

Research



Cite this article: Al Borno M, Hicks JL, Delp SL. 2020 The effects of motor modularity on performance, learning and generalizability in upper-extremity reaching: a computational analysis. *J. R. Soc. Interface* **17**: 20200011. <http://dx.doi.org/10.1098/rsif.2020.0011>

Received: 6 January 2020

Accepted: 6 May 2020

Subject Category:

Life Sciences—Engineering interface

Subject Areas:

biomechanics, computational biology, biophysics

Keywords:

synergies, motor control, modularity, reaching, predictive simulation

Author for correspondence:

Mazen Al Borno

e-mail: malborno@stanford.edu

Electronic supplementary material is available online at <https://doi.org/10.6084/m9.figshare.c.4979909>.

The effects of motor modularity on performance, learning and generalizability in upper-extremity reaching: a computational analysis

Mazen Al Borno, Jennifer L. Hicks and Scott L. Delp

Department of Bioengineering and Mechanical Engineering, Stanford University, Stanford, CA, USA

MAB, 0000-0003-2208-9934

It has been hypothesized that the central nervous system simplifies the production of movement by limiting motor commands to a small set of modules known as muscle synergies. Recently, investigators have questioned whether a low-dimensional controller can produce the rich and flexible behaviours seen in everyday movements. To study this issue, we implemented muscle synergies in a biomechanically realistic model of the human upper extremity and performed computational experiments to determine whether synergies introduced task performance deficits, facilitated the learning of movements, and generalized to different movements. We derived sets of synergies from the muscle excitations our dynamic optimizations computed for a nominal task (reaching in a plane). Then we compared the performance and learning rates of a controller that activated all muscles independently to controllers that activated the synergies derived from the nominal reaching task. We found that a controller based on synergies had errors within 1 cm of a full-dimensional controller and achieved faster learning rates (as estimated from computational time to converge). The synergy-based controllers could also accomplish new tasks—such as reaching to targets on a higher or lower plane, and starting from alternative initial poses—with average errors similar to a full-dimensional controller.

1. Introduction

Every movement we make, from picking up a glass of water to opening a door, requires the control of a large number of muscles, which is complex due to non-linearities and redundancies in the musculoskeletal system. It has been hypothesized that the central nervous system (CNS) uses modularity to simplify the control of our high-dimensional musculoskeletal system, limiting motor commands to a small set of modules known as muscle synergies [1]. A muscle synergy is a group of muscles that have a fixed ratio of excitations. Evidence from experiments with lower vertebrates, cats, and monkeys suggests that muscle synergies are encoded in the spinal circuitry [2–4]. Previous investigators have also shown that the muscle activity measured from electromyography (EMG) can be reconstructed accurately from a linear combination of a small number of muscle synergies for a variety of tasks, including walking [5,6] and reaching [7,8]. However, other studies have presented evidence that synergies may be a by-product of task and performance constraints [9–11]. A better understanding of the advantages and limitations of muscle synergies may help us evaluate their evolutionary role across different species [12] as well as their role in motor development and learning [13].

More evidence is needed to test the hypothesis that the CNS uses muscle synergies to simplify control, particularly for movements of the upper extremity. Several investigators have argued for the need to assess muscle synergies based on task achievement rather than on their ability to reconstruct EMG, because small errors in muscle excitations could cause large errors in task

performance [10,14,15]. In simulation studies, Berniker *et al.* [16] and Kargo *et al.* [17] showed that in the movement of the frog hindlimb, a low-dimensional controller implemented with muscle synergies causes negligible performance degradation. In another simulation study, Meyer *et al.* [18] showed that muscle synergy controls result in more accurate walking kinematic and kinetic predictions compared to independent muscle control. By contrast, for the upper extremity, previous computational work indicates that muscle synergies can introduce substantial aiming errors and limitations in terms of endpoint stiffness and energy consumption during isometric tasks [10,15]. However, this assessment was made with simple biomechanical models with little redundancy, where any dimensionality reduction is likely to introduce substantial errors. For instance, in de Rugy *et al.* [10], synergies introduce larger aiming errors in their simplified wrist model compared to their complex elbow model.

Evidence that modularity aids motor learning could provide additional support for the muscle synergy hypothesis. For example, recent experiments with rats that underwent spinal transections provides evidence that synergies may be formed early postnatally, and the authors suggest that these synergies could be evolutionarily advantageous by simplifying skill acquisition through adulthood [19]. Muscle synergies might speed the learning process by restricting the size of the action space (i.e. the space of possible muscle excitation patterns) [20]. On the other hand, it is possible that a large action space could facilitate skill acquisition as it would provide an abundance of ways to produce the desired behaviour [21]. In Berger *et al.* [22], after a virtual tendon rearrangement surgery, subjects had difficulty adapting to perturbations that were incompatible with synergies identified pre-virtual-surgery, whereas they quickly adapted to perturbations compatible with the synergies. More information about the role of motor modularity in learning new or more complex tasks is needed.

An additional compelling case for the synergy hypothesis would be if a few synergies could generate a wide variety of behaviours. There is some evidence for the generalizability of synergies from experiments in the frog hindlimb during jumping, swimming, and walking movements [1]. Researchers have also demonstrated that the synergies computed from EMG measured in one task can reconstruct EMG measured for other tasks, including postural perturbations in cats [23] and humans [24,25], and isometric force generation in humans [26]. Similarly, it has been shown that only six synergies are necessary to reconstruct EMG accurately for planar reaching movements with different directions and speeds [7], and through visuomotor adaptations [27]. These studies provide promising evidence for the generalizability of synergies, but as discussed above, even small errors in muscle excitations could significantly degrade task performance. In simulation studies, researchers have shown that synergies computed for one walking task facilitate the generation of walking dynamics at different speeds [18] and body weights [28], but it is not known if these results extend to the upper extremity and non-periodic tasks. Giszter suggests that a possible hybrid control architecture, where the CNS has access to both synergies and the control of individual muscles, could potentially avoid the limitations of a reduced control set, while maintaining its advantages [29].

