma = 9$, and $\gamma = 30$ for the AICO (b–d) and the LQR (f–h) controller. The link lengths are optimized with a stochastic optimizer and could vary up to 50% from the initial configuration shown in (a) and (e). With increasing friction coefficients, the link lengths change to compensate for the larger control costs. The final costs $L(\tau)$ for these morphologies are given in the labels.

which are drastically reduced compared to the final costs using the four-link model without springs (AICO: 146.9; LQR: 246.9) shown in Figure 3a.

The optimal spring constants strongly depend on the friction coefficient used. This is illustrated in Figure 10, where we compare the learned spring constants using AICO and LQR as control methods. With increasing friction, the spring constants increase to compensate for larger control costs.

3.6 Concluding Summary

We have shown that by using optimal control methods an optimal morphology could be found for the four-link robot model. This was evaluated on three tasks with increasing complexity: optimizing only the friction coefficient, optimizing the friction coefficient and the link lengths, and finally, optimizing the friction coefficient and the four spring constants. We can now also ask how beneficial it was to optimize




Figure 8. The four-link robot model endowed with linear springs, which are mounted at the midpoints of the joints. These springs support the control of the robot as they push it back into its upright position. Due to the calculation of the spring length displacements, the springs act nonlinearly on the joint torques.




Figure 9. The results for optimizing the spring constants and the friction coefficient of the four-link model. We used (a) the final costs, (b) the total jerk, and (c) the variance of the time-varying feedback controller as complexity measures.

the link lengths compared to optimizing the spring constants. This is best reflected by the achieved final costs in Figure 3a, Figure 6a, and Figure 9a, where the most complex task also performs best.

4 Conclusion

In this article, we have shown that optimal control and morphological computation are two complementary approaches that can benefit from each other. The search for an optimal morphology is simplified if we can calculate an optimal controller for a given morphology. This calculation can be done by new approaches from probabilistic inference for motor control, such as the approximate inference control algorithm [16]. By the use of optimal control methods, we have shown that for dynamic nonlinear balancing tasks, an appropriate setting of the friction, the link lengths, or the spring constants of the incorporated springs of the compliant robot can simplify the control problem.

We have demonstrated that there are optimal values for the physical properties for a given control task, which can only be found by the use of optimal control laws. With naive control architectures such as the evaluated constant linear feedback controller, which does not adapt to the morphology, different, suboptimal morphologies would be chosen.

In our studies, we assumed that we can always compute an optimal controller given the current morphology. If this is possible, then the morphological computation can be quantified using the proposed method. However, for computing optimal control laws (using, e.g., AICO), we need to




Figure 10. The optimal spring constants for different friction coefficients, using AICO and LQR.

know the costs at each time step, as well as the dynamics model. For more complex tasks, these costs and the model are usually unknown. However, as argued in the related work in Section 1.1, for model learning, many types of function approximators can be applied [9, 10, 18]. For learning the cost function, inverse reinforcement learning methods such as those of Abbeel and Ng [1], and Boularias et al. [3] could be used.

In the future, we plan to investigate more complex and more nonlinear tasks. In these cases the benefit of AICO in comparison with LQR controllers should be even more prominent. In the end we are planning to simultaneously evolve walking controllers based on AICO and the morphology of bipedal robots, such as a model of a planar walker.

Acknowledgments

This article was written with partial support by the European Union projects FP7-248311 (AMARSI) and FP7-506778 (PASCAL2). The funders had no role in the study design, data collection and analysis, decision to publish, or preparation of the manuscript.

