

2014-10

Development of joint stiffness and learnability

Kuppuswamy, N

<http://hdl.handle.net/10026.1/10037>

10.1109/devlrm.2014.6983016

4th International Conference on Development and Learning and on Epigenetic Robotics

IEEE

All content in PEARL is protected by copyright law. Author manuscripts are made available in accordance with publisher policies. Please cite only the published version using the details provided on the item record or document. In the absence of an open licence (e.g. Creative Commons), permissions for further reuse of content should be sought from the publisher or author.

Development of Joint Stiffness and Learnability

Naveen Kuppaswamy

Dept. of Robot. Brain and Cognitive Sci. (RBCS),
Istituto Italiano di Tecnologia, Genova, Italy
Email: naveen.kuppaswamy@iit.it

Christopher M. Harris

Center for Robotics and Neural Systems (CNRS),
Plymouth University, United Kingdom

Abstract—One of the fundamental problems in developmental robotics relates to the progressive spontaneous acquisition of motor abilities by an organism. Throughout this process, the speed of acquiring abilities, which we term ‘*learnability*’, is strongly limited by the dimensionality of the sensori-motor space; this in turn could affect the survival of an organism. In this paper, we tackle the problem of dimensional change during development using a framework of control dimensionality reduction based on nonlinear system balancing. Using a set of internal models of behaviour of increasing dimensionality, we show that joint-stiffness regulation can be used to ensure optimal development of motor skills. This is quantified as a maximisation of internal model accuracy at intermediate stages of learning. We test our approach in a simulation of a human arm modelled as a 2 link kinematic chain performing point-to-point and via-point reaching tasks. We then analyse optimal joint-stiffness development towards facilitating effective dimensional change and compare two strategies, (i) uniform development and (ii) proximo-distal development, i.e. variation of only the distal joint stiffness. Our results indicate that latter strategy, although lower in accuracy is a simpler approach towards learnability regulation. The implications of the model and the results for biological motor control and robotics are then discussed.

I. INTRODUCTION

In the study of motor behaviour development in nature, one of the key questions is that of how the neural control mechanisms of organisms cope with the large neuro-mechanical Degrees of Freedom (DoF). While large dimensionality affords an organism evolutionary advantages in tackling with novel environmental contexts, it is unclear how

organisms cope with the difficulties of real-time control and motor learning in the resulting high-dimensional state spaces. Given the importance of accurate (optimal) motor behaviours for survival it seems likely that there should be a premium on learning as quickly as possible for a given level of competence and task complexity. This speed of learning, which we call ‘*learnability*’ [8], depends ultimately on the dimensionality of the state-space and there is increasingly a consensus that some form of manipulation of dimensionality is taking place essential in order to regulate learnability [11]. Naturally, any control over dimensionality must be manifest in the neural and/or structural architecture of the organisms, and must be inheritable itself (i.e. coded in DNA) [6]. It is therefore important to understand the mechanisms underlying dimensional change during motor development.

Bernstein[1] was one of the earliest to identify the necessity for some form of developmental strategy of dimensionality regulation to reduce the complexity of motor skill learning. His *three stage learning model*, consisted of (i) initial DoF freezing, (ii) later DoF unfreezing and (iii) exploitation of reactive phenomena. Several investigators have since sought to validate this theory in natural [11], [2] and synthetic contexts [3] and the strategy is usually modelled as progressive state-space exploration for acquiring task-specific kinematic transformations [12]. Despite some experimental evidence [13], several open questions remain. First, the terms “freezing”, “unfreezing” and “exploitation of reactive phenomena” all are open to interpretation and

joint stiffness regulation is one possible mechanisms suggested. Secondly, no single study has so far managed to examine all three stages of learning in a skill, mainly due to the time scales within which such a mechanism seems to function. Lastly, focus is only paid to resulting variations in the kinematics of behaviours – this approach does not explain how optimal human motor behaviours are acquired in a dynamic context. The dynamical systems perspective [11], [10] represents an alternative modelling approach however the nature of the dimensional change (increase vs. decrease) remains inconclusive and the role of neuro-mechanical effects such as joint stiffness is not easy to model. In contrast, we aim to build a synthetic theory of the dimensional change phenomenon by employing a control-theoretic approach based on state-space models and model/control order (dimensionality) reduction. The advantages of this approach is that neuro-mechanical properties can be explicitly modelled in the dynamic equations and thus analysed for their effects [9].

