Biological Actuators Are Not Just Springs

Investigating Muscle Dynamics and Control Signals

Thomas Buehrmann and Ezequiel Di Paolo

Centre for Computational Neuroscience and Robotics (CCNR),
Department of Informatics, University of Sussex, Brighton, BN1 9QG, UK
{T.Buehrmann, Ezequiel}@sussex.ac.uk

Abstract. While there is a trend in current robotics towards more biologically inspired actuators, most work emphasizes the elastic property of muscles and tendons. Although elasticity plays a major role in many forms of movements, particularly walking and running, other features of animal muscles might also affect or even dominate movement dynamics. In this paper we use the Hill-type muscle model, common in biomechanics, to investigate the relationship between muscle dynamics and control signals in simple goal-directed movements. We find that the various non-linearities of the model lead to desirable properties with regard to controllability, such as increased stability and robustness to noise, independence of position and stiffness, or near linearity in search space. We conclude that in our attempt to create robots exhibiting the same flexibility and robustness as animals we have to seek a balance between the complexity of actuators and the extent to which their natural dynamics can be exploited in a given task.

1 Introduction

It is not unreasonable to assume that the qualitative properties found in the majority of animal muscles have been selected for by evolution for their adaptive advantage with respect to the generation of movement. The benefits of elasticity for instance have been recognized for a while and robotic systems have been built that incorporate springs in order to gain higher force fidelity, low impedance, low friction, good force control bandwidth or energy conservation [9]. More generally, the concept of 'preflexes' [5] summarizes the idea that the intrinsic dynamics of a musculoskeletal system alone can be sufficient for self-stabilization [11] and that they could be tuned such that higher level motor control becomes easier because joint level dynamics don't have to be accounted for. One of the tenets of the embodied approach to behavior, Pfeifer's notion of 'morphological computation' [7], also expresses the idea that materials can sometimes take over some of the processes normally attributed to control mechanisms. In view of this embodied perspective, in this paper we propose to investigate the different qualitative properties of muscles and their effect on motor control. We are interested in the question to what extent it would be useful to deal with the complexities of biological muscles and whether it makes the task of designing robotic control systems

S. Nolfi et al. (Eds.): SAB 2006, LNAI 4095, pp. 89–100, 2006.

easier or harder. Specifically, we use a simulation of an antagonistic muscle pair acting on a hinge-joint to ask how the non-linear visco-elastic properties of muscle affect stability, how they allow for fast but appropriately damped movements and how they relate to control signals.

2 Muscle Model

Muscles are different from any current robotic actuator in many ways. Not considering effects such as hysteresis, the instantaneous force a muscle produces is a complex non-linear function not only of its activation but also of its length and velocity. In biomechanics, the most commonly used model of skeletal muscle is the so-called Hill-type model [12]. Its basic configuration (fig.1) uses dimensionless constitutive relationships to describe the muscle's visco-elastic properties. Specifically, it consists of a contractile element (CE) producing force as a function of length and velocity; a parallel elastic element (PE) exponentially resisting stretch beyond resting length; and a series elastic element (SE), the tendon. Figure 2 depicts the non-linear nature of these elements. As is shown, active force production (F_a) peaks at an optimal length of L_0, and decreases to either side. Passive tension (F_p) is the result of an exponential spring with a slack length also around L_0. The dependence of muscle force on velocity (F_v) is described by Hill's relationship. If a muscle is shortening, its force generation ability drops as velocity increases and reaches zero at velocity v_{max} (heavier loads can be lifted less rapidly). If it is lengthening in contrast, force increases when compared to the maximum in statics F_0^M.

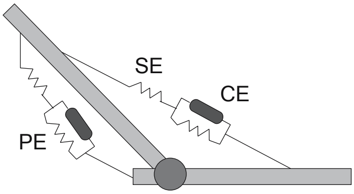

Fig. 1. Schematic of antagonistic muscle pair acting on hinge-joint

The overall force output of the muscle model is specified by

$$F^M = aF_0^M F_v(v^M)F_a(L^M) + F_p(L^M) \tag{1}$$

where activation a scales both F_a and F_v, which themselves combine in a multiplicative way. The term activation refers to the low-pass filtered (neural) input to the muscle and describes its slow activation dynamics. F_0^M is the muscle's maximal force at zero velocity.