References

1. Abbeel, P., & Ng, A. Y. (2004). Apprenticeship learning via inverse reinforcement learning. In *Proceedings of the Twenty-First International Conference on Machine Learning (ICML 2004)* (pp. 1–8).
2. Atkeson, C. G., & Stephens, B. (2007). Multiple balance strategies from one optimization criterion. In *Proceedings of the 7th International Conference on Humanoid Robots* (pp. 57–64).
3. Boularias, A., Kober, J., & Peters, J. (2011). Relative entropy inverse reinforcement learning. *Journal of Machine Learning Research—Proceedings Track*, 15, 182–189.
4. Hansen, N., Muller, S. D., & Koumoutsakos, P. (2003). Reducing the time complexity of the derandomized evolution strategy with covariance matrix adaptation (CMA-ES). *Evolutionary Computation*, 11(1), 1–18.
5. Hauser, H., Ijspeert, A. J., Fuchslin, R. M., Pfeifer, R., & Maass, W. (2011). Towards a theoretical foundation for morphological computation with compliant bodies. *Biological Cybernetics*, 105(5–6), 355–370.
6. Kappen, H. J. (2007). An introduction to stochastic control theory, path integrals and reinforcement learning. In *Proceedings of the 9th Granada Seminar on Computational Physics: Computational and Mathematical Modeling of Cooperative Behavior in Neural Systems* (pp. 149–181).
7. Kappen, H. J., Gómez, V., & Opper, M. (2012). Optimal control as a graphical model inference problem. *Machine Learning Journal*, 87(2), 159–182.
8. Iida, F., & Pfeifer, R. (2006). Sensing through body dynamics. *Robotics and Autonomous Systems*, 54(8), 631–640.
9. Nguyen-Tuong, D., Peters, J., Seeger, M., & Schölkopf, B. (2008). Learning inverse dynamics: A comparison. In *16th European Symposium on Artificial Neural Networks (ESANN 2008)* (pp. 13–18).
10. Nguyen-Tuong, D., Seeger, M., & Peters, J. (2008). Local Gaussian process regression for real time online model learning and control. In *Proceedings of the 22nd Annual Conference on Neural Information Processing Systems (NIPS 2008)* (pp. 1193–1200).
11. Pfeifer, R., & Bongard, J. (2006). *How the body shapes the way we think: A new view of intelligence*. Cambridge, MA: MIT Press.
12. Pfeifer, R., Lungarella, M., & Iida, F. (2007). Self-organization, embodiment, and biologically inspired robotics. *Science*, 318(5853), 1088–1093.
13. Rawlik, K., Toussaint, M., & Vijayakumar, S. (2010). An approximate inference approach to temporal optimization in optimal control. In *Proceedings of the 24th Annual Conference on Neural Information Processing Systems (NIPS 2010)* (pp. 2011–2019).
14. Tedrake, R., Zhang, T. W., & Seung, H. S. (2005). Learning to walk in 20 minutes. In *Proceedings of the Fourteenth Yale Workshop on Adaptive and Learning Systems*.
15. Todorov, E., & Li, W. (2005). A generalized iterative LQG method for locally-optimal feedback control of constrained nonlinear stochastic systems. In *Proceedings of the 24th American Control Conference (ACC 2005)*, Vol. 1 (pp. 300–306).

16. Toussaint, M. (2009). Robot trajectory optimization using approximate inference. In *Proceedings of the 26th International Conference on Machine Learning (ICML 2009)* (pp. 1049–1056).
17. Toussaint, M. (2011). *Lecture notes: Gaussian identities* (Technical report). Freie Universität Berlin. <http://userpage.fu-berlin.de/mtoussai/notes/gaussians.pdf>.
18. Vijayakumar, S., D’Souza, A., & Schaal, S. (2005). Incremental online learning in high dimensions. *Neural Computation*, 17(12), 2602–2634.
19. Wood, R. J. (2007). Design, fabrication, and analysis of a 3 DOF, 3 cm flapping-wing MAV. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2007)* (pp. 1576–1581).
20. Ziegler, M., Iida, F., & Pfeifer, R. (2006). “Cheap” underwater locomotion: Roles of morphological properties and behavioral diversity. In *9th International Conference on Climbing and Walking Robots (CLAWAR 2006)*.

Appendix I: Extension of the Approximate Inference Control Method

The original formulation of the AICO method [16] does not consider a linear term for the control costs. However, this is needed to encode torque limits, which are important for our dynamic balancing experiments, and hence we needed to extend the AICO algorithm.

The introduction of a linear term for the control costs yields not only a modified cost function, but also different update equations for the messages, and finally different equations for the optimal feedback controller. For completeness we will first recap the main steps in deriving the AICO method, and will then discuss the modifications to implement control constraints.

