

Synthesis and Adaptation of Effective Motor Synergies for the Solution of Reaching Tasks

Cristiano Alessandro¹, Juan Pablo Carbajal¹, and Andrea d'Avella²

¹ Artificial Intelligence Laboratory, University of Zürich, Switzerland.
{alessandro, carbajal}@ifi.uzh.ch

² Laboratory of Neuromotor Physiology, Fondazione Santa Lucia, Italy.
a.davella@hsantalucia.it

Abstract. Taking inspiration from the hypothesis of muscle synergies, we propose a method to generate open loop controllers for an agent solving point-to-point reaching tasks. The controller output is defined as a linear combination of a small set of predefined actuations, termed synergies. The method can be interpreted from a developmental perspective, since it allows the agent to autonomously synthesize and adapt an effective set of synergies to new behavioral needs. This scheme greatly reduces the dimensionality of the control problem, while keeping a good performance level. The framework is evaluated in a planar kinematic chain, and the quality of the solutions is quantified in several scenarios.

Keywords: motor primitives, motor control, development

1 Introduction

Humans are able to perform a wide variety of tasks with great flexibility; learning new motions is relatively easy, and adapting to new situations (e.g. change in the environment or body growth) is usually dealt with no particular effort. The strategies adopted by the central nervous system (CNS) to master the complexity of the musculoskeletal apparatus and provide such performance are still not clear. However, it has been speculated that an underlying modular organization of the CNS may simplify the control and provide the observed adaptability. There is evidence that the muscle activity necessary to perform various tasks (e.g. running, walking, keeping balance, reaching and other combined movements) may emerge from the combination of predefined muscle patterns, the so-called *muscle synergies* [1]. This organization seems to explain muscle activity across a wide range of combined movements [2–4].

The scheme of muscle synergies is inherently flexible and adaptable. Different actions are encoded by specific combinations of a small number of predefined synergies; this reduces the computational effort and the time required to learn new useful behaviors. The learning scheme can be regarded as developmental since information previously acquired (i.e. synergies) can be reused to generate new behaviors[5]. Finally, improved performance can be easily achieved by introducing additional synergies. Thus, the hypothetical scheme of muscle synergies would contribute to the autonomy and the flexibility observed in biological

systems, and it could inspire new methods to endow artificial agents with such desirable features.

In this paper we propose a method to control a dynamical system (i.e. the agent) in point-to-point reaching tasks by linear combinations of a small set of predefined actuations (i.e. synergies). Our method initially solves the task in state variables by interpolation; then, it identifies the combination of synergies (i.e. actuation) that generate the closest kinematic trajectory to the computed interpolant. Additionally, we propose a strategy to synthesize a small set of synergies that is tailored to the task and the agent. The overall method can be interpreted in a developmental fashion; i.e. it allows the agent to autonomously synthesize and update its own synergies to increase the performance of new reaching tasks.

Other researchers in robotics and control engineering have recently proposed architectures inspired by the concept of muscle synergies. In [6] the authors derive an analytical form of a set of primitives that can drive a feedback linearized system (known analytically) to any point of its configuration space. In [7] the authors present a numerical method to identify synergies that optimally drive the system over a set of desired trajectories. This method does not require an analytical description of the system, and it has the advantage of assessing the quality of the synergies in task space. However, it is computationally expensive as it involves heavy optimizations. In [8] muscle synergies are identified by applying an unsupervised learning procedure to a collection of sensory-motor data obtained by actuating a robot with random signals. In [9] the architecture of the dynamic movement primitives (DMP) is proposed as a novel tool to formalize control policies in terms of predefined differential equations. Linear combinations of Gaussian functions are used as inputs to modify the attractor landscapes of these equations, and to obtain the desired control policy.

In contrast to these works, our method to synthesize synergies does not rely on feedback linearization, nor on repeated integrations of the dynamical system. The method is grounded on the input-output relation of the dynamical system (as in [8]), and it provides a computationally fast method to obtain the synergy combinatorics to solve a given task. Furthermore, our method is inherently adaptable as it allows the on-line modification of the set of synergies to accommodate to new reaching tasks.

2 Definitions and Methods

In this section we introduce the mathematical details of the method we propose. After some definitions, we present the core element of our method: a general procedure to compute actuations that solve point-to-point reaching tasks (see Sec. 2.1). Subsequently, in Section 2.2, we propose a framework for the synthesis and the development of a set of synergies.