The result of eqn.(1) is a normalized force-length-velocity surface scaled by activation as depicted in figure 3. Any movement in this picture corresponds to a 'walk' on that surface. Already complex, in any realistic scenario muscle length,

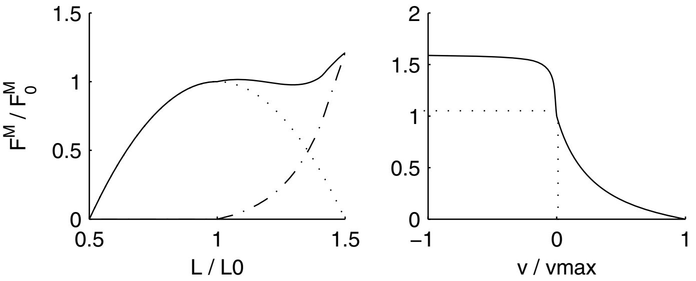

Fig. 2. Dimensionless muscle force as a function of length (left) and velocity (right). The former is the sum of a passive exponential-to-linear elasticity resisting lengthening of the muscle (dash-dotted) and a quadratic function with a maximum at resting length describing the active generation of force (dotted).

velocity and force would probably feed back into the control signal, thereby creating a continuously changing surface. The complete system simulated in the following experiments consists of two antagonistic muscles around a single hinge-joint. Its dynamics are completely determined by the torque produced at the joint through the muscles, as well as the physical properties of the limbs that the joint connects (such as inertia). For simplicity and because we're mainly interested in the effects of the various non-linearities of the muscle model on its dynamics we have made several abstractions. Moment arms of the muscles relative to the joint they actuate are assumed to be constant; gravity in some experiments is zero, which allows us to remove static parts from the control signals and focus on the dynamics instead; no tendon is present. The latter can be justified in elbow movements for example, where the ratio of tendon to muscle lengths is such that tendons have only minimal effect [12]. [1]

2.1 Muscle Dynamics

With regards to the behavior of the muscle model several observations can be made from fig.2 and 3. Firstly, in Hill's force-velocity relationship the slope, that is the rate of change of force, is highest around zero velocity. Hence muscles have the desirable property of being damped the most at rest, while being less damped when moving fast. The slope, and thus damping, will also increase with the level of activation. The same holds for the active force-length relation, hence stiffness will increase with activation as well. It should be noted that the parallel elasticity (the 'spring') of two antagonistic muscles could in principle provide for stability if those muscles were arranged such that over most of the joint's workspace their lengths would be greater than the spring's slack length. Obviously though, this

---

[1] Because we're only interested in qualitative results here, and not particularly simulations of human movements, details about the implementation of the non-linear functions as well as values for the muscle specific parameters F_0, L_0 and v_{max} and limb dimensions are ommitted.