Computational models can help reveal the interplay between synergies and task space errors, motor learning

and generalization, allowing us to quantitatively assess these debated relationships. While computational studies have begun to unravel these relationships, previous models have often been simplified due to the challenges of modelling complex musculoskeletal structures and synthesizing realistic movements. Detailed models of the upper extremity are available [30], along with computational models that synthesize kinematics instead of track experimental data. These computational models are being used to study cause-effect relationships between muscular deficits and movement abnormalities and to test hypotheses in motor neuroscience [31–33].

We implemented muscle synergies in a biomechanically realistic upper extremity model [30] and used optimal control to synthesize reaching movements to a specified target. We focus on point-to-point movements because many voluntary upper extremity movements consist of moving from one pose to another, and being able to hold one or more joints in a static pose [34]. We performed computational experiments to determine whether a modular architecture based on muscle synergies introduces task performance deficits, as defined by a cost term that incorporates target accuracy and muscle effort. We examined the rate the optimization converged to a set performance threshold for models with varying numbers of synergies to assess whether synergies facilitate motor skill acquisition. We also tested whether synergies generated from one set of reaching tasks generalized to new tasks, such as reaching from alternative starting postures or while holding a weight. Finally, we implemented a hybrid controller that included both synergy and individual muscle control to determine whether this architecture improved generalizability to new tasks without sacrificing convergence rates. Overall, our aim was to assess how muscle synergies influence the ability to acquire and generate a repertoire of movements.

2. Methods

2.1. Modelling and simulation

We used an upper extremity biomechanical model developed by Saul *et al.* [30] (available at https://simtk.org/projects/upex_dyn). It consists of 47 Hill-type muscle-tendon actuators [35] with parameters and paths derived from experimental and anatomical studies. The skeletal model has three degrees-of-freedom for the shoulder, one degree-of-freedom for the elbow, and one degree-of-freedom for the forearm to allow for pronation and supination. We removed the wrist degree-of-freedom. The biomechanical modelling and simulation was performed with OpenSim 3.3 [36] using a semi-explicit Euler integrator with an accuracy of 10^{-2} .

2.2. Trajectory optimization

We used a numerical optimal control method to synthesize motion [37]. We optimized for muscle excitations that minimized the following cost function:

$$F = w_1 \|h_N - h_N^d\|^2 + w_2 \sum_{i=1}^N \|a_i\|^2 + w_3 \sum_{i=1}^N l(q_i) + w_4 \|\dot{q}_N\|^2 + w_5 \rho(q_N). \quad (2.1)$$

The terms h_N and h_N^d denote the actual and desired position of the centre of the hand at the last timestep N of the trajectory, a_i denotes the vector of muscle activations at timestep i , $l(q_i)$ is a quadratic penalty on joint limit violations given pose q_i , \dot{q}_N denotes the joint velocities at the last timestep, $\rho(q_N)$ is a

quadratic term to encourage pronation of the forearm at the last timestep (as is typical in reaching movements to a target), and $\|\cdot\|$ denotes the 2-norm. The muscle activations are related to the muscle excitations through a first-order differential equation [35]. We empirically tuned the weights w_1, w_2, w_3, w_4 and w_5 to 5, $\frac{1}{N} \times 10^{-3}$, 10^{-5} , 0.15 and 0.5. We refer to $w_1 \|h_N - h_N^d\|^2$ and $w_2 \sum_{i=1}^N \|a_i\|^2$ as the *task* and *effort* terms, respectively. The cost function encourages the hand to be as close as possible to a target with a small final velocity and the forearm in pronation, while minimizing muscle activations and joint limit violations throughout the movement. We specified the movement duration T and the number of timesteps N in the trajectory was chosen to achieve the desired integration accuracy.

We optimized for the muscle excitations to produce movement with the covariance matrix adaption evolution strategy (CMA-ES) [38]. The free variables were the values of the muscle excitations at every 0.1 s interval in the movement. The excitations were held constant for the 0.1 s interval. The CMA-ES algorithm used a Gaussian distribution over the free variables to sample a number of possible control solutions (referred to as the population size). Given each sample of control solutions (i.e. muscle excitations), we used a forward simulation to obtain a trajectory which we evaluated according to the cost function described above. A new Gaussian distribution was then formed as a function of a subset of the samples with the lowest costs. The CMA-ES algorithm iterated through this process to steer the distribution towards a low-cost region. The number of iterations and the population size in CMA-ES determined the exhaustiveness of the search; we set both to 100 based on empirical testing. We tested a more exhaustive search by doubling the population size and the number of iterations to ensure that it yielded little change to the optimized solutions. We initialized the algorithm with a Gaussian distribution with mean 0.1 and a diagonal covariance matrix with variance 0.09 in all dimensions. We ran our parallelized code on two Intel Xeon CPU E5-4640 processors. The optimization for a single movement took about 6 hours of computation time. In Al Borno *et al.* [31], we show that this method can synthesize movements that replicate kinematic features reported in motor control studies and in experimental three-dimensional reaching data.