AI.1 Approximate Inference Control Without Torque Limits

For motor planning we consider the stochastic process

$$P(\mathbf{x}_{1:T}, \mathbf{u}_{1:T-1}, \mathbf{z}_{1:T}) = P(\mathbf{x}_1) \prod_{t=1}^{T-1} P(\mathbf{u}_t | \mathbf{x}_t) \prod_{t=1}^T P(\mathbf{x}_t | \mathbf{x}_{t-1}, \mathbf{u}_{t-1}) \prod_{t=1}^{T-1} P(z_t | \mathbf{x}_t, \mathbf{u}_t),$$

where $P(\mathbf{u}_t | \mathbf{x}_t)$ is the state-dependent prior for the controls, the distribution $P(\mathbf{x}_t | \mathbf{x}_{t-1}, \mathbf{u}_{t-1})$ is the state transition model, and $P(\mathbf{x}_1)$ is the initial state distribution. Here, we assume that the prior of the controls is independent of the states, and thus we will simply write $P(\mathbf{u}_t | \mathbf{x}_t) = P(\mathbf{u}_t)$ for the rest of this appendix. The time horizon is fixed at T time steps. The binary task variable z_t denotes a reward event; its probability is expressed by $P(z_t = 1 | \mathbf{x}_t, \mathbf{u}_t) \propto \exp(-c_t(\mathbf{x}_t, \mathbf{u}_t))$, where $c_t(\mathbf{x}_t, \mathbf{u}_t)$ is the immediate cost function¹ for time step t . It expresses a performance criterion (like avoiding a collision, or reaching a goal).

We want to compute the posterior $P(\mathbf{x}_{1:T}, \mathbf{u}_{1:T} | z_{1:T} = 1)$ over trajectories, conditioned on observing a reward ($z_t = 1$) at each time step t . This posterior can be computed by using message passing in the given graphical model of Figure 1. To simplify the computations, we integrate out the controls:

$$P(\mathbf{x}_{t+1} | \mathbf{x}_t) = \int_{\mathbf{u}_t} P(\mathbf{x}_{t+1} | \mathbf{x}_t, \mathbf{u}_t) P(\mathbf{u}_t | \mathbf{x}_t) d\mathbf{u}_t. \quad (3)$$

The marginal belief $b_t(\mathbf{x}_t)$ of a state at time t is given by

$$b_t(\mathbf{x}_t) = \alpha_t(\mathbf{x}_t) \beta_t(\mathbf{x}_t) \phi_t(\mathbf{x}_t), \quad (4)$$

¹ In this article the immediate cost function is composed of the intrinsic costs and the constraint costs: $c_t(\mathbf{x}_t, \mathbf{u}_t) + c_p(\mathbf{x}_t, \mathbf{u}_t)$.

where $\alpha_t(\mathbf{x}_t)$ is the forward message, $\beta_t(\mathbf{x}_t)$ is the backward message, and $\phi_t(\mathbf{x}_t)$ is the current task message. The messages are given by

$$\alpha_t(\mathbf{x}_t) = \int_{\mathbf{x}_{t-1}} P(\mathbf{x}_t|\mathbf{x}_{t-1})\alpha_{t-1}(\mathbf{x}_{t-1})\phi_{t-1}(\mathbf{x}_{t-1})d\mathbf{x}_{t-1}, \quad (5)$$

$$\beta_t(\mathbf{x}_t) = \int_{\mathbf{x}_{t+1}} P(\mathbf{x}_{t+1}|\mathbf{x}_t)\beta_{t+1}(\mathbf{x}_{t+1})\phi_{t+1}(\mathbf{x}_{t+1})d\mathbf{x}_{t+1}, \quad (6)$$

$$\phi_t(\mathbf{x}_t) = P(\zeta_t|\mathbf{x}_t). \quad (7)$$

We consider discrete-time, nonlinear stochastic systems with zero-mean Gaussian noise

$$P(\mathbf{x}_{t+1}|\mathbf{x}_t, \mathbf{u}_t) = \mathcal{N}(\mathbf{x}_{t+1}|f_{\text{Dyn}}(\mathbf{x}_t, \mathbf{u}_t), \mathbf{Q}_t).$$