Our approach is based on the following observations. Firstly, dimensional change must be viewed as not only a computational mechanism for state-space exploration, but as a method for progressively varying controller complexity (through dimensionality) in order to achieve the goals of motor learning. Secondly, it has been suggested by some form of internal models [14], [7] are towards optimal adaptation and control of behaviour. Lastly, joint stiffness regulation can also be viewed as a method for suitably modify dynamic behaviour during learning, i.e. thereby control accuracy of internal models during intermediate stages. Here, dimensionality refers to the size of the state-space of internal models. Dimensional change is denoted as a progressive variation in internal model dimensionality. We use the method of nonlinear system balancing [4] to obtain a progression of forward models of increasing dimensionality. The accuracy of such models in intermediate stages of learning then becomes a measure of effectiveness of the dimensional change strategy for control. We use this accuracy measure to propose an optimal stiffness development strategy for regulating learnability.

Our model is tested on a simulated 2-link kinematic chain with muscle-like actuation. Through comparison of state trajectories of a reduced dimensional forward model and the full order system, we study effect of joint stiffness development on two kinds of tasks : point-to-point reaching and via-point reaching. We then compare two kinds of joint stiffness development strategies, (i) uniform development, and (ii) proximo-distal development. Interestingly our results indicate that while uniform stiffness development is better for overall accuracy, proximo-distal stiffness development could be easier to regulate due to the uniformity of the forward model accuracy landscape.

This paper is organised as follows. In Sec. II we introduce our mathematical framework in modelling dimensional change using control dimensionality reduction. Sec. III presents the simulation model we employ and details the experiments performed. The results are then analysed in Sec. IV, followed by the discussion in Sec. V.

II. FRAMEWORK AND BACKGROUND

Consider the following representation of neuromuscular dynamics,

$$\dot{x} = f(x) + g(x, u), \quad y = h(x), \quad (1)$$

in which, $x \in \mathbb{R}^S$ is the denoted the state, $u \in \mathbb{R}^I$ is denoted the input, $y \in \mathbb{R}^O$ denoted as the output. In this representation, $f(\cdot)$ represents the state transition dynamics (or ‘natural’ dynamics) and $g(\cdot)$ the input dynamics. The output y relates to the full order state x through the output transformation $h(\cdot)$. Let us denote this system as \mathcal{F} . The input reflects the activity of motor neurons, which for vertebrates can number in the thousands, but may be correlated. Here, I reflects the number of independent components. In the representation in Eq. (1), the joint stiffness affects the mechanical component of the natural dynamics.

Organisms control task trajectories by manipulating the neural control input $u(t)$. The key question in both the biological and robotic context is how task control problem is solved given redundancy and high dimensionality. The problem can be re-cast using optimal control theory (OCT) by assuming

that different trajectories have different *fitness* for the organism; where this fitness is quantified by some form of a cost function J . The trajectory $y^*(t)$ that minimises cost J^* , and its corresponding control $u^*(t)$ are optimal and can be found by various means.

The OCT models of human behaviour have suggested several possibilities for the cost function J with energy, time, accuracy, control input norm, joint torque changes [5] being popular candidates depending upon the motor behaviour being tested. Regardless of the actual cost function, the investigations do not address the neural acquisition of this optimal control $u^*(t)$. Due to phenotypic variations, it can be expected that individual organisms must somehow acquire this optimal solution through some process of reinforcement learning [6] in order to achieve the task goals. Key to this learning process is the suitable exploration of the state-space and therefore large state-spaces can severely affect the tractability of the iterative optimal control problem. In biological systems, the exploration of state-action solutions is carried out using the so-called internal models [14], [7], comprising of forward and inverse models.