Let us consider a differential equation modeling a physical system $\mathcal{D}(\mathbf{q}(t)) = \mathbf{u}(t)$, where $\mathbf{q}(t)$ represents the time-evolution of its configuration variables (their derivatives with respect to time are $\dot{\mathbf{q}}(t)$), and $\mathbf{u}(t)$ is the actu-

ation applied. Inspired by the hypothesis of muscle synergies³ [1], we formulate the actuation as a linear combination of predefined motor co-activation patterns:

$$\mathbf{u}(t) = \sum_{i=1}^{N_\phi} \phi_i(t) \mathbf{b}_i := \mathbf{\Phi}(t) \mathbf{b}, \quad (1)$$

where the functions $\phi_i(t) \in \mathbf{\Phi}$ are called *motor synergies*. The notation $\mathbf{\Phi}(t)$ describes a formal matrix where each column is a different synergy. If we consider a time discretization, $\mathbf{\Phi}(t)$ becomes a $N \dim(\mathbf{q})$ -by- N_ϕ matrix, where N is the number of time steps, $\dim(\mathbf{q})$ the dimension of the configuration space and N_ϕ the number of synergies.

We define *dynamic responses* (DR) of the set of synergies as the responses $\theta_i(t) \in \mathbf{\Theta}$ of the system to each synergy (i.e. forward dynamics):

$$\mathcal{D}(\theta_i(t)) = \phi_i(t) \quad i = 1 \dots N_\phi. \quad (2)$$

with initial conditions chosen arbitrarily.

2.1 Solution to point-to-point reaching tasks

A general point-to-point reaching task consists in reaching a final state $(\mathbf{q}_T, \dot{\mathbf{q}}_T)$ from an initial state $(\mathbf{q}_0, \dot{\mathbf{q}}_0)$ in a given amount of time T :

$$\begin{aligned} \mathbf{q}(0) &= \mathbf{q}_0, & \dot{\mathbf{q}}(0) &= \dot{\mathbf{q}}_0, \\ \mathbf{q}(T) &= \mathbf{q}_T, & \dot{\mathbf{q}}(T) &= \dot{\mathbf{q}}_T. \end{aligned} \quad (3)$$

Controlling a system to perform such tasks amounts to finding the actuation $\mathbf{u}(t)$ that fulfills the point constraints⁴ (4). Specifically, assuming that the synergies are known, the goal is to identify the appropriate synergy combinator \mathbf{b} . In this paper we consider only the subclass of reaching tasks that impose motionless initial and final postures, i.e. $\dot{\mathbf{q}}_T = \dot{\mathbf{q}}_0 = 0$.

The procedure consists of, first, solving the problem in kinematic space (i.e. finding the appropriate $\mathbf{q}(t)$), and then computing the corresponding actuations. From the kinematic point of view, the task can be seen as an interpolation problem; i.e. $\mathbf{q}(t)$ is a function that interpolates the data in (4). Therefore, a set of functions is used to build the interpolant trajectory that satisfy the constraints imposed by the task; these functions are herein the dynamic responses of the synergies:

$$\mathbf{q}(t) = \sum_{i=1}^{N_\theta} \theta_i(t) \mathbf{a}_i := \mathbf{\Theta}(t) \mathbf{a}, \quad (4)$$

where the vector of combinator \mathbf{a} is chosen such that the task is solved. As mentioned earlier, if time is discretized, $\mathbf{\Theta}(t)$ becomes a $N \dim(\mathbf{q})$ -by- N_θ matrix, where N_θ is the number of dynamic responses. The quality of the DR as interpolants is evaluated in sections 3.

³ With respect to the model of time-varying synergies, in this paper we neglect the synergy onset times.