The nonlinear stochastic system f_{Dyn} is approximated by a linear dynamics quadratic costs, and Gaussian noise (LQG) system, by Taylor expansion [15, 16]:

$$P(\mathbf{x}_{t+1}|\mathbf{x}_t, \mathbf{u}_t) = \mathcal{N}(\mathbf{x}_{t+1}|\mathbf{A}_t\mathbf{x}_t + \mathbf{a}_t + \mathbf{B}_t\mathbf{u}_t, \mathbf{Q}_t). \quad (8)$$

Thus, the system is linearized along a given trajectory $\langle \hat{\mathbf{x}}_{1:T}, \hat{\mathbf{u}}_{1:T-1} \rangle$ at every point in time. We will use f_t as shorthand for $f_{\text{Dyn}}(\mathbf{x}_t, \mathbf{u}_t)$. Then, the state transition matrix \mathbf{A}_t is given by $\mathbf{A}_t = \mathbf{I} + \frac{\delta f_t}{\delta \mathbf{x}_t} \Delta t$, the control matrix \mathbf{B}_t is $\mathbf{B}_t = \frac{\delta f_t}{\delta \mathbf{u}_t} \Delta t$, and the linear term reads

$$\mathbf{a}_t = \left(f_t - \frac{\delta f_t}{\delta \mathbf{x}_t} \mathbf{x}_t - \frac{\delta f_t}{\delta \mathbf{u}_t} \mathbf{u}_t \right) \Delta t.$$

In the original formulation of the AICO algorithm, the cost function c_t is approximated as

$$c_t(\mathbf{x}_t, \mathbf{u}_t) = \mathbf{x}_t^T \mathbf{R}_t \mathbf{x}_t - 2\mathbf{r}_t^T \mathbf{x}_t + \mathbf{u}_t^T \mathbf{H}_t \mathbf{u}_t.$$

Note that there is no linear term for the control costs, as we only punish quadratic controls. We can now write $P(\zeta_t = 1|\mathbf{x}_t, \mathbf{u}_t) = P(\zeta_t = 1|\mathbf{x}_t)P(\mathbf{u}_t)$ as

$$P(\zeta_t = 1|\mathbf{x}_t) \propto \mathcal{N}[\mathbf{x}_t|\mathbf{r}_t, \mathbf{R}_t], \quad (9)$$

$$P(\mathbf{u}_t) = \mathcal{N}[\mathbf{u}_t|\mathbf{0}, \mathbf{H}_t], \quad (10)$$

where the distributions are given in canonical form. The canonical form of a Gaussian is used because numerical operations such as products or integrals are easier to calculate in this notation. The canonical form is indicated by the square-bracket notation and given by

$$\mathcal{N}[\mathbf{x}|\mathbf{a}, \mathbf{A}] = \frac{\exp\left(-\frac{1}{2}\mathbf{a}^T \mathbf{A}^{-1} \mathbf{a}\right)}{|2\pi \mathbf{A}^{-1}|^{1/2}} \exp\left(-\frac{1}{2}\mathbf{x}^T \mathbf{A} \mathbf{x} + \mathbf{x}^T \mathbf{a}\right).$$

A Gaussian in normal form can always be transformed into the canonical form by $\mathcal{N}(\mathbf{x}|\mathbf{a}, \mathbf{A}) = \mathcal{N}[\mathbf{x}|\mathbf{A}^{-1}\mathbf{a}, \mathbf{A}^{-1}]$. For more details we refer to the Gaussian identities in Toussaint [17].

We can see in Equation 10 that our prior for applying the control \mathbf{u}_t is given by the control costs, that is, $\mathcal{N}[\mathbf{u}_t | \mathbf{0}, \mathbf{H}_t]$. By integrating out the controls from our system dynamics we get the following state transition probabilities:

$$P(\mathbf{x}_{t+1} | \mathbf{x}_t) = \int_{\mathbf{u}_t} \mathcal{N}(\mathbf{x}_{t+1} | \mathbf{A}_t \mathbf{x}_t + \mathbf{a}_t + \mathbf{B}_t \mathbf{u}_t, \mathbf{Q}_t) \mathcal{N}[\mathbf{u}_t | \mathbf{0}, \mathbf{H}_t] d\mathbf{u}_t \quad (11)$$

$$= \mathcal{N}(\mathbf{x}_{t+1} | \mathbf{A}_t \mathbf{x}_t + \mathbf{a}_t, \mathbf{Q}_t + \mathbf{B}_t \mathbf{H}_t^{-1} \mathbf{B}_t^T), \quad (12)$$

where the integral was evaluated using a reformulation of the propagation rule in Toussaint [17].