Within this control-theoretic state-space framework, we interpret the notion of dimensional change simply as a progressively increase in dimensionality of the forward model. Now, the full-dimensional forward model is simply a formulation of the dynamics in the form of Eq. (1), since it relates the input signals $u(t)$ to the state $x(t)$ and output trajectories $y(t)$. This is task-specific since it is dependent on specification of the task variables y .

A reduced dimensional state $x_r \in \mathbb{R}^{S_r}$ is then taken to be the map $x_r = Px$ defined by the *projection* operator $P \in S_r \times S$. We use the method of nonlinear system balancing [4] in order to compute the projection matrix P . This process is based on collecting datasets of system behaviour and then computing controllability and observability Gramians – measures for quantifying the importance of each of the state variables to the control problem. A dimension reduction of this forward model can

then be obtained by the Galerkin projection,

$$\begin{aligned} \dot{x}_r &= Pf(P^T x_r) + Pg(P^T x_r, u), \\ y &= h(P^T x_r), \end{aligned} \quad (2)$$

we can then denote this system as \mathcal{F}_r . In order to investigate a dimensional change, we assume that at intermediate stages of learning, the dimensionality of forward model state-space x_r progressively *increases* in the range $S_r = [1 \dots S]$. This therefore yields a family of forward models of increasing complexity which can then be represented by $\mathcal{F}_r \subset [\mathcal{F}_1, \mathcal{F}_2, \dots \mathcal{F}]$.

Now, since the reduced dimensional forward models only correspond to a projection of the dynamics, an iterative optimal control computed on such a system only yields a suboptimal solution, and control. However through progressive increase of state-space complexity, the optimal solutions may be found in a tractable manner as the dimensionality is increased to a sufficiently large degree [9]. Nevertheless, it is important to ensure that solutions computed at intermediate stages of learning minimise the error. This entails maximising the accuracy of the various intermediate forward models \mathcal{F}_r .

We examine the hypothesis that through careful and appropriate regulation of the joint stiffness K , the accuracy of the intermediate stages of learning can be maximised. This approach then yields a dimensional change strategy that ensures optimality of the motor behaviours along with a speed up of the learnability. We examine this hypothesis by quantifying the forward model error by,

$$\epsilon_r = \|\mathcal{F} - \mathcal{F}_r\| = \|y(t) - y_r(t)\|_2, \quad (3)$$

where we assume $\|\cdot\|$ is some error operator while $\|\cdot\|_2$ is the standard 2 norm on the sequences $y(t)$ and $y_r(t)$ that are obtained in response to some (task-specific) control input $u(t)$. Consequently, at intermediate stages of learning, the optimal stiffness development K is chosen in order to minimise,

$$K^* = \underset{K}{\operatorname{argmin}} \quad \epsilon_r, \quad (4)$$

where $K_{min} \leq K^* < K_{max}$. Since the forward model accuracy now depends on joint stiffness as well as dimensionality S , we analyse the accuracy

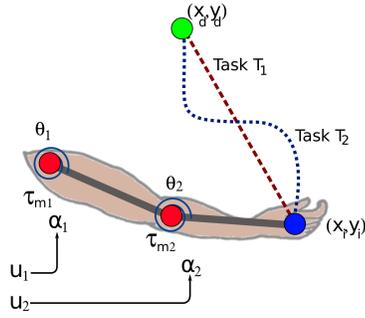


Fig. 1. Simulation setup : A two-link kinematic chain actuated by muscle-like actuators generating torque from a first-order response to activation. The output is the Cartesian position at the end of the chain. Tasks modelled are straight line point-to-point, and viapoint reaching, with zero terminal velocities.

as a $2D$ surface that we denote as the forward model *accuracy landscape*. Finding an optimal stiffness development strategy therefore is reduced to a question of finding an optimal path in this $2D$ surface as seen in the results in Sec. IV.

The simulation setup and experiments performed are introduced next.

