

Is Growing Good for Learning?

Steve Heim, Alexander Spröwitz

Dynamic Locomotion Group, Max Planck Institute of Intelligent Systems, Germany

{heim,sprowitz}@is.mpg.de

1 Introduction

In both living organisms and human-made machines, we are interested in how morphology (or mechanical design) and neural control (or control) are combined to reach a desired behavior. Engineers tend to split the mechanical and controller design into two separate steps, since it is difficult to analyze the effect of a controller without already having model for the mechanical dynamics. In recent years, control engineers have tried more and more to *exploit the natural dynamics*: this means, the controller should not overpower what the passive mechanical system naturally tends to do, but rather simply push it in the right direction then let the natural dynamics take over. This concept was pioneered in the engineering field by Tad McGeer[5], and also shown to be prevalent in biology[3][7]. This concept highlights the importance of mechanical design on controller design[6], especially for controller-design algorithms such as reinforcement learning where stable natural dynamics can provide a smooth gentle slope to help the algorithm to converge[9].

2 Growth and Natural Dynamics

We aim to address not only the neuro-mechanical relationship between controller design and natural dynamics, but also highlight the related aspect of growth. Indeed most animals go through a critical phase of motor-control learning during early development, a period when their morphology and therefore natural dynamics change drastically and very rapidly. How this change affects the learning process is not yet well understood. We take loose inspiration from this and define growth as *time-dependent natural dynamics*, as illustrated in the following canonical equations of motion:

$$M(q(t),t)\ddot{q}(t) = C(q(t),\dot{q}(t),t) + B(q(t),t)u(t) \quad (1)$$

Where $q(t)$ is the vector of generalized coordinates and $u(t)$ are the control inputs, in this case forces and torques. M , C and B are the inertia matrix, the differentiable force vector (including gravitational, coriolis and other forces) and control matrix, respectively. Note that in this case M , C and B are time-dependent. While algorithms such as reinforcement learning are known to be able to handle changing system dynamics [8], these changes are generally viewed as an additional challenge the controller has to handle, outside the

influence of the engineer. Instead we are interested in how an ad-hoc pre-designed change over time of the natural dynamics can help improve the convergence of learning motor control tasks. In this sense, we are not only interested in natural growth in which an individual generally gains in mass, power-output and complexity, but a generalized notion of growth, as any directed change over time of the natural dynamics. We specify here three initial types of growth we will initially test, for which we plan to show preliminary results at AMAM.

2.1 Dimensionality Reduction

The curse of dimensionality is a common challenge in motor control. Starting with a low-dimensional system and gradually "growing" the additional dimensions therefore has a lot of potential to simplify the learning process. This is related to the freezing-and-freeing phenomenon[1] observed in human learning. While reducing the dimensionality of a problem has clear benefits in reducing the search space, a naive reduction might not be as beneficial as expected. We expect to find strategies for choosing the order in which dimensions should be removed/added to the system.

2.2 Decoupling

A second concept is to maintain the same total dimensionality but *decouple* certain degrees of freedom from others, as discussed in [4]. This essentially splits up the problem into multiple problems, such as body-pitch balance and forward movement, of lower dimensionality that can be solved individually. In a second step, the coupling terms in the dynamics can be gradually introduced, and the the required motor control necessary to deal with this new coupling can be learned. This approach sounds very promising in theory, however it is not trivial to determine which dimensions can effectively be decoupled, nor how or if the coupling term can be learned after a sub-dimensional controller has been learned in the decoupled case.

2.3 Actuator Bounds

As animals grow, their muscle mass distribution, inertia, etc. change drastically. These changes influence the effective mechanical advantage[2] as well as the needed the required power output. In a simplistic view this can be approximated as changing the bounds of the actuator output. While this is typically seen as a limitation, it could also provide unforeseen benefits, especially when performing exploratory learning in actual hardware.

3 Conclusion

Through systematic testing of different "growth cases" we aim to identify important details on how morphology influences controller design, as well as the dynamics of learning in general. This could shed light on developmental motor learning, as well as provide ideas for more effective robot controller designs, especially in model-free learning. We plan to show initial results in simulation by the time of AMAM, and thereafter also test these out in actual hardware.

References

- [1] Luc Berthouze and Max Lungarella. "Motor skill acquisition under environmental perturbations: On the necessity of alternate freezing and freeing of degrees of freedom". In: *Adaptive behavior* 12.1 (2004), pp. 47–64.
- [2] Andrew A Biewener. "Biomechanical consequences of scaling". In: *Journal of Experimental Biology* 208.9 (2005), pp. 1665–1676.
- [3] Reinhard Blickhan. "The spring-mass model for running and hopping". In: *Journal of biomechanics* 22.11-12 (1989), pp. 1217–1227.
- [4] Steve W Heim et al. "On designing an active tail for legged robots: simplifying control via decoupling of control objectives". In: *Industrial Robot: An International Journal* 43.3 (2016), pp. 338–346.
- [5] Tad McGeer et al. "Passive dynamic walking". In: *I. J. Robotic Res.* 9.2 (1990), pp. 62–82.
- [6] Daniel Renjewski et al. "Exciting engineered passive dynamics in a bipedal robot". In: *IEEE Transactions on Robotics* 31.5 (2015), pp. 1244–1251.
- [7] G Sachs et al. "Experimental verification of dynamic soaring in albatrosses". In: *Journal of Experimental Biology* 216.22 (2013), pp. 4222–4232.
- [8] Richard S Sutton and Andrew G Barto. *Reinforcement learning: An introduction*. Vol. 1. 1. MIT press Cambridge, 1998.
- [9] Russ Tedrake, Teresa Weirui Zhang, and H Sebastian Seung. "Learning to walk in 20 minutes". In: *Proceedings of the Fourteenth Yale Workshop on Adaptive and Learning Systems*. Vol. 95585. 2005, pp. 1939–1412.