As we can see, all distributions in the approximated LQG system in Equation 12 are Gaussian, and thus also all messages are Gaussians and can be calculated analytically. The resulting messages are given in Toussaint [16].

AI.2 Approximate Inference Control with Torque Limits

In order to implement torque and joint limits, we introduce an additional cost function c_p , which punishes the violation of the given constraints. The function c_p is just added to the current immediate costs. We use separate cost terms for control constraints c_t'' and joint constraints c_t^x , that is, $c_p(\mathbf{x}_t, \mathbf{u}_t) = c_t^x(\mathbf{x}_t) + c_t''(\mathbf{u}_t)$. Here, we will only discuss how to implement the function $c_t''(\mathbf{u}_t)$ for the torque constraints; joint constraints are implemented similarly.

The cost function c_t'' is quadratic in \mathbf{u} and punishes leaving the valid control limits of \mathbf{u} . In order to implement the upper bound \mathbf{u}_{\max} for the torques, we use the following cost function:

$$\begin{aligned} c_t''(\mathbf{u}_t) &= \mathbf{u}_t^T \mathbf{H}_t \mathbf{u}_t + (\mathbf{u}_t - \mathbf{u}_{\max})^T \mathbf{H}_t^U (\mathbf{u}_t - \mathbf{u}_{\max}) \\ &= \mathbf{u}_t^T \mathbf{H}_t \mathbf{u}_t + \mathbf{u}_t^T \mathbf{H}_t^U \mathbf{u}_t - 2\mathbf{u}_{\max}^T \mathbf{H}_t^U \mathbf{u}_t + \mathbf{u}_{\max}^T \mathbf{H}_t^U \mathbf{u}_{\max} \\ &= \mathbf{u}_t^T \mathbf{H}_t \mathbf{u}_t + \mathbf{u}_t^T \mathbf{H}_t^U \mathbf{u}_t - 2\mathbf{u}_{\max}^T \mathbf{H}_t^U \mathbf{u}_t + \text{const}, \end{aligned}$$

where the matrix \mathbf{H}_t contains the quadratic control costs. The constrained costs are only imposed for the control variable u_i if the torque value exceeds the upper bound $u_{\max, i}$. In order to do so, \mathbf{H}_t^U is taken as a diagonal matrix where the i th diagonal entry is zero if $u_i \leq u_{\max, i}$ and nonzero otherwise. The lower bound \mathbf{u}_{\min} is implemented likewise, using an individual diagonal matrix \mathbf{H}_t^L .

We can again implement $c_t''(\mathbf{u}_t)$ as a prior distribution $P(\mathbf{u}_t)$ for the controls:

$$P(\mathbf{u}_t) \propto \mathcal{N}[\mathbf{u}_t | \mathbf{h}_t, \hat{\mathbf{H}}_t], \quad (13)$$

where $\mathbf{h}_t = \mathbf{u}_{\max}^T \mathbf{H}_t^U + \mathbf{u}_{\min}^T \mathbf{H}_t^L$ and the precision $\hat{\mathbf{H}}_t = \mathbf{H}_t + \mathbf{H}_t^U + \mathbf{H}_t^L$. As we can see, the linear term \mathbf{h}_t of the prior distribution $P(\mathbf{u}_t)$ is now nonzero. This yields different message equations.

Joint limits can be imposed similarly by using additional terms as costs for $c_t^x(\mathbf{x}_t)$. However, for joint limits, the update equations stay the same, because $P(z_t = 1 | \mathbf{x}_t)$ already has a nonzero mean, denoted by \mathbf{r}_t in Equation 9.