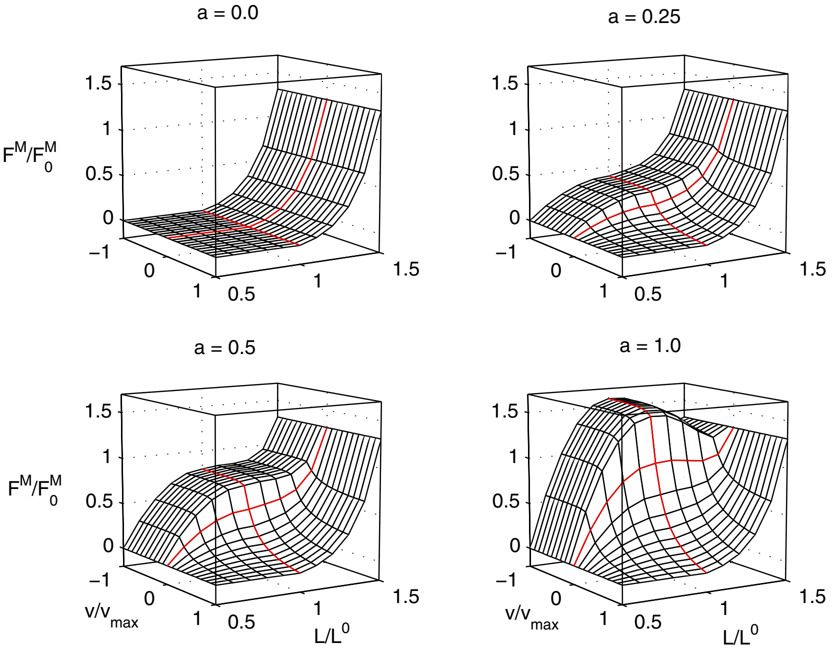

Fig. 3. Force-length-velocity surfaces resulting from eqn.1 for activations of 0.0, 0.25, 0.5 and 1.0. Horizontal axes correspond to normalized length (length divided by resting length) and normalized velocity (velocity divided by maximum velocity).

would make movements rather inefficient because the system would always pull against the opposing muscle's resistance. Indeed, in nature many muscles seem to operate mostly on the ascending limb of the force length curve [3,6], indicating lengths being in a range where the passive elasticity won't dominate the muscle's dynamics. Stability therefore must have a different origin.

Control of position and stiffness. Figure 4 shows how the force-length relationship of two antagonistic muscles interact to create an equilibrium position (EP). The muscles are arranged symmetrically and such that their length, measured over the range of joint positions θ, varies between 0.6 and 1.1. Because the resting positions of both muscles are shifted towards the joint extremes, no elastic resisting forces are created in the midrange. Hence, without activation the system is truly passive and and does not behave like a spring. Activation however creates a stable equilibrium position. While increased co-contraction stiffens the joint at the EP, differential activation will shift its position. The model thus allows for independent control of joint position and stiffness. Note that if the two muscles were springs whose force production depended linearly on displacement and whose resting length was shifted by activation, no increase in stiffness would result from co-activation. The net force produced would only depend on the individual spring stiffnesses and the amount of perturbation from equilibrium.

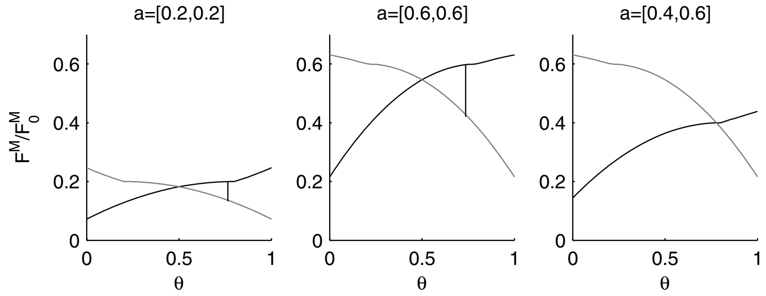

Fig. 4. Force-length relationships of an antagonistic muscle pair. Intersections correspond to the systems equilibrium point (EP). Shifts from the EP will be resisted by a net force equal to the difference between the two curves (vertical line). While co-activation doesn't change the EP, the slope of both force-length curves increases, corresponding to increased stiffness of the joint. A change in the difference between activations will move the EP position.

Rejection of perturbations. From the model it is easily predicted that the viscous element will enhance the system's stability with respect to perturbations. This effect is illustrated in fig.5, where transient loads are applied to the joint at increasing levels of co-activation. Not surprisingly, the model including viscosity shows damping without oscillations and resistance to loads that increases with co-contraction. This rejection of perturbations results from an interplay of both the F_v as well as F_a relationships. While co-contraction creates the EP and stiffens the joint as described above, the steeper slopes in F_v increase its viscosity. Hence the net rejecting force will generally be comprised of a dynamic F_v contribution as well as a static F_a element.