2.3. Computing synergies

We incorporated time-invariant muscle synergies in our computational model. According to the time-invariant formulation, the muscle excitations \mathbf{e} can be decomposed as linear combinations of synergies \mathbf{w}_i

$$\mathbf{e}(t) = \sum_{i=1}^K c_i(t) \mathbf{w}_i, \quad (2.2)$$

where c_i denote the combination coefficients, t denotes time and K denotes the number of synergies. When performing trajectory optimization (§2.2) with synergies, the free variables were the combination coefficients c_i instead of the muscle excitations.

We computed muscle synergies when reaching to a target on a horizontal plane from a chosen initial starting pose (figure 1). The plane was placed approximately 20 cm below the top of the clavicle. We randomly chose a target position on the plane, within a squared area 65 cm on a side and solved a trajectory optimization problem to steer the hand as close as possible to the target. We empirically found that 15 trajectory optimizations with random targets was sufficient to cover the workspace (i.e. the domain of possible targets). The muscle excitations were stored in a matrix $\mathbf{E} = [\mathbf{e}_1, \dots, \mathbf{e}_L]$, where L denotes the total number of timesteps from all the synthesized trajectories. We performed dimensionality reduction on \mathbf{E} with principal component analysis (PCA) to compute the muscle synergies. It has been argued that non-negative matrix factorization (NNMF)




Figure 1. Reaching for a target on a horizontal plane. We solved trajectory optimization problems to reach random targets (grey sphere) on a horizontal plane.

computes more physiologically realistic muscle synergies than PCA because it does not allow negative components [39]. In this work, we were mainly interested in studying the computational implications of motor modularity, which could be implemented in different and currently unknown ways. For this reason, we used PCA to compute the modules in our experiments. We used NNMF to assess the sensitivity of our results to the dimensionality reduction technique.

2.4. Evaluating synergies

We next solved trajectory optimization problems to reach new random targets on the plane, except that we constrained the excitations to be linear combinations of synergies by setting the free variables in the optimization to be the synergy combination coefficients, rather than individual muscle excitations. We evaluated how well a task was achieved based on the final value of the cost function in equation (2.1). In addition to the overall cost value, we evaluated the performance of the full and low-dimensional controllers in reaching the target (i.e. the Euclidean distance between the final hand position and the target) and minimizing muscle effort (i.e. the sum of muscle activations squared). We chose every third number between 5 and 20 as the number of synergies to study how performance varies with the amount of dimensionality reduction.

We also investigated the potential consequences of muscle synergies for motor skill acquisition by analysing how muscle synergies impacted the time required to achieve a desired performance level with trajectory optimization. Thus we assumed that optimization time is correlated with motor skill acquisition difficulty (i.e. how much practice is required to achieve a given task). In particular, we compared the number of samples required to achieve the desired performance because sample evaluation was the most computationally expensive operation in the optimization. We said that a desired performance had been achieved when the value of the cost function for the current optimization reached a threshold value, which we set to 0.1. At each iteration in the trajectory optimization, we used the same number of samples as the control dimension (e.g. with eight synergies, we had eight samples). We tried using fewer samples per iteration for the full-dimensional optimization (specifically, 20 instead of 47), but the optimization failed to converge to the desired motor performance.

We then performed computational simulations to determine whether the synergies computed in the reaching task above generalized to different tasks with performance comparable to the full-dimensional optimization. First, we varied the height of the plane by placing it 60 cm higher and 30 cm lower than its initial position (figure 2). We also tested four different initial upper extremity poses chosen to assess the spectrum from which synergies can generalize (figure 2). We next increased the movement speed by making the motion 1.5 times faster. Finally, we tested adding a light (1.75 kg) and heavy (8.7 kg) weight to the hand to model the effect of holding a load. Note that we did not recompute the synergies, as these were fixed from the simulation with the initial plane position and model pose (§2.3).

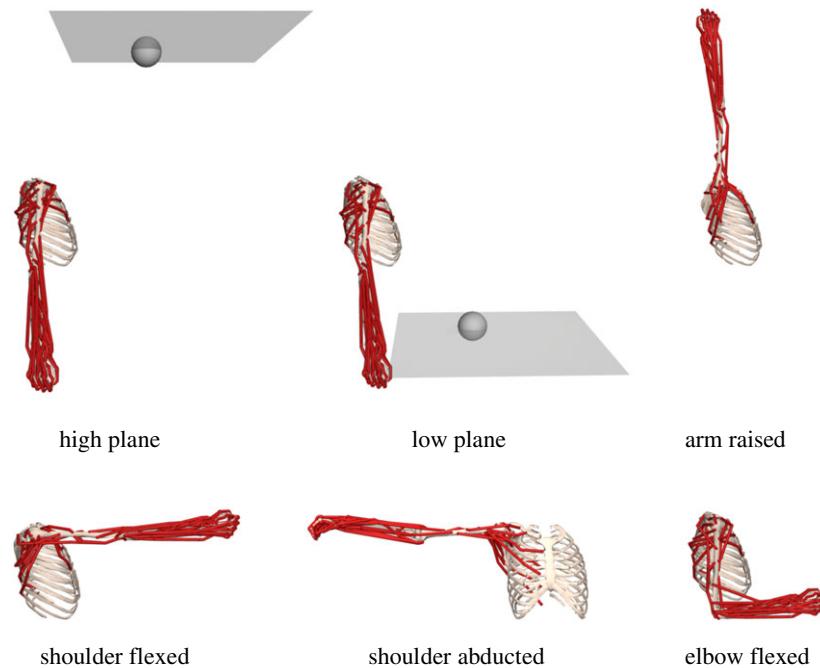


Figure 2. Varying target planes and initial model positions to test the generalizability of synergy-based controllers. We tested the controllers in tasks with a high plane and low plane, with the same initial pose as in the nominal task. We also tested the controllers with varying initial poses, including with the arm raised, shoulder flexed, shoulder abducted and elbow flexed. For the varying initial poses, the targets were in the same plane as in the nominal pose.