⁴ In this paper we assume that the initial conditions of the systems are equal to $(\mathbf{q}_0, \dot{\mathbf{q}}_0)$

Once a kinematic solution has been found (as linear combination of DRs), the corresponding actuation can be obtained by applying the differential operator; i.e. $\mathcal{D}(\Theta(t)\mathbf{a}) = \tilde{\mathbf{u}}(t)$. Finally, the vector \mathbf{b} can be computed by projecting $\tilde{\mathbf{u}}(t)$ onto the synergy set Φ . If $\tilde{\mathbf{u}}(t)$ does not belong to the linear span of Φ , the solution can only be approximated in terms of a defined norm (e.g. Euclidean):

$$\mathbf{b} = \arg \min_{\mathbf{b}} \|\tilde{\mathbf{u}}(t) - \Phi(t)\mathbf{b}\|. \quad (5)$$

When the time is discretized, all functions of time becomes vectors and this equation can be solved explicitly using the psuedoinverse of the matrix Φ ,

$$\Phi^+ \tilde{\mathbf{u}} = \Phi^+ \mathcal{D}(\Theta\mathbf{a}) = \mathbf{b}. \quad (6)$$

This equation highlights the operator $\Phi^+ \circ \mathcal{D} \circ \Theta$ (\circ denotes operator composition) as the mapping between the kinematic combinators \mathbf{a} (kinematic solution) and the synergy combinators \mathbf{b} (dynamic solution). Generically, this operator represents a nonlinear mapping $\mathcal{M} : \mathbb{R}^{N_\theta} \rightarrow \mathbb{R}^{N_\phi}$, and it will be discussed in Section 4.

To assess the quality of the solution we define the following measures:

Interpolation error: Measures the quality of the interpolant $\Theta(t)\mathbf{a}$ with respect to the task. Strictly speaking, only the case of negligible errors corresponds to interpolation. A non-zero error indicates that the trajectory $\Theta(t)\mathbf{a}$ only approximates the task

$$\text{err}_I = \sqrt{\|\mathbf{q}_T - \Theta(T)\mathbf{a}\|^2 + \|\dot{\Theta}(T)\mathbf{a}\|^2}, \quad (7)$$

where $\|\cdot\|$ denotes the Euclidean norm, and the difference between angles are mapped to the interval $(-\pi, \pi]$.

Projection error: Measures the distance between the actuation that solves the task $\tilde{\mathbf{u}}(t)$, and the linear span of the synergy set Φ

$$\text{err}_P = \sqrt{\int_0^T \|\tilde{\mathbf{u}}(t) - \Phi(t)\mathbf{b}\|^2 dt}. \quad (8)$$

Forward dynamics error: Measures the error of a trajectory $\tilde{\mathbf{q}}(t, \lambda)$ generated by an actuation $\Phi(t)\lambda$, in relation to the task.

$$\text{err}_F = \sqrt{\|\tilde{\mathbf{q}}(T, \lambda) - \mathbf{q}_T\|^2 + \|\dot{\tilde{\mathbf{q}}}(T, \lambda) - \dot{\mathbf{q}}_T\|^2}. \quad (9)$$

Replacing $\tilde{\mathbf{q}}(t, \lambda)$, \mathbf{q}_T and $\dot{\mathbf{q}}_T$ with their corresponding end-effector values provides the forward dynamics error of the end-effector.

2.2 Synthesis and Development of Synergies

The synthesis of synergies is carried on in two phases: exploration and reduction. The exploration phase consists in actuating the system with an extensive set of motor signals Φ_0 in order to obtain the corresponding DRs Θ_0 . The reduction

phase consists in solving a small number of point-to-point reaching tasks in kinematic space (that we call *proto-tasks*) by creating the interpolants using the elements of set Θ_0 , as described in Eq. (5). These solutions are then taken as the elements of the reduced set Θ . Finally, the synergy set Φ is computed using relation (3), i.e. inverse dynamics. As a result, there will be as many synergies as the number of the proto-tasks (i.e. $N_\phi = N_\theta$). The intuition behind this reduction is that the synergies that solve the proto-tasks may capture essential features both of the task and of the dynamics of the system. Despite the non-linearities of \mathcal{D} , linear combination of these synergies might be useful to solve point-to-point reaching tasks that are similar (in terms of Eq. (4)) to the proto-tasks (see Sec. 3).

The number of proto-tasks as well as their specific instances determine the quality of the synergy-based controller. To obtain good performance in a wide variety of point-to-point reaching tasks, the proto-tasks should cover relevant regions of the state space (see Sec. 3). Clearly, the higher the number of different proto-tasks, the more regions that can be reached with good performance. However, a large number of proto-tasks (and the corresponding synergies) increases the dimensionality of the controller. In order to tackle this trade-off, we propose a procedure that parsimoniously adds a new proto-task only when and where it is needed: if the performance in a new reaching task is not satisfactory, we add a new proto-task in one of the regions with highest projection error or we modify existing ones.