Fast movements. A desirable feature of muscle models including Hill's viscosity term is that they allow for the generation of movements that are fast but appropriately damped so as not to produce significant overshoot. This can be demonstrated for example within the context of the equilibrium point hypothesis (EP-hypothesis) [2]. The general idea is that the neuro-musculoskeletal system creates an equilibrium point that can be shifted using central commands in a simple, e.g. linear, fashion without needing to take into account the actual dynamics. Specifically, in the λ-formulation of the EP-hypothesis, central commands are believed to set the reflex threshold for motoneuron recruitement, such that muscles are activated proportionally to the amount of muscle stretch beyond the setpoint. In this formulation, muscle activation a is described by:

$$a = [L - \lambda + \mu v]^+ \tag{2}$$

where λ is the commanded reflex threshold, v the muscle's velocity and μ the gain of reflex damping. Figure 6 presents a trajectory for the antagonistic muscle pair controlled using this proportional-derivative mechanism (PD). Note that for both, step and ramp shaped commands, the muscle's non-linear relations and

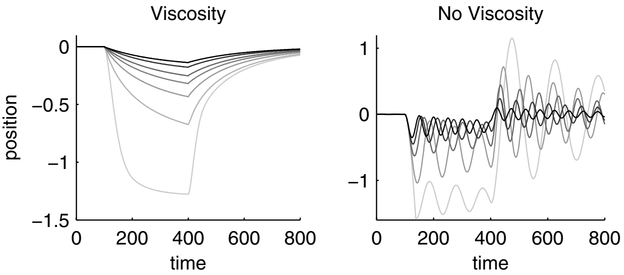

Fig. 5. Response of antagonistic muscle pair to transient external load (between t=200 and t=400) at different levels of co-activation. Left: joint position of model including viscosity, Right: without viscosity. Increased co-activation (darker plots) leads to increased level of stiffness and damping. Only when viscosity is present however, is the resisting force appropriate for self-stabilization. Otherwise oscillations arise.

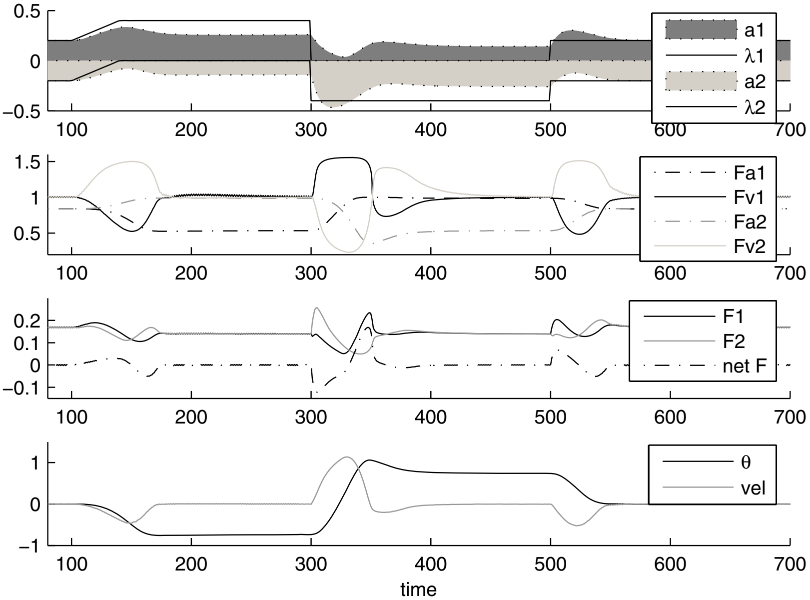

Fig. 6. Trajectories over time of muscles and joint controlled by λ-model. Indices in legend indicate agonist (1) and antagonist (2). Top: input signal λ and resulting activations. Second row: values of F_a and F_v. Third plot: force output of muscles as well as net force. Bottom: joint position and velocity. Commanded reflex thresholds λ change in form of a ramp at time t=100 and as step-functions of different amplitudes at t=300 and t=500. The system starts with initial co-activation of 0.2.

low-pass characteristic combine to sculpt a bi-phasic force trajectory with one peak for acceleration and another for deceleration of the joint. While a pure PD controller would only produce net breaking forces in the case of overshoot, here the breaking starts significantly before overshoot occurs. It is thus fair to say that the muscle model can generate complex trajectories from simple control signals that produce fast but smooth shifts in joint position (similar results with a different model are presented in [4]).