2.5. Combining synergies and independent muscle control

Finally, we evaluated a hybrid controller that combined synergies with independent control of each muscle in the model. We compared a model with a hybrid control architecture (i.e. 12 synergies and independent muscle control for a total of 59 control variables) to independent muscle control (i.e. the full-dimensional system with 47 variables) and a model with 12 synergies (and thus 12 control variables). We chose to have 12 synergies in our hybrid architecture to balance between flexibility and learning speed. We used the same number of synergies in the hybrid architecture as in the synergies architecture to aid the comparison. The optimization procedure followed the descriptions in §§2.2 and 2.3. We set the initial variance in the optimization to 0.0025 for the independent muscles and to 0.09 for the synergies. This means that the optimization prioritized finding a control solution with muscle synergies and started exploring using individual muscles as a secondary step to improve control.

The advantage of the hybrid architecture can be observed when optimizing for a movement outside the space where the synergies are computed. To illustrate this case, we took the initial position of the arm to be as in the experiment with the arm raised (figure 2), while the synergies were computed with the arm in a neutral position (figure 1). To determine the effect of the hybrid architecture on convergence rates (and thus the potential effects on skill acquisition), we also investigated the number of samples required to achieve a desired motor performance (i.e. a cost function value less than 0.1) for the task of reaching to targets in the same plane on which the synergies were created.

3. Results

3.1. Synergies do not degrade task performance

Twenty synergies were required to account for 88% of the variance in 47 muscle excitation patterns for the reaching task in a single horizontal plane (figure 3a). While five synergies explained only 49% of the variance, with five or more synergies,

the low-dimensional optimization achieved a performance within 0.011 or 10% of the full-dimensional optimization for new, random reaching targets in the same plane for which synergies were extracted (figure 3b). Differences in the overall cost function values of this order represented solutions that are not easily visually distinguishable (see electronic supplementary material, video 1) and are near the variation from run-to-run due to the stochastic nature of the optimization. In particular, over 10 runs reaching to the same target, there was variability of the order of 5% of the total cost, 0.5 cm in the final hand position, and of 10% in the movement effort (defined as the sum of muscle activations squared).

To compare the task performance of the synergy-based controllers to the full-dimensional controller, we set the population size and the number of iterations in CMA-ES to 100 for the synergies and we doubled the population size and the number of iterations for the full-dimensional controller (label '47E' in figure 3b). On average, the final hand position with five and eight synergies was 0.6 (1.0) cm and 0.2 (0.8) cm further from the target, respectively, than the hand position with the full-dimensional solution (we provide the standard deviation in parentheses throughout the paper). While the average error was low for the five-synergy solution, there were instances where the five-synergy controller performed poorly. For example, the target error for the worst run with five synergies was 2.8 cm further away from the target than the full-dimensional solution. With three synergies, the low-dimensional optimization was on average 4.2 (4.7) cm further away from the target than the full-dimensional optimization, an order of magnitude worse than the five and eight synergy solutions. The full-dimensional solution required on average 1.6 (0.3) and 1.3 (0.2) times more effort than the five and eight synergies solutions, respectively. In fact, the low-dimensional optimizations often achieved a lower cost function value since the full-dimensional optimization was more prone to fall in local minima. We found that the optimization tended to converge to movements requiring more effort when more




Figure 3. Synergies and task space performance. (a) Variance in muscle excitations explained versus the number of computed synergies. (b) Average cost (equation (2.1)) for 10 random targets in the plane of figure 1 when the trajectory optimization was performed with increasing control dimensionality. The cost includes weighted terms to represent task performance (blue), effort (green) and other (yellow) criteria such as final hand velocity near 0 and joint limit violations. Error bars represent one standard deviation of the cost function value. We set the same population size and the number of iterations in the optimization for the different controllers, except for 47E where we conducted a more exhaustive search by doubling the population size and the number of iterations.

control variables were allowed. This local minimum explains why the cost can increase with the number of control variables. When comparing the kinematics of the movements, qualitatively, the synergies produced similar, but simplified curves compared to the full-dimensional system (figure 4). We show the movements being performed by the different controllers in electronic supplementary material, video S2.

We also conducted a search with the full-dimensional controller using the same population size and number of iterations (100) as for the synergy-based controllers (label ‘47’ in figure 3b). This less exhaustive search did not have a large effect on performance. The cost function for the exhaustive search was 0.016 (0.01) lower. On average, this caused the final hand position to be 0.6 (0.4) cm closer to the target, while requiring 16 (20)% less effort than the less exhaustive optimization. The more exhaustive optimization required 25 h of computation on an Intel Xeon machine with 20 cores for a single movement (as opposed to 6 h) and yielded relatively small improvements. This indicates that the parameters for the full-dimensional optimization were adequate to converge to good local minima.