3 Results

We apply the methodology described in Section 2 to a simulated planar kinematic chain (see [10] for model details) modeling a human arm[11]. In the exploration phase, we employ an extensive set of motor signals Φ_0 to actuate the arm model and generate the corresponding dynamic responses Θ_0 . The panels in the first row of Fig. 1 show the end-effector trajectories resulting from the exploration phase. We test two different classes of motor signals: actuations that generate minimum jerk end-effector trajectories (100 signals), and low-passed uniformly random signals (90 signals). In order to evaluate the validity of the general method described in Sec. 2.1, we use the sets Φ_0 and Θ_0 to solve 13 different reaching tasks without performing the reduction phase. The second row of Fig. 1 depicts the trajectories drawn by the end-effector when the computed mixture of synergies are applied as actuations (i.e. forward dynamics of the solution). It has to be noted how the nature of the solutions (as well as that of the responses), depends on the class of actuations used. The maximum errors are reported in Table 1. The results are highly satisfactory for both the classes of actuations, and show the validity of the method proposed. Since the reduction phase has not been performed, the dimension of the combinator vectors \mathbf{a} and \mathbf{b} equals the number of actuations used in the exploration.

The objective of the reduction phase is to generate a small set of synergies and DRs that can solve desired reaching tasks effectively. As described in Section 2.2, this is done by solving a handful of proto-tasks. The number (and the instances) of these proto-tasks determines the quality of the controller. Figure

| | Min. | Jerk | Random |
|------------------|-------------------|------|-------------------|
| err _I | 10 ⁻¹⁵ | | 10 ⁻¹⁵ |
| err _P | 10 ⁻⁵ | | 10 ⁻³ |
| err _F | 10 ⁻⁴ | | 10 ⁻³ |

Table 1. Order of the maximum errors obtained by using Φ_0 and Θ_0 (no reduction phase).

2 shows the projection error as a function of the number of proto-tasks. The reduction is applied to the low-passed random signal set. Initially, two targets are chosen randomly (top left panel); subsequent targets are then added on the regions characterized by higher projection error. As it can be seen, the introduction of new proto-tasks leads to better performance on wider regions of the end-effector space, and eventually the whole space can be reached with reasonable errors. In fact, the figure shows that this procedure decreases the average projection error to 10^{-3} (comparable to the performance of the whole set Φ_0 , see Tab. 1) and reduces the dimension of the combinator vector to 6, a fifteen-fold reduction. This result shows that a set of “good” synergies can drastically reduce the dimensionality of the controller, while maintaining similar performance. The bottom right panel of the figure shows the forward dynamics error of the end-effector obtained with the 6 proto-tasks. Comparing this panel with the bottom left one, it can be seen that the forward dynamics error of the end-effector reproduces the distribution of the projection error, rendering the latter a good estimate for task performance.

To further demonstrate that the reduction phase we propose is not trivial, we compare the errors resulting from the set of 6 synthesized synergies, with the errors corresponding to 100 random subsets of size 6 drawn from the set of low-passed random motor signals. Figure 3 shows this comparison. The task consists in reaching the 13 targets in Fig. 1. The boxplots correspond to the errors of the random subsets, and the filled circles to the errors of the synergies resulting from the reduction phase. Observe that, the order of the error of the reduced set is, in the worst case, equal to error of the best random subset. However, the mean error of the reduced set is about 2 orders of magnitude lower. Therefore, the reduction by proto-tasks can produce a parsimonious set of synergies out of a extensive set of actuations. Evaluating the performance with different classes of proto-tasks (e.g. catching, hitting, via-points) is postponed to future works.

4 Discussion

The results shown in the previous section justify the interpretation of the methodology as a developmental framework. Initially, the agent explores its sensory-motor system employing a variety of actuations. Later, it attempts to solve the first reaching tasks (proto-tasks), perhaps obtaining weak performance as the exploration phase may not have produced enough responses yet (see the boxplots in Fig. 3). If the agent finds an acceptable solution to a proto-task, it is used to generate a new synergy (populating the set Φ), otherwise it continues with the exploration. The failure to solve tasks of importance for its survival, could motivate the agent to include additional proto-tasks; Figure 2 illustrates