3 Control Signal Optimization

The gain in stability described above also has interesting consequences for the optimization of control signals. To illustrate this point we implemented a simple control strategy that activates each of the two muscles using a square pulse of a given amplitude (a_1, a_2) and duration (d_1, d_2). A fifth parameter (t_2) specifies the time between the onset of agonist activation (fixed) and that of the antagonist. This allows us to look at the space created by evaluating all possible strategies against an optimality criterion, hereafter called the fitness landscape.

Flexibility. Figure 7 shows a fitness landscape in which the criterion consisted of reaching and stopping at a target position of 45° flexion (starting from 0°) at any point during a 2 second trial. In order to show the whole search space we somewhat arbitrarily fixed the amplitudes a_1 and a_2 to values of 0.2. Several interesting observations can be made from this case of unconstrained goal-directed movement. Firstly, the region of good performance spans a considerable range in each of the three remaining dimensions. One can pick almost any value for one of the parameters and will find a combination for the other two that produces a good strategy. In other words, there is a continuum of valid strategies all of which will move the joint towards the desired position, but each having different kinetic or kinematic properties. Movements will differ in terms of velocity, stiffness or energy required. In fact, the point marked B corresponds to the fastest movement in this space, while point C marks the one using least energy (measured as the integral over muscle activation). Thus, compared to the stereotypical behavior of e.g. a PD controller, by using this model one gains immense flexibility with respect to details of a movement, while introducing only few parameters to be chosen (by either a controller or a more constrained optimization procedure). Secondly, although the model is highly non-linear in all its properties, good performance within the fitness landscape is found along rather linear regions. This simple relationship between parameters should make it easy to create a controller that finds and moves along the range of all optimal strategies.

Robustness. In terms of control signal optimization the viscous property of the Hill-type muscle model also shows as increased robustness to noise or increased ‘searchability’ of the fitness landscape, a property of interest for evolutionary robotics for instance. Fig.8 compares the fitness landscapes of the muscle model with and without the viscosity term for an opimization that maximizes velocity while minimizing overshoot. The slices shown were produced by finding for each model the global peak in the 5-dimensional parameter space $(a_1, a_2, d_1, d_2, t_2)$ and subsequently fixing two of them (amplitudes a_1, a_2) to the found values. The resulting slices show the fitness landscape around the optimum in the remaining

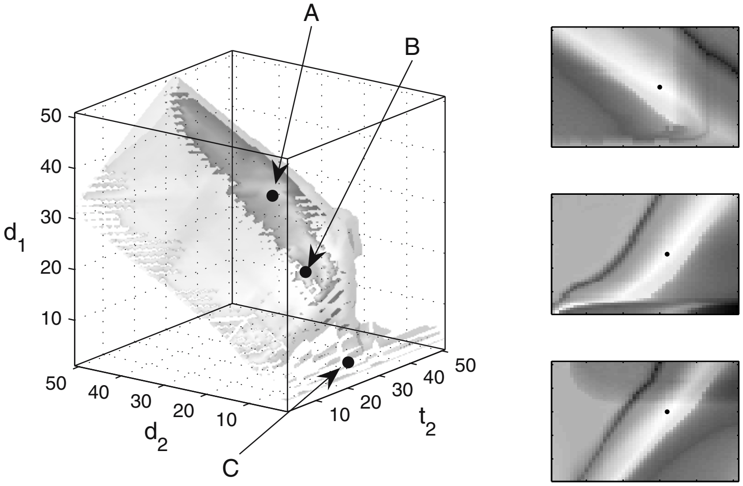

Fig. 7. Left: isosurface of fitness landscape at fitness levels of 95% (dark) and 80% (bright). Point A shows the overall peak of the surface. B corresponds to a movement that maximizes velocity, while C minimizes energy. Right: slices through the peak of the same fitness landscape.