To derive the messages, we will first integrate out the controls to get the state transition probabilities:

$$\begin{aligned} P(\mathbf{x}_{t+1} | \mathbf{x}_t) &= \int_{\mathbf{u}_t} \mathcal{N}(\mathbf{x}_{t+1} | \mathbf{A}_t \mathbf{x}_t + \mathbf{a}_t + \mathbf{B}_t \mathbf{u}_t, \mathbf{Q}_t) \mathcal{N}[\mathbf{u}_t | \mathbf{h}_t, \hat{\mathbf{H}}_t] d\mathbf{u}_t \\ &= \mathcal{N}(\mathbf{x}_{t+1} | \mathbf{A}_t \mathbf{x}_t + \mathbf{a}_t + \mathbf{B}_t \hat{\mathbf{H}}_t^{-1} \mathbf{h}_t, \mathbf{Q}_t + \mathbf{B}_t \hat{\mathbf{H}}_t^{-1} \mathbf{B}_t^T). \end{aligned} \quad (14)$$

Note that, since the cost function $c_t^h(\mathbf{u}_t)$ contains a nonzero linear term \mathbf{h}_t , we get a new linear term $\hat{\mathbf{a}}_t = \mathbf{a}_t + \mathbf{B}_t \mathbf{H}_t^{-1} \mathbf{h}_t$ in the transition dynamics. The forward and backward messages are the same as Toussaint [16] except that \mathbf{a}_t is replaced by $\hat{\mathbf{a}}_t$, and \mathbf{H}_t by $\hat{\mathbf{H}}_t$.

Like Toussaint [16], we use the canonical representations for the forward and the backward message:

$$\alpha_t(\mathbf{x}_t) = \mathcal{N}[\mathbf{x}_t | \mathbf{s}_t, \mathbf{S}_t],$$

$$\beta_t(\mathbf{x}_t) = \mathcal{N}[\mathbf{x}_t | \mathbf{v}_t, \mathbf{V}_t],$$

$$\phi_t(\mathbf{x}_t) = P(z_t | \mathbf{x}_t) = \mathcal{N}[\mathbf{x}_t | \mathbf{r}_t, \mathbf{R}_t].$$

The messages are represented by Gaussians in canonical form, for which mathematical operations like products are simply performed by adding the linear terms and the precisions. The mean of the belief is given by $b_t(\mathbf{x}_t) = (\mathbf{S}_t + \mathbf{V}_t + \mathbf{R}_t)^{-1}(\mathbf{s}_t + \mathbf{v}_t + \mathbf{r}_t)$ (multiplying three canonical messages and a subsequent transformation to normal form). Furthermore, we use the shorthand $\tilde{\mathbf{Q}}_t = \mathbf{Q}_t + \mathbf{B}_t \hat{\mathbf{H}}_t^{-1} \mathbf{B}_t^T$ for the covariance in Equation 14.

The messages are computed by inserting the state transition probabilities given in Equation 14 in the message-passing Equations 5 and 6. Subsequently, the integrals are evaluated using the propagation rule in Toussaint [17]. The final equations in canonical form are

$$\mathbf{S}_t = (\mathbf{A}_{t-1}^{-T} - \mathbf{K}_t) \mathbf{S}_{t-1} \mathbf{A}_{t-1}^{-1}, \quad (15)$$

$$\mathbf{s}_t = (\mathbf{A}_{t-1}^{-T} - \mathbf{K}_t) (\tilde{\mathbf{s}}_{t-1} + \mathbf{S}_{t-1} \mathbf{A}_{t-1}^{-1} (\hat{\mathbf{a}}_{t-1} + \mathbf{B}_{t-1} \hat{\mathbf{H}}_{t-1}^{-1} \mathbf{h}_{t-1})), \quad (16)$$

$$\mathbf{K}_t = \mathbf{A}_{t-1}^{-T} \mathbf{S}_{t-1} \left(\mathbf{S}_{t-1} + \mathbf{A}_{t-1}^{-T} \tilde{\mathbf{Q}}_{t-1}^{-1} \mathbf{A}_{t-1}^{-1} \right)^{-1}. \quad (17)$$

And for the backward messages we have

$$\mathbf{V}_t = (\mathbf{A}_t^T - \mathbf{K}_v) \tilde{\mathbf{V}}_{t+1} \mathbf{A}_t, \quad (18)$$

$$\mathbf{v}_t = (\mathbf{A}_t^T - \mathbf{K}_v) (\tilde{\mathbf{v}}_{t+1} - \tilde{\mathbf{V}}_{t+1} (\hat{\mathbf{a}}_t + \mathbf{B}_t \hat{\mathbf{H}}_t^{-1} \mathbf{h}_t)), \quad (19)$$

$$\mathbf{K}_v = \mathbf{A}_t^T \tilde{\mathbf{V}}_{t+1} (\tilde{\mathbf{V}}_{t+1} + \tilde{\mathbf{Q}}_t^{-1})^{-1}. \quad (20)$$

To obtain the same mathematical form as in Toussaint [16] one needs to apply the Woodbury identity and reformulate the equations. In contrast to the update message in normal form [16], direct inversions of $\tilde{\mathbf{S}}_{t-1}$ and $\tilde{\mathbf{V}}_{t+1}$ are not necessary in the canonical form, and therefore the iterative updates are numerically more stable.