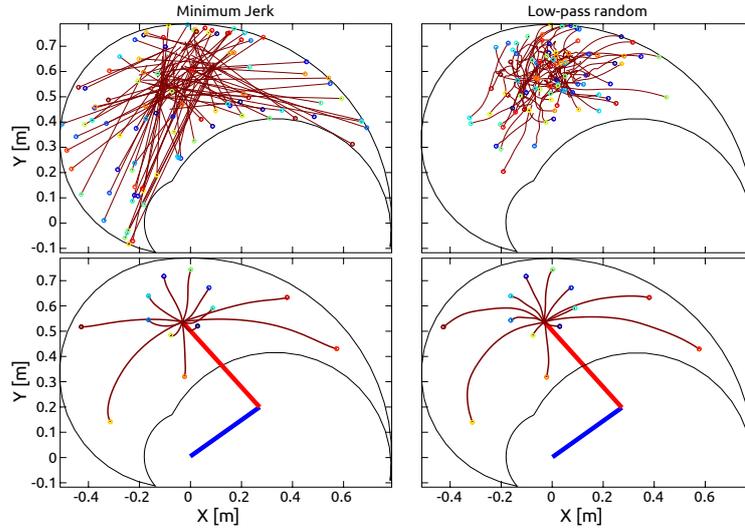


Fig. 1. Comparison of explorations with two different classes of actuation: minimum jerk and low-passed random signal. Each panel shows the kinematic chain in its initial posture (straight segments). The limits of the end-effector are shown as the boundary in solid line.

this mechanism. As it can be seen, the development of the synergy set incrementally improves the ability of the agent to perform point-to-point reaching. Alternatively, existing proto-tasks could be modified by means of a gradient descent or other learning algorithms. In a nutshell, the methodology we propose endows the agent with the ability to autonomously generate and update a set of synergies (and dynamic responses) that solve reaching tasks effectively.

Despite the difficulty of the mathematical problem (i.e nonlinear differential operator), our method seems to generate a small set of synergies that span the space of actuations required to solve reaching tasks. This is not a trivial result, since these synergies over-perform many other set of synergies randomly taken from the set Φ_0 (see Fig. 3). It appears as if the reduction phase builds features upon the exploration phase, that are necessary to solve new reaching tasks. To verify whether solving proto-tasks plays a fundamental role, our synergies could be compared with the principal components extracted from the exploration set. This verification goes beyond the scope of this paper.

An important aspect of our method is the relation between Θ and Φ (see Eq. (3)). This mapping makes explicit use of the body parameters (embedded in the differential operator \mathcal{D}), hence the synergies obtained can always be realized as actuations. The same cannot be said, in general, for synergies identified from numerical analyses of biomechanical data. Though some studies have verified the feasibility of extracted synergies as actuations [12], biomechanical constraints are not explicitly included in the extraction algorithms. Additionally, Eq. (3) provides an automatic way to cope with smooth variations of the morphology of the agent. That is, both the synergies and their dynamic responses evolve