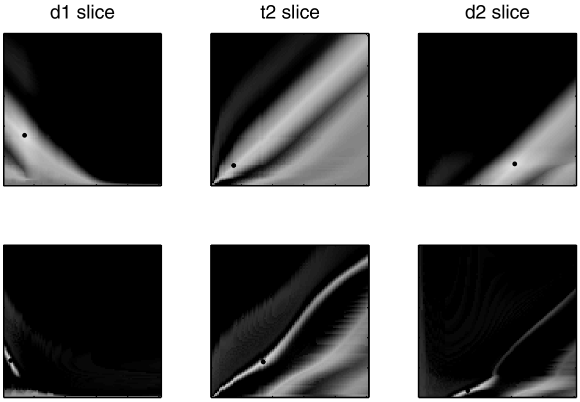

Fig. 8. Slices through peaks of fitness landscapes for maximizing velocity while minimizing overshoot. Top row includes viscosity, bottom row does not.

three dimensions. As can easily be seen, without viscosity the regions of good fitness become much more narrow. For the optimization procedure this means increased difficulty of finding the global optimum. It can also be interpreted though as robustness to noise in the control signal. In the viscous model a slight perturbation away from the optimum will still produce relatively good results, while in the non-viscuous case performance is easily lost completely. Intuitively this is easy to understand. In the non-viscous case, the antagonist activation has

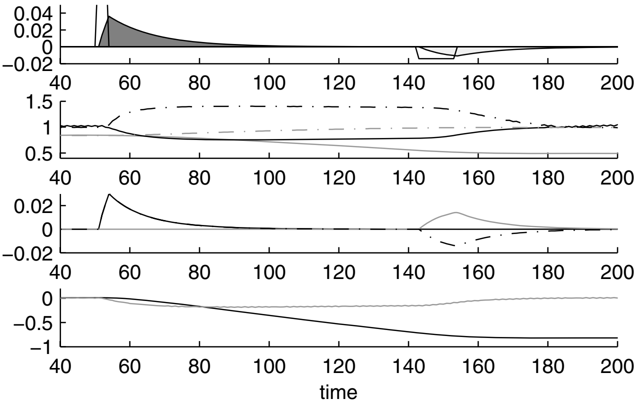

Fig. 9. Trajectory resulting from typical control signals evolved to minimize energy. Plots correspond to fig.6, i.e. top: activations, 2nd plot: F_a, F_v, 3rd plot: muscle forces, bottom: position and velocity.

to be precisely timed and scaled such that at the target position forces cancel out exactly and the joint comes to a stop. Any remaining forces not counteracted completely by the antagonist will move the joint away from the target. In the viscous model however, because of its damping effect, small remaining forces will fade quickly and the joint will come to a stop near the target position.

Efficiency. Motorized actuators have to be powered throughout a movement. Even compliant actuators will have to make motors move to simulate a zero force trajectory, i.e. a purely passive swing. Muscles, however, allow for more efficient movement through bi/tri-phasic pulse patterns. Minimal muscle activations are sufficient to accelerate and decelerate the joint towards a desired position. This is possible, however, only because antagonistic muscles don't work like springs. That is, in their passive state they don't have to work against each other's resistance. Figure 9 is an example of control signals optimized for minimal energy use. Clearly, throughout a large part of the movement neither muscle produces any force and the joint is passively swinging towards its desired position.

Multijoint movement. The movements and control signals presented so far are clearly oversimplified when compared to natural movements involving many interacting joints. It is striking though that simple square pulses, appropriately scaled and timed, allow for well-behaved movement trajectories when combined with non-linear muscle properties. In order to investigate if the increased robustness and flexibility also translates to more complex scenarios we used the same approach of control signal optimization to generate motions of two joints (elbow and shoulder). We also enabled gravity and included a static part in the control signals that could compensate for its effect. Figure 10 presents optimized trajectories in two different conditions. The elbow joint is always required to flex to a position of 45°. However, in scenario 1 the shoulder moves in the opposite