Finally, in order to compute the optimal feedback controller, we calculate the joint state-control posterior

$$\begin{aligned} P(\mathbf{u}_t, \mathbf{x}_t) &= P(\mathbf{u}_t, \mathbf{x}_t | z_t = 1) \\ &= \int_{\mathbf{x}_{t+1}} \alpha_t(\mathbf{x}_t) \phi_t(\mathbf{x}_t) P(\mathbf{x}_{t+1} | \mathbf{x}_t, \mathbf{u}_t) P(\mathbf{u}_t) \beta_{t+1}(\mathbf{x}_{t+1}) \phi_{t+1}(\mathbf{x}_{t+1}) d\mathbf{x}_{t+1} \\ &= P(\mathbf{x}_t) P(\mathbf{u}_t) \int_{\mathbf{x}_{t+1}} P(\mathbf{x}_{t+1} | \mathbf{x}_t, \mathbf{u}_t) \mathcal{N}[\mathbf{x}_{t+1} | \tilde{\mathbf{v}}_{t+1}, \tilde{\mathbf{V}}_{t+1}] d\mathbf{x}_{t+1}. \end{aligned}$$

The conditional distribution is given by $P(\mathbf{u}_t|\mathbf{x}_t) = P(\mathbf{u}_t, \mathbf{x}_t)/P(\mathbf{x}_t)$, and the solution is

$$P(\mathbf{u}_t|\mathbf{x}_t) = \mathcal{N}\left(\mathbf{u}_t|\mathbf{M}_t^{-1}\left(\mathbf{B}_t^T\mathbf{V}_*\left(\bar{\mathbf{V}}_{t+1}^{-1}\bar{\mathbf{v}}_{t+1} - \mathbf{A}_t\mathbf{x}_t - \hat{\mathbf{a}}_t\right) + \mathbf{h}_t\right), \mathbf{M}_t^{-1}\right),$$

where $\mathbf{V}_* = (\mathbf{Q} + \bar{\mathbf{V}}_{t+1}^{-1})^{-1}$ and $\mathbf{M}_t = \mathbf{B}_t^T\mathbf{V}_*\mathbf{B}_t + \hat{\mathbf{H}}_t$. After a reformulation we can obtain an optimal feedback controller of the form $\mathbf{u}_t = \mathbf{o}_t + \mathbf{O}_t\mathbf{x}_t$ with

$$\mathbf{o}_t = \mathbf{M}_t^{-1}\left(\mathbf{B}_t^T\mathbf{V}_*\bar{\mathbf{V}}_{t+1}^{-1}\bar{\mathbf{v}}_{t+1} - \mathbf{B}_t^T\mathbf{V}_*\mathbf{a}_t + \mathbf{h}_t\right), \quad (21)$$

$$\mathbf{O}_t = -\mathbf{M}_t^{-1}\mathbf{B}_t^T\mathbf{V}_*\mathbf{A}_t. \quad (22)$$

Similarly to Toussaint [16], we use an iterative message-passing approach where we approximate the nonlinear system by linear dynamics, quadratic costs, and a Gaussian noise (LQG) system at the new mode of the trajectory. In Toussaint [16], this is done by using a learning rate on the current modes of the belief. However, in contrast to Toussaint [16], we also need an estimate of the optimal action \mathbf{u}_t in order to impose the control constraints. Using a learning rate on the control action \mathbf{u}_t turned out to be very ineffective, because feedback is attenuated. For this reason we will instead use a learning rate on the feedback controller.