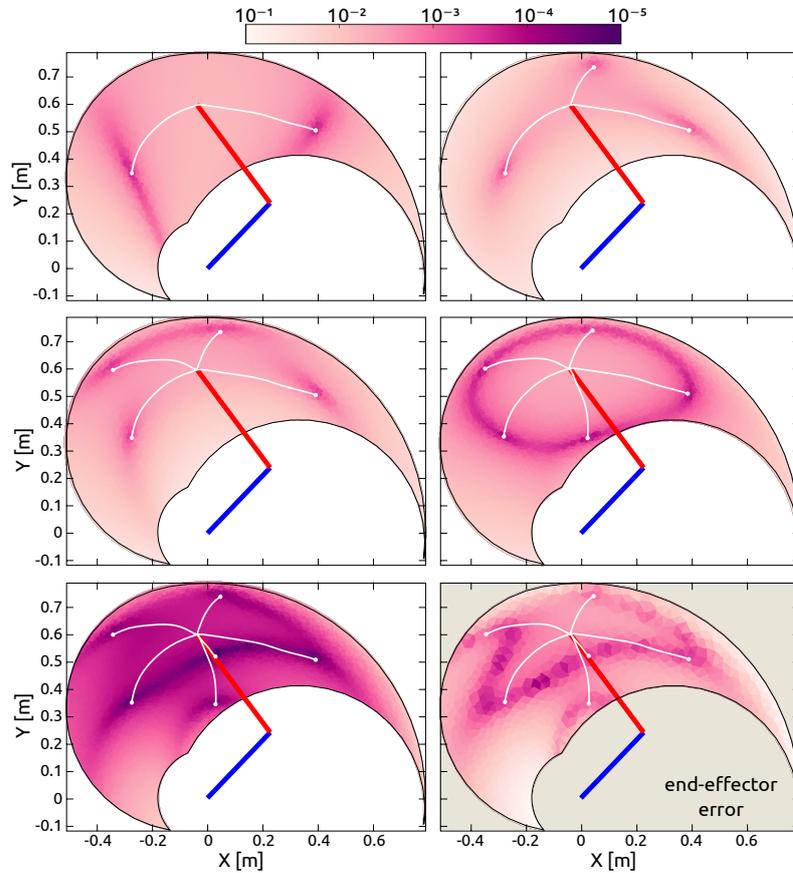


Fig. 2. Selection of targets based on projection error. Each panel shows the kinematic chain in its initial posture (straight segments). The limits of the end-effector are the boundary of the colored regions. The color of each point indicates the projection error produced to reach a target in that position. The bottom right diagram shows the forward dynamics error of the end-effector using 6 proto-tasks (6 synergies).

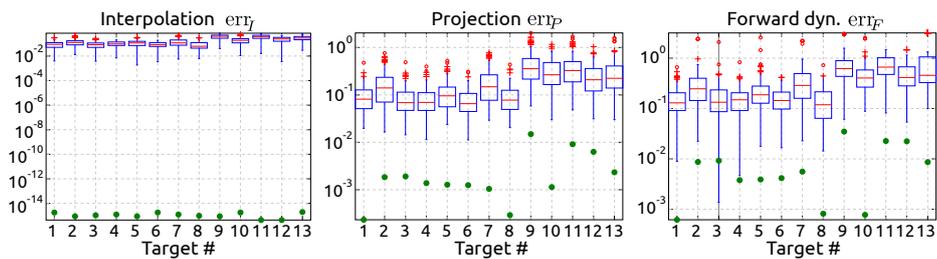


Fig. 3. Evaluation of the reduction phase. Errors produced by subsets randomly selected from the exploration-actuations (boxplots) are compared with the errors obtained after the reduction phase (filled circles).

together with the body. In line with [6, 7], these observations highlight the importance of the body in the hypothetical modularization of the CNS.

Once the task is solved in kinematic space, the corresponding actuation can be computed using the explicit inverse dynamical model of the system (i.e. the differential operator \mathcal{D}). It might appear that there is no particular advantage in projecting this solution onto the synergy set. However, the differential operator might be unknown. In this case, a synergy-based controller would allow to compute the appropriate actuation by evaluating the mapping \mathcal{M} on the vector \mathbf{a} , hence obtaining the synergy combinator \mathbf{b} . Since \mathcal{M} is a mapping between two finite low-dimensional vector spaces, estimating this map may turn to be easier than estimating the differential operator \mathcal{D} . Furthermore, we believe that the explicit use of \mathcal{D} may harm the biological plausibility of our method. In order to estimate the map \mathcal{M} , the input-output data generated during the exploration phase (i.e. Φ_0 and Θ_0) could be used as learning data-set. Further work is required to test these ideas. Additionally, preliminary theoretical considerations (not reported here) indicate that the synthesis of synergies without the explicit knowledge of \mathcal{D} is also feasible.

Finally, the current formulation of the method does not include joint limits explicitly. The interpolated trajectories are valid, i.e. they do not go beyond the limits, due to the lack of intricacy of the boundaries. In higher dimensions, especially when configuration space and end-effector are not mapped one-to-one, this may not be the case anymore. Nevertheless, joint limits can be included by reformulating the interpolation as a constrained minimization problem. Another solution might be the creation of proto-tasks with a tree-topology, relating our method to tree based path planning algorithms[13].

5 Conclusion and Future Work

The current work introduces a simple framework for the generation of open loop controllers based on synergies. The framework is applied to a planar kinematic chain to solve point-to-point reaching tasks. Synergies synthesized during the reduction phase over-perform hundreds of arbitrary choices of basic controllers taken from the exploration motor signals. Furthermore, our results confirm that the introduction of new synergies increases the performance of reaching tasks. Overall, this shows that our method is able to generate effective synergies, greatly reducing the dimensionality of the problem, while keeping a good performance level. Additionally, the methodology offers a developmental interpretation of the emergence of task-related synergies that could be validated experimentally.