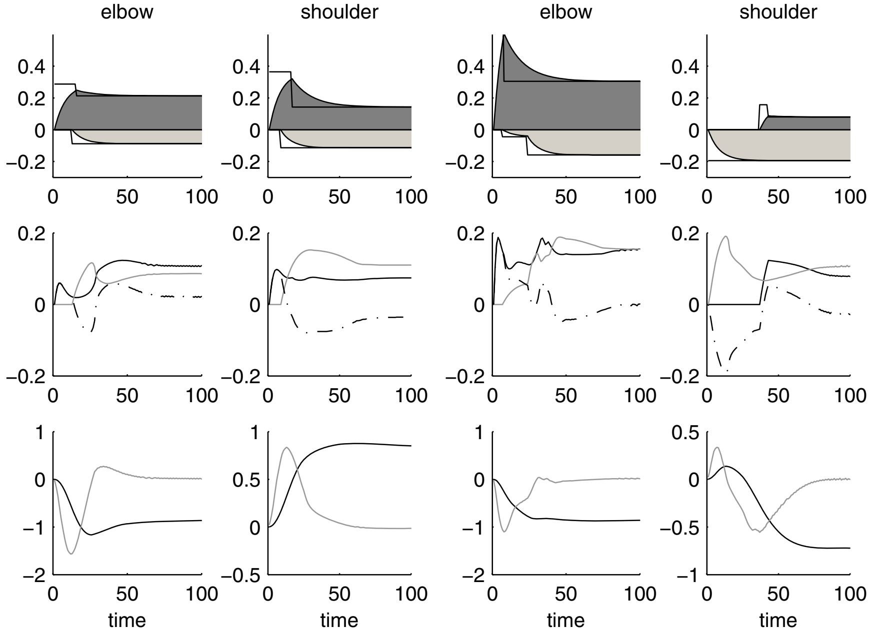

Fig. 10. Shoulder and elbow trajectories optimized for goal-directed 2-joint movements maximizing velocity while minimizing energy. First two columns correspond to scenario 1 (synergistic), the columns on the right to scenario 2. Top: activations, middle: muscle forces, bottom: position and velocity.

direction while in scenario 2 it moves in the same direction. Both cases were easily evolved and produced trajectories whose final positions corresponded to the desired targets. The figure shows that in the first case the velocity profiles resemble smooth bell-shapes, while they are more jerky in the second case. The reason for this effect are the interaction torques arising from the mechanical coupling of the two joints. In the first scenario movement of the shoulder creates interaction torques in the elbow that are ‘synergistic’, i.e. support the intended movement, while in the second case the torques counteract movement in the desired direction. It is thus obvious that the simple open-loop control used here is insufficient in some circumstances. In fact, it is one of the big open questions in motor control whether the (human) central nervous system uses an internal model of the body to calculate control signals that account for its dynamics or if a well-designed neuro-musculoskeletal system itself could perform the neccessary ‘morphological computations’. Independent of the case of human movements, both a feedback controller such as the λ-model as well as appropriate open-loop signals created using internal models could generate the desired movements. While the former stands out for its simplicity, the latter more easily could exploit interaction torques rather than resist them. In both cases though a muscle’s intrinsic stability should be beneficial.

4 Conclusion

As shown above, the non-linear behavior of an antagonistic muscle pair produces many desirable properties with respect to embodied motor control; properties which at first could seem surprising, given the complexity of the model. It allows for independent control of position and stiffness; its flexibility allows for a continuum of control strategies differing in various kinematic and kinetic features; it creates 'nice' search spaces for the optimization of control signals; non-linear damping allows for fast movements with precise breaking; it can produce efficient movements by using only phasic activation; and simple control signals can create complex but well behaved trajectories. In addition, the global dynamics of the system can be tuned with only few parameters that describe the shape of the muscle's length and velocity dependence as well as their geometry. We thus argue that versions of the Hill-type or similar muscle models strike an ideal 'ecological balance' [8]. Robots equipped with similar actuators should be able to trade an increase in the complexity of morphology for a reduction in the complexity of the control system. Several attempts exist at building such muscle-like actuators for robots and prostheses. Most prominent are series elastic actuators [9], McKibben-style pneumatic actuators [10] and electroactive polymer actuators [1]. More research seems necessary though to achieve the right combination of viscoelastic properties in an efficient package applicable to multiarticulate robots.