The complete message-passing algorithm, considering state and control constraints, is listed in Algorithm 1. This is a straightforward implementation of Gaussian message passing in linearized systems, similar to an extended Kalman smoother:

ALGORITHM 1: Approximate inference control for constrained systems

Data: initial trajectory $\hat{\mathbf{x}}_{1:T}$, learning rate η

Result: $\mathbf{x}_{1:T}$ and $\mathbf{u}_{1:T-1}$

initialize $\mathbf{S}_1 = 1e10 \cdot \mathbf{I}$, $\mathbf{s}_1 = \mathbf{S}_1\mathbf{x}_1$, $k = 0$, $\hat{\mathbf{o}}_{1:T-1} = \mathbf{0}$, $\hat{\mathbf{O}}_{1:T-1} = \mathbf{0} \cdot \mathbf{I}$;

while $L(\tau)$ *not converged* **do**

for $t \leftarrow 1$ **to** T **do**

 Linearize model: $\mathbf{A}_t, \mathbf{a}_t, \mathbf{B}_t$ using Equation 8

 Compute costs: $\hat{\mathbf{H}}_t, \mathbf{h}_t, \mathbf{R}_t, \mathbf{r}_t$ using Equations 9, 13

for $t \leftarrow 1$ **to** T **do**

 Forward messages: $\alpha_t(\mathbf{x}_t)$ using Equations 15 – 17

for $t \leftarrow T - 1$ **to** 1 **do**

 Backward messages: $\beta_t(\mathbf{x}_t)$ using Equations 18 – 20

for $t \leftarrow 1$ **to** $T - 1$ **do**

 Feedback controller: $\mathbf{o}_t, \mathbf{O}_t$ using Equations 21, 22

if $k == 0$ **then**

$\mathbf{u}_t = \mathbf{o}_t + \mathbf{O}_t\mathbf{x}_t$

else

$\hat{\mathbf{o}}_t = (1 - \eta)\hat{\mathbf{o}}_t + \eta\mathbf{o}_t$

$\hat{\mathbf{O}}_t = (1 - \eta)\hat{\mathbf{O}}_t + \eta\mathbf{O}_t$

$\mathbf{u}_t = \hat{\mathbf{o}}_t + \hat{\mathbf{O}}_t\mathbf{x}_t$

$\mathbf{x}_{t+1} = \mathbf{A}_t\mathbf{x}_t + \mathbf{a}_t + \mathbf{B}_t\mathbf{u}_t$

$k = k + 1$

In Toussaint [16] and Rawlik et al. [13], more time-efficient methods are presented, where for each time step the belief is updated until convergence, in contrast to updating all messages and iterating until the intrinsic costs $L(\tau)$ converge. The computational benefits of such an approach still need to be evaluated for our messages.

Appendix 2: Four-Link Robot Model Specifications

For the four-link model, the eight-dimensional state $x = [\phi_1, \dot{\phi}_1, \phi_2, \dot{\phi}_2, \phi_3, \dot{\phi}_3, \phi_4, \dot{\phi}_4]$ is composed of the ankle, the knee, the hip, and the arm positions and velocities. In our experiments, the velocities are instantaneously affected by the applied force F [2]. These initial velocities and the valid joint angle ranges are shown in Table 1.

Table 1. Joint angle configurations where a robot gets pushed by a force F .

Joint	Initial velocities	Lower bound	Upper bound
Ankle	$+1.2 \times 10^{-2}F$	-0.8	0.8
Knee	$-7.4 \times 10^{-2}F$	-0.05	2.5
Hip	$+5.1 \times 10^{-2}F$	-2.0	0.1
Arm	$-4.2 \times 10^{-2}F$	-0.6	3.0

In addition, we use multiple initial joint angles. The ankle angle is given by $\phi_1 = -\alpha$, the knee angle by $\phi_2 = -\alpha f$, the hip angle by $\phi_3 = \alpha f$, and the arm joint angle by $\phi_4 = 2\alpha f$, where α is sampled from $\alpha \sim \mathcal{U}[0, 7]\pi/180$ and the linear factor is uniformly distributed with $f \sim \mathcal{U}[-3.3, -2.3]$. Four example initial states are illustrated in Figure 2.

Copyright of Artificial Life is the property of MIT Press and its content may not be copied or emailed to multiple sites or posted to a listserv without the copyright holder's express written permission. However, users may print, download, or email articles for individual use.