Due to the nonlinear nature of the operator \mathcal{D} , the theoretical grounding of the method poses a difficult challenge, and it is the focus of our current research. Another interesting line of investigation is the validation of our method against biological data, paving the way towards a predictive model for the hypothesis of muscle synergies. Similarly, the development of an automatic estimation process for the mapping \mathcal{M} would further increase the biological plausibility of the model.

The inclusion of joint limits into the current formulation must be prioritized. Solving this problem will allow to test the method on higher dimensional redun-

dant systems. Tree-based path planning algorithms may offer a computationally effective way to approach the issue.

Funds: The research leading to these results has received funding from the European Community’s Seventh Framework Programme FP7/2007-2013-Challenge 2 - Cognitive Systems, Interaction, Robotics- under grant agreement No 248311-AMARSi, and from the EU project RobotDoC under 235065 from the 7th Framework Programme (Marie Curie Action ITN).

Authors Contribution: **CA** and **JPC** worked on the implementation of the algorithm and the generation of the results reported here. The method was born during **JPC**’s visit to **AD**’s laboratory. **AD** provided material support for this development and uncountable conceptual inputs. All three authors have contributed to the creation of the manuscript. The authors list follows an alphabetical order.

References

- [1] Andrea d’Avella, Philippe Saltiel, and Emilio Bizzi. Combinations of muscle synergies in the construction of a natural motor behavior. *Nat. Neurosci.*, 6:300–308, 2003.
- [2] Yuri P Ivanenko, Germana Cappellini, Nadia Dominici, Richard E Poppele, and Francesco Lacquaniti. Coordination of locomotion with voluntary movements in humans. *J. Neurosci.*, 25(31):7238–53, August 2005.
- [3] G Cappellini, Y P Ivanenko, R E Poppele, and F Lacquaniti. Motor patterns in human walking and running. *J. Neurophysiol.*, 95(6):3426–37, June 2006.
- [4] Andrea D’Avella, Laure Fernandez, Alessandro Portone, and Francesco Lacquaniti. Modulation of phasic and tonic muscle synergies with reaching direction and speed. *J. Neurophysiol.*, 100(3):1433–54, September 2008.
- [5] Nadia Dominici, Yuri P Ivanenko, Germana Cappellini, Andrea D’Avella, Vito Mondì, Marika Cicchese, Adele Fabiano, Tiziana Silei, Ambrogio Di Paolo, Carlo Giannini, Richard E Poppele, and Francesco Lacquaniti. Locomotor primitives in newborn babies and their development. *Science*, 334(6058):997–9, November 2011.
- [6] Francesco Nori. *Symbolic Control with Biologically Inspired Motion Primitives*. PhD thesis, University of Genova, 2005.
- [7] Cristiano Alessandro and Francesco Nori. Identification of synergies by optimization of trajectory tracking tasks. Submitted, 2012.
- [8] E Todorov and Z Ghahramani. Unsupervised learning of sensory-motor primitives. In *Proc. 25th Int. Conf. IEEE Eng. Med. & Biol Soc.*, pages 1750–1753. IEEE, 2003.
- [9] Stefan Schaal, Jan Peters, Jun Nakanishi, and Auke Ijspeert. Learning movement primitives. In Paolo Dario and Raja Chatila, editors, *Robotics Research*, volume 15 of *Tracts in Adv. Rob.*, pages 561–572. Springer, 2005.
- [10] John M. Hollerbach and Tamar Flash. Dynamic interactions between limb segments during planar arm movement. *Biol. Cybern.*, 44(1):67–77, May 1982.
- [11] Silvia Muceli, Andreas Trø llund Boye, Andrea D’Avella, and Dario Farina. Identifying representative synergy matrices for describing muscular activation patterns during multidirectional reaching in the horizontal plane. *J. Neurophysiol.*, 103(3):1532–42, March 2010.
- [12] Richard R Neptune, David J Clark, and Steven A Kautz. Modular control of human walking: a simulation study. *J. Biomech.*, 42(9):1282–7, June 2009.
- [13] Alexander Shkolnik and Russ Tedrake. Sample-based planning with volumes in configuration space. [arXiv:1109.3145v1](https://arxiv.org/abs/1109.3145v1), September 2011.