III. EXPERIMENTS AND RESULTS

The dimension change problem was studied using a simulated 2 link planar kinematic chain model of a human limb with passive joint stiffness as represented by,

$$M(\Theta)\ddot{\Theta} + N(\Theta, \dot{\Theta})\dot{\Theta} + K_j(\Theta - \Theta_0) = \mu_\alpha \alpha, \quad \dot{\alpha} = \left(\frac{1}{\tau_s}\right)(u - \alpha) \quad (5)$$

where $\Theta = [\theta_1, \theta_2]^T$ is a vector of joint angles, $M(\Theta)$ is the mass-inertia matrix, $N(\Theta, \dot{\Theta})$ the Coriolis and viscous damping matrix, K_j the passive joint stiffness, μ_α is a muscle force scaling, $\alpha = [\alpha_1, \alpha_2]^T$ is a vector of muscle activations, Θ_0 is a vector of rest angles, and τ_s is the muscle activation time constant.

Each joint is supplied with torques from individual muscle-like actuators with linear first order dynamics driven by the input signals $u = [u_1, u_2]^T$. The parameters of the model are chosen to loosely

mimic an adult human arm [15] with the following parameters. The link masses are chosen as $m_1 = 0.75m$, $m_2 = 0.5m$, link lengths as $l_1 = l_2 = 0.4m$. The joints are affected by viscous friction with a coefficient of $0.15N/m/sec$ and the muscles use a time constant of $\tau_s = 0.01sec$. The muscle torques are scaled by a magnitude of 1.6 and 0.4 at the first and second joint respectively. The passive joint compliance has a rest length of $\Phi_0 = [\pi/8, \pi/4]^T$ respectively, and the arm is initialised to these joint angles to start at rest.

The arm model and the experimental scenario can be seen in Fig. 1. Two kinds of tasks are studied, that of point-to-point reaching (T_1) and via-point reaching (T_2). In both cases, the final objective is to reach the point $P = [0.4, 0.2]^T$ in space within a duration of $2.5secs$. The experimental methodology is as follows :

- 1) First using the full dimensional plant model, an inverse controller of full order is derived in the form of $u_d(t)$. This is used to benchmark the forward model accuracy.
- 2) Next using the approach of system balancing [4] reduced dimensional (forward) models of the plant are derived. Since the full order plant is of order $R = 6$, this entails a set of 6 models of dimensionality $S \subset [1 \dots 6]$. The obtained forward models belong to the set $\mathcal{F}_r \subset [\mathcal{F}_1, \mathcal{F}_2, \dots \mathcal{F}_6]$
- 3) The forward model accuracy ϵ_r is estimated in each case from the response to the controls signal $u_d(t)$ from the normed difference in the outputs as described by Eq. (3).
- 4) The forward model accuracy landscape is then calculated as a $2D$ interpolate in order to uniformly represent the region under analysis.
- 5) This process is iterated for the joint stiffness $K_{min} \leq K^* < K_{max}$.

The dimensional change problem was then studied for two kinds of stiffness change strategies : (i) Uniform stiffness change : the stiffness of both the joints were uniformly varied within the range; (ii) Proximo-distal stiffness change : the stiffness of the proximal joint was fixed to a nominal (final) value and the distal joint stiffness was varied within the full range ($K_{min} = 0.1$, $K_{max} = 0.6$).