3.2. Synergies facilitate motor learning

A controller with 8 or 20 synergies required about 20 and 14 times fewer samples, respectively, to converge to a cost value of 0.1 or less, compared to the optimization with 47 muscles (figure 5a). Note that even computing synergies without dimensionality reduction (i.e. having 47 synergies) helped motor learning as the control variables partially encoded the solution, requiring approximately three times fewer samples. The eight synergies optimization required about 18 iterations to achieve the same motor performance (evaluated as the current optimization minimum) that the full-dimensional optimization achieved at about 50 iterations, while requiring fewer samples at each iteration (figure 5b). Furthermore, the full-dimensional optimization converged to a worse local minimum than the optimization with eight synergies. The low-dimensional optimization achieved a final hand position that was 0.21 cm closer to the target and was 2.4 times less costly in terms of the sum of muscle activations squared. We also tested whether randomly generated synergies would speed up learning. None of the five randomly generated, low-dimensional controllers we

tested converged to the desired performance (see electronic supplementary material, figure S2). This indicates that use of a low-dimensional basis alone does not speed up learning.

3.3. Synergies generalize to new tasks

The optimizations with eight or more muscle synergies converged to solutions that were similar in cost function value (i.e. on average within 0.005 or better) to the full-dimensional optimization for many, although not all, of the new tasks that we simulated (figure 6). When the plane of targets was raised 60 cm higher, the final hand position with eight synergies was within 1.4 (1.9) cm of the full-dimensional optimization and had 23 (15)% times less effort (figure 6a). When the plane was lowered 30 cm and for three of the alternative initial poses (elbow flexed, shoulder flexed and shoulder abducted), all the low-dimensional controllers (i.e. 5 or greater synergies) also achieved an average performance within 0.005 (or better) of the full-dimensional controller, which corresponds to an average difference of 0.36 (0.79) cm in the final hand positions and of 28 (13)% in the total effort (figure 6a). While the differences in the percentage of the total effort were large, the effort required for these movements was relatively small; hence, the effort has a small contribution to the overall cost function value. For example, when reaching to the lower plane, the effort terms accounted on average for only 20% of the overall cost. For the initial pose with the arm raised (180 degrees shoulder flexion), the controller needed at least 17 synergies to achieve a comparable (or better) performance to the full-dimensional controller (figure 6b). Note that the initial pose for this task was quite different from the initial pose where the synergies were computed (figure 1). The model was able to execute the reaching task at higher speed and with the 1.75 kg weight using low-dimensional controllers with eight or more synergies (figure 6c). The low-dimensional solution with eight synergies was on average 0.25 (0.9) cm closer to the target and requires 33 (9)% less total effort. The reaching task with the 8.7 kg weight was a case where the task-specific synergies did not generalize effectively (figure 6d), as the full-dimensional solution performed substantially better on equation (2.1). In this case, the full-dimensional optimization was on average 13.0 (7.7) cm closer to the target than the solution with eight synergies. This is due to the fact that the synergies did not produce sufficient joint torques to hold and move the heavy object because the