The fact that hill-type muscle models are generic in the sense of needing only few parameters to implement particular types of muscles make them particularly interesting for evolutionary robotics. In fact, if one is looking at more complex motor behaviors such as multijoint reaching or walking, it is not only possible but necessary to co-evolve or co-adapt muscle morphology, skeletal geometry and the control system. Only then is it possible to tune the system's dynamics to be appropriate for a given task.

We'd like to emphasize that the experiments presented herein are not meant to be models of how animals actually control their movements. In particular, we do not make any claim about the relative importance of the brain or the musculoskeletal system in motor control, but emphasize that the latter exhibits intrinsic dynamics which, unless entirely suppressed by the brain, do play a role in natural movements. The forms of control have been chosen only for their simplicity to show the dynamics of the muscle model in the most general context. However, in ongoing work we are concerned with neural control in more realistic tasks. To this end we are evolving dynamical neural 'reflex' networks that try to exploit the natural dynamics of the musculoskeletal system in order to simplify control as seen from higher levels. It can be shown that the right muscle model coupled to a controller in a closed loop can function as a source of force, as a spring, a servo of position or velocity or combinations of these. We believe that well-designed neuro-musculoskeletal systems enable higher levels to set up the right global dynamics for the task at hand, effectively choosing between different modes of control, such that subsequent control signals can make use of the particularities of task, environment and body.

References

1. Bar-Cohen, Y. (2001). Electroactive Polymer (EAP) Actuators as Artificial Muscles - Reality, Potential and Challenges. SPIE Press, Vol. PM98.
2. Feldman, A., Ostry, D., et al. (1998). Recent tests of the equilibrium-point hypothesis (lambda model). Motor Control 2: 189205.
3. Garner, B. A. and Pandy, M. G. (2003). Estimation of Musculotendon Properties in the Human Upper Limb. Annals of Biomedical Engineering 31(2): 207-220.
4. Gribble, P. L., Ostry, D. J., Sanguineti, V. and Laboissiere, R. (1998). Are Complex Control Signals Required for Human Arm Movement? J. Neurophysiology 79: 1409-1424.
5. Brown, I.E. and G.E. Loeb. A reductionist approach to creating and using neuromusculoskeletal models. In: Biomechanics and Neural Control of Posture and Movement. J.M. Winters and P.E. Crago (ed.) Springer Verlag, New York, 2000.
6. Murray, W. M., Buchanan, T. S., et al. (2000). The isometric functional capacity of muscles that cross the elbow. Journal of Biomechanics 33(8): 943-952.
7. Pfeifer, R. and Iida, F. (2005). Morphological computation: Connecting body, brain and environment. Japanese Scientific Monthly 58(2): 48-54.
8. Pfeifer, R. (1996). Building ‘fungus eaters’: Design principles of autonomous agents. In: From Animals to Animats 4: Proceedings of the Fourth International Conference on Simulation of Adaptive Behavior. Maes, P., Mataric, M.J., Meyer, J.A., Pollack, J. and Wilson, S.W. (ed.) MIT Press.
9. Pratt, J.E. and Krupp, B.T. (2004). Series Elastic Actuators for Legged Robots. SPIE 2004.
10. Verrelst, B. et al. (2005). The pneumatic biped ‘lucy’ actuated with pleated pneumatic artificial muscles. Autonomous Robots, vol. 18, pp. 201213, 2005.
11. Wagner, H. and Blickhan, R. (1999). Stabilizing function of skeletal muscles: an analytical investigation. J Theor Biol 199: 163-179.
12. Zajac, F. E. (1989). Muscle and tendon: properties, models, scaling, and application to biomechanics and motor control. CRC Crit. Revs biomed. Engng 17: 359-411.