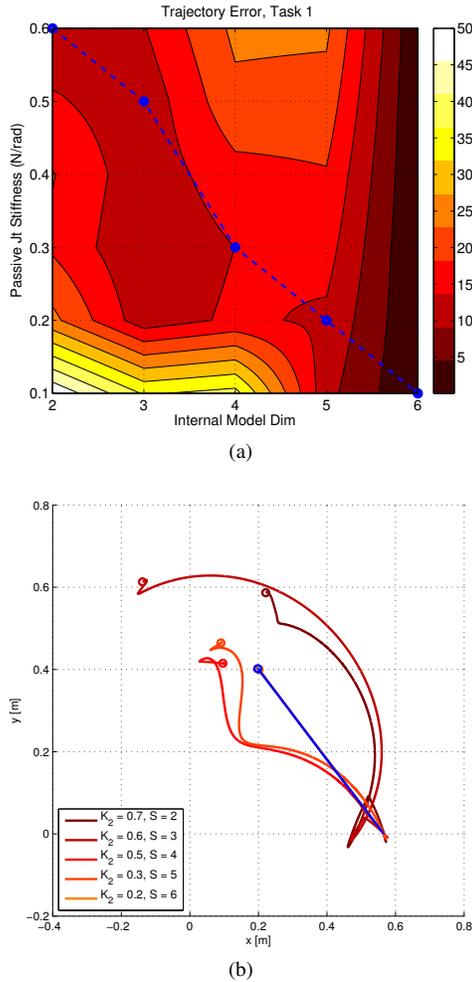


Fig. 2. Joint stiffness variation and dimensional change in point-to-point reaching : Joint stiffness was varied on the distal joint, and fixed on proximal joint to $0.3N/rad$ (Proximo-distal development); The resulting (a) Forward model accuracy landscape, and the optimal joint stiffness development (blue line); (b) Cartesian trajectories of the optimal stiffness candidates at individual stages of dimensional change (i.e. corresponding to the blue circles in (a)).

IV. RESULTS

The proximo-distal experiment results are presented in Fig. 2. The forward model accuracy landscape Fig. 2a as explained earlier plots the difference between the full dimensional system response and the respective reduced dimensional forward models. The errors are large to the left

of this area, since a 2 dimensional system cannot accurately capture the dynamics. To the right edge, the errors drop to 0 magnitude since the source plant is of dimensionality 6, i.e. the model is at full accuracy. From the distribution of errors in this landscape, the optimal stiffness development is indicated by the blue line. This result is qualitative identical to the some experimental observations [13], [2]. The Cartesian trajectories of the optimal development strategy at various stages of learning (corresponding to dimensional increase) and the full dimensional plant can be seen in Fig. 2b. Clearly, as dimensionality increases, the model performance improves until the 6th dimension which is identical to the full order system in its response.

The task independent nature of the optimal solution can be seen in the results from the second experiment depicted in Fig. 3. In this case, even though the task is that of via-point reaching, the obtained accuracy landscape in Fig. 3a is qualitatively similar to the point-to-point case. This is because even though the tasks are different, the chosen task output variables are identical, i.e. that of the endpoint position (x, y) . It must however be noted that the magnitude of errors is greater than the previous case (especially towards the left half of the landscape), as seen in Fig. 3b, due to the dynamic nature of the trajectory. Clearly a great number of dimensions are necessary in order to suitably capture the dynamics.

In the next experiment, we compare the the alternative joint-stiffness variation strategy of uniform development. In this case, the stiffness of both joints are simultaneously varied and the forward model accuracy landscape is presented in Fig. 4a. In this case, the landscape is clearly different from the results of the proximo-distal joint stiffness development strategy in Fig. 2a. For the stiffness range that was evaluated, a different optimal joint stiffness development profile is obtained as a result as seen in the blue line. Also, in this scenario, the peak magnitude of the accuracy errors is less than half of the proximo-distal strategy, implying that this strategy is preferable if only accuracy in reaching at intermediate stages is desired. Nevertheless, the landscape is clearly not as uniform as the proximo-