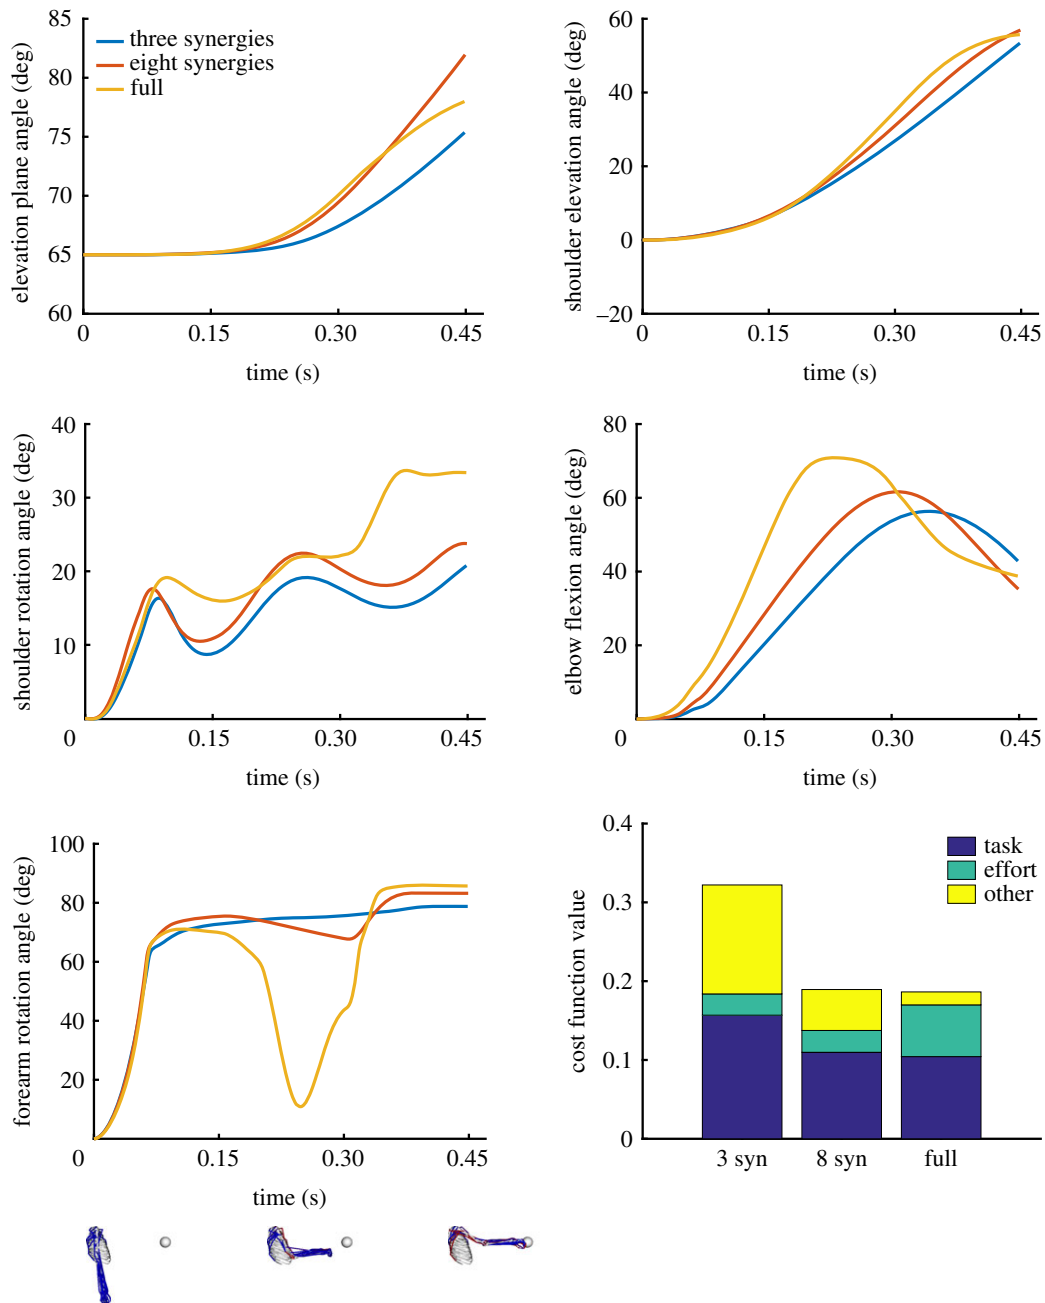


Figure 4. Example kinematics when optimizing reaching to a target with three and eight synergies, and the full-dimensional system. Kinematics are shown for the model's five degrees-of-freedom, which include the elevation plane, shoulder elevation (or thoracohumeral angle), shoulder rotation (where internal rotation is positive), elbow flexion and forearm rotation (where pronation is positive) [30]. Kinematics are from a representative run, but the multiple runs were qualitatively similar. We also show some key poses in the movement for the full-dimensional controller. In the bottom rightmost panel, we compare the performance on the cost for each controller for the kinematics in the plots above. We show the proportion of the cost due to the task term, the effort term and all other terms (final velocity near 0, joint limits, etc.; see equation (2.1)).

database of movements from which the synergies were computed did not require large joint torques. We again verified that our results were robust to more exhaustive optimizations (i.e. doubling the population size and the number of iterations in §2.2 yielded the same trend for the new tasks).

3.4. Synergies and independent muscle control

The synergy-only architecture had relatively low performance on the task that required reaching from an arm raised position; however, a hybrid architecture with 12 synergies, along with independent control of the model's 47 muscles achieved a lower cost than the full-dimensional optimization, while still being nearly as sample-efficient as the synergy

solution. On average, the 12 synergies solution was 6.9 (2.5) cm further away from the target than the full-dimensional solution, while the hybrid solution was 1.8 (2.5) cm further away (figure 7a). On average over 10 runs, the hybrid architecture achieved the desired motor performance for the same-plane reaching task using 8.5 fewer samples than the full-dimensional optimization, while the synergies optimization used 20 times fewer samples (figure 7b).

4. Discussion

We developed a computational model of upper extremity movement to study the muscle synergy hypothesis for

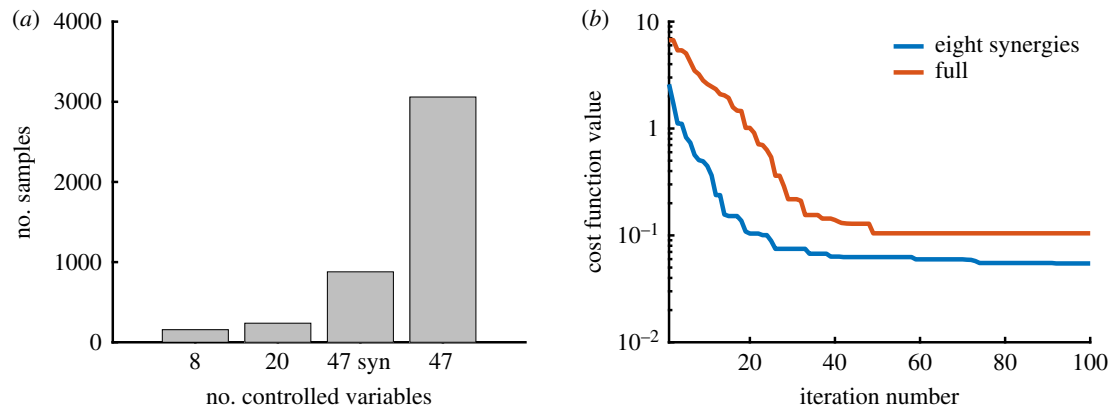


Figure 5. Effect of synergies on search complexity. (a) The number of samples required to achieve a cost function value less than 0.1 based on the number of controlled variables, averaged over ten random runs. We use '47 syn' to indicate controlling 47 synergies, instead of controlling the 47 muscles independently. (b) Convergence rate for the controller with eight synergies (blue) compared to the full-dimensional system (red).

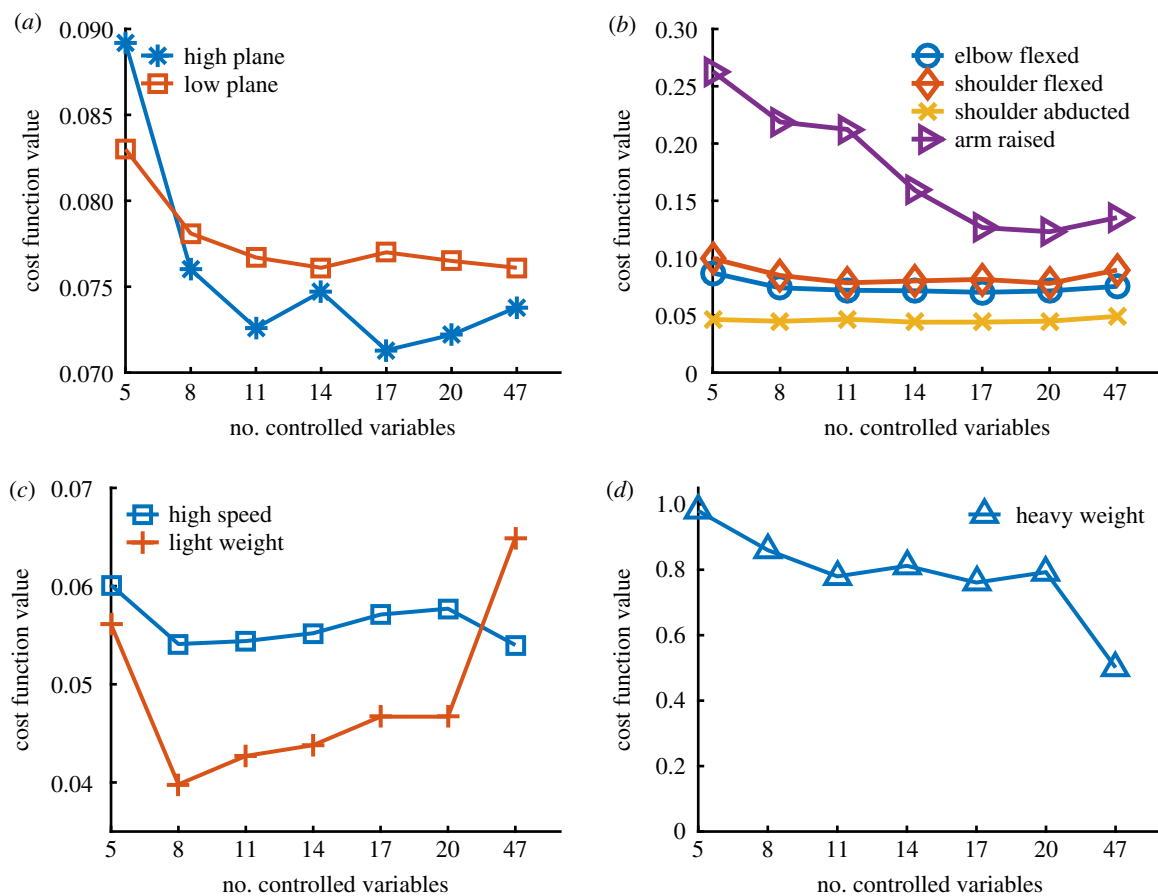


Figure 6. Performance of synergy-based controllers for new tasks. We show the average final cost depending on the number of controlled synergies for several new tasks. (a) We varied the height of the horizontal plane. The horizontal plane was 60 cm higher and 30 cm lower in the high plane and low plane conditions, respectively. (b) We varied the initial arm position and the goal was to reach a target on the same horizontal plane as in the nominal task. (c) The movement was performed 1.5 times faster and we placed a load of 1.75 kg on the hand. (d) We placed a load of 8.7 kg on the hand. All of these results are the average of 10 trajectory optimizations with random targets.

voluntary movement generation. The upper extremity model is more biomechanically realistic than previous work, but it is also more complex to control due to its many muscles and degrees-of-freedom. Our computational experiments indicate that a low-dimensional controller based on muscle synergies does not lead to a significant loss of performance, generates a wide variety of behaviours, and supports the acquisition of skills at a faster rate than a full-dimensional controller (up to a 20-fold increase in some cases). A study indicating the converse would pose a serious challenge for the synergy hypothesis. Our results do not imply the existence of

synergies, but highlight and quantify their potential benefits. Further speed-ups in motor learning could potentially be achieved with the time-variant formulation of muscle synergies by capturing the time dependence of muscle activations [40], but one would also need to study its effects on movement generalization. We also showed that muscle synergies can be combined with the control of individual muscles to maintain the ability to quickly acquire motor skills, while being able to generalize to more novel tasks.

For our nominal reaching task, 20 synergies were needed to explain 88% of the variance in the muscle excitations. By

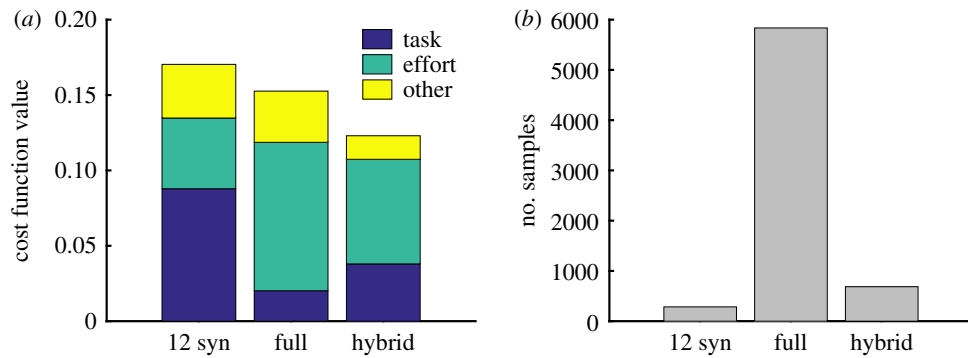


Figure 7. Performance of the hybrid controller. (a) Average cost function value for 12-synergy, full-dimensional and hybrid architectures across five random targets with the arm raised initial pose (figure 2). We show the proportion of the total cost due to the task term (blue), the effort term (green) and all other terms (yellow) in equation (2.1). (b) Number of samples required to achieve a cost function value less than 0.1 based on the number of controlled variables, averaged over 10 random runs. Although the hybrid control architecture had 59 variables (12 synergies and 47 independent muscles), it still had a more than eightfold reduction in the number of samples required to achieve a desired performance when reaching to a target on a horizontal plane (figure 1).

contrast, experiments with rhesus monkeys indicate that 10 synergies are sufficient to capture 95% of the variance in the reaching EMG data [41]. In human reaching movements on a plane, 4–5 time-varying synergies capture 73–82% of the variance observed in experiments [42]; a larger number of time-invariant synergies is expected to achieve the same variance explained. In experimental studies, EMG is typically only recorded from at most 19 muscles of the upper extremity due to measurement difficulties. We thus expect a larger number of synergies would be needed to explain most of the variance with our simulations that used 47 muscles to control reaching.

Inouyue *et al.* [15] argue that muscle synergies can cause a drastic reduction in task space performance, and thus per cent variance explained may not be a good metric of synergy performance. However, this prior analysis was performed with a simplified biomechanical model, and the authors acknowledge that ‘having many more muscles will naturally allow the implementation of synergies without such a drastic reduction’. This is the conclusion we have reached in our work. We did not see a significant performance degradation when using eight synergies that only explained 62% of the variance. Another difference that could explain the divergent results is that the synergies in Inouyue *et al.* were hand-tuned (and possibly less effective), while our synergies were obtained through optimization. We have also shown that our synergy-based controller often achieved better performance than the full-dimensional optimization, which was more prone to fall in local minima.

The muscle synergies hypothesis is appealing in part because it potentially simplifies the control and learning of movements [43]. Loeb *et al.* [21] hypothesized that a large action space could facilitate skill acquisition as it would provide an abundance of ways of producing the desired behaviour. Our computational study does not support this hypothesis as the large action space made learning more costly by a numerical optimizer. This result is consistent with the experiments by Berger *et al.* [22] that indicate that it is more difficult (i.e. requires more trials and time) to learn movements that are incompatible with existing experimentally determined muscle synergies than to learn movements that are consistent with measured synergies. Hagio *et al.* [44] previously showed that modularity speeds up motor adaptation to rotational perturbations in an isometric force production task. We extend their result by studying reaching movements and by studying the effects of modularity on motor learning, not just adaptation.

Human motor learning can be framed in terms of optimizing a movement policy or reinforcement learning [45]. Our evolutionary algorithm for trajectory optimization (§2.2) shares similarities with reinforcement learning and can be seen as a scalable alternative [46]. Hence, the number of samples required to achieve a motor performance in our model may provide a first-order assessment of the motor learning difficulty of the task in humans. If future research shows that humans learn in ways fundamentally different from our current reinforcement learning models, then our assessment will need to be re-evaluated. State-of-the-art deep reinforcement learning approaches in motor control currently require a prohibitive number of samples to learn new movements, with limited fidelity [47]. It would be interesting to investigate whether modular architectures can help scale these approaches from torque-driven to muscle-based models to achieve greater motion realism. Synergy-based control architectures may also offer an alternative to recent work in computer graphics that attempt to simplify the computational problem by learning realistic torque-driven models instead of muscle-based models directly [48].

There is growing evidence that synergies can generalize across tasks [7,25–27], and our study provides further support. We go beyond showing the ability of synergies to reconstruct EMG across tasks and demonstrate that synergies from one reaching task enable task performance in other reaching tasks with accuracy equivalent to a full-dimensional controller. Sohn & Ting [49] present evidence that generalizability does not arise directly from musculoskeletal or optimality constraints and causes an increase in effort. We complement this study by showing that synergies are generalizable for upper extremity reaching tasks, although we did not see large cost increases for our synergy solutions, likely because we included an effort term for all of our optimizations. The generalization tests designed in our study limit basis generation to a subset of the workspace (i.e. reaching to random targets on a horizontal plane). In an organism, the optimal basis (as might be discovered on evolutionary timescales) can likely only be well identified from a richer dataset, involving a larger workspace (i.e. given the common tasks that an organism must perform). Nevertheless, our study suggests that synergies generalize well outside the workspace, which is valuable independently of how the workspace is defined.

We determined the muscle synergies with PCA as we were interested in analysing the computational implications of

motor modularity without regard to the actual physiological implementation (e.g. in the spinal circuitry [50] or in cortex [51]), which is unknown. We observed a decrease in performance when performing the experiments with NNMF, using the solver provided in Kim *et al.* [52]. For instance, the average cost over five runs when reaching in the horizontal plane was 0.04 higher with NNMF. In task space, this difference amounts to the hand being 2.6 cm further from the target with twice the effort (as measured by the sum of muscle activations squared) compared to the PCA solution. When reaching on the 60 cm higher horizontal plane (outside the space where the synergies were computed), the NNMF cost is 0.069 higher. This reduced performance with NNMF could be due to the fact that PCA generalizes better to data outside the subspace where dimensionality reduction occurs [