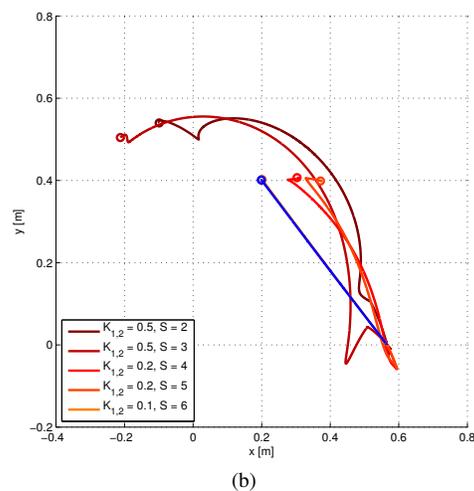
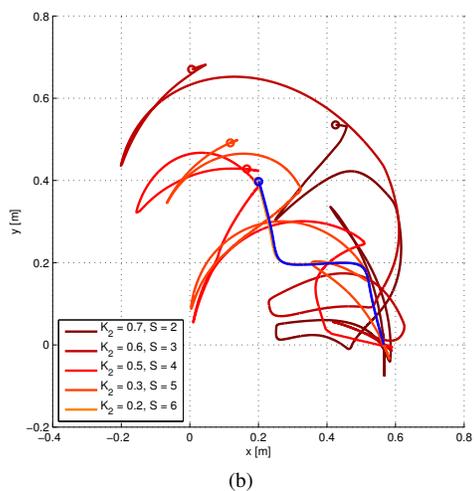
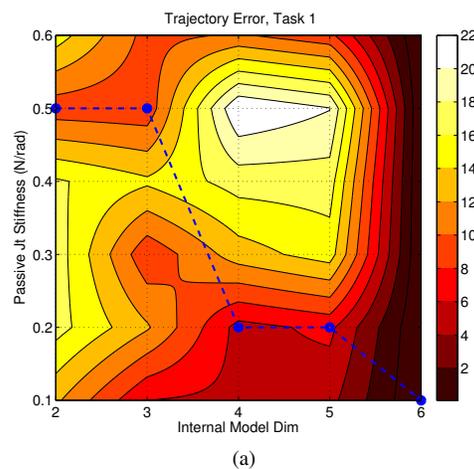
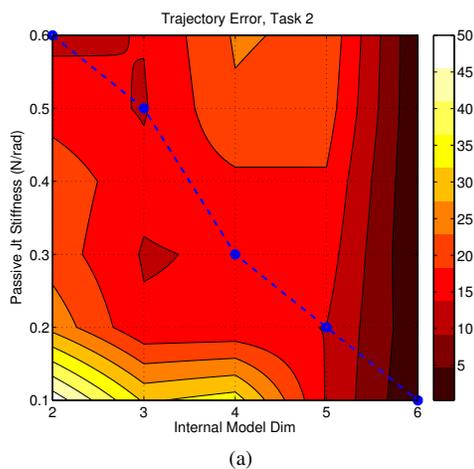


Fig. 3. Joint stiffness variation and dimensional change in via-point reaching : Joint stiffness was varied on the distal joint, and fixed on proximal joint (Proximo-distal development); The resulting (a) Forward model accuracy landscape, and the optimal joint stiffness development (blue line) which is identical to point-to-point reaching; (b) Cartesian trajectories of the optimal stiffness candidates at various stages of dimensional change (i.e. corresponding to the blue circles in (a)).

Fig. 4. Joint stiffness variation and dimensional change on point-to-point reaching : Joint stiffness varied uniformly across both joints (Uniform development); (a) Forward model accuracy landscape has lower magnitudes than proximo-distal development, but the error surface is non-uniform. (b) Cartesian trajectories of the optimal stiffness candidates at various stages of dimensional change (i.e. corresponding to the blue circles in (a)).

distal case. It can be expected that with parametric variation corresponding to individual phenotypes, this strategy may not be that conducive to a smooth well-defined joint-stiffness development. We believe that this gives additional support to the predominating observations of proximo-distal strategies in both children and in adults in learning

new skills [13], [2]. Nevertheless, it can be expected that the proximo-distal development approach need not be the best on a task-independent basis and further analysis on various tasks might be needed.

V. DISCUSSION

In this paper, we presented an approach for quantifying the nature of dimensional change occurring during development in an organism from the control-theoretic viewpoint of dimensionality reduction. We then analysed the role played by joint stiffness towards dimensional change as a mechanism for regulating learnability. Using the approach of nonlinear system balancing, we modelled dimensional change as a progressive variation in dimensionality of a forward model of the sensorimotor behaviour. Using the accuracy of the model at intermediate stages of learning as a benchmark, we compared two strategies for joint stiffness development, that of uniform and proximo-distal development. Our results indicate that on one hand, uniform joint stiffness development is more beneficial from the perspective of maximising accuracy while regulating learnability, whereas, proximo-distal development might be easier due to the regularity of the accuracy landscape.

Developmental plasticity is concerned with the irreversible changes that the newborn phenotype undergoes as it matures and develops, which typically takes a considerable fraction of the phenotype's lifetime. In either case, it seems likely that there should be a premium on learning as quickly as possible for a given level of competence and task complexity. It can be expected that any strategy that has evolved for speeding up of learning exploits the nature of the embodiment through the morphological properties. The joint stiffness development mechanism we present in this paper is an effective method for controlling the learnability.

The results we present have two key implications for artificial systems. Firstly, model-based reinforcement or adaptive learning strategies could be facilitated in their tractability for complex robot morphologies. Secondly, our method presents a strategy for harnessing variable compliance at the joints to expedite learning. Such a strategy can then be used in conjunction with existing methods focussing on utilising muscle-like actuation for energy harvesting and control robustness. This might pave the way towards developmental acquisition of motor abilities in complex anthropomorphic robots.

ACKNOWLEDGEMENT

This paper was supported by the FP7 EU projects CoDyCo (No. 600716 ICT 2011.2.1 Cognitive Systems and Robotics), and Koroibot (No. 611909 ICT-2013.2.1 Cognitive Systems and Robotics).

REFERENCES

- [1] N. Bernstein. *The Co-ordination and Regulation of Movements*. Oxford, UK: Pergamo, 1967.
- [2] N. E. Berthier, R. K. Clifton, D. D. McCall, and D. J. Robin. Proximodistal structure of early reaching in human infants. *Experimental Brain Research*, 127(3):259–269, 1999.
- [3] L. Berthouze and M. Lungarella. Motor skill acquisition under environmental perturbations: On the necessity of alternate freezing and freeing of degrees of freedom. *Adaptive Behavior*, 12(1):47–64, 2004.
- [4] J. Hahn and T. F. Edgar. An improved method for nonlinear model reduction using balancing of empirical gramians. *Computers & chemical engineering*, 26(10):1379–1397, 2002.
- [5] C. M. Harris. On the optimal control of behaviour: a stochastic perspective. *Journal of Neuroscience Methods*, 83(1):73–88, 1998.
- [6] C. M. Harris. *The Oxford Handbook of Eye Movements*, chapter Oculomotor Development Pathology : An 'Evo-Devo' Perspective, pages 663–686. Number 37. Oxford University Press, 2011.
- [7] Kawato. Internal models for motor control and trajectory planning. *Curr Opin Neurobiol*, 9(6):718–27, Dec. 1999.
- [8] N. Kuppawamy and C. M. Harris. Developing learnability : The case for reduced dimensionality. In *(Development and Learning and Epigenetic Robotics (ICDL), 2013 IEEE International Conference on*, pages 1–2. IEEE, 2013.
- [9] N. Kuppawamy and C. M. Harris. Do muscle synergies reduce the dimensionality of behaviour? *Frontiers in Computational Neuroscience*, 8(63), 2014.
- [10] S. Mitra, P. G. Amazeen, and M. Turvey. Intermediate motor learning as decreasing active (dynamical) degrees of freedom. *Human Movement Science*, 17(1):17–65, 1998.
- [11] K. M. Newell and D. E. Vaillancourt. Dimensional change in motor learning. *Human Movement Science*, 20(4):695–715, 2001.
- [12] O. Sporns and G. M. Edelman. Solving bernstein's problem: A proposal for the development of coordinated movement by selection. *Child development*, 64(4):960–981, 1993.
- [13] B. Vereijken, R. E. v. Emmerik, H. Whiting, and K. M. Newell. Free (z) ing degrees of freedom in skill acquisition. *Journal of motor behavior*, 24(1):133–142, 1992.
- [14] D. M. Wolpert, R. C. Miall, and M. Kawato. Internal models in the cerebellum. *Trends in cognitive sciences*, 2(9):338–347, 1998.
- [15] F. E. Zajac. Muscle and tendon: properties, models, scaling, and application to biomechanics and motor control. *Critical reviews in biomedical engineering*, 17(4):359–411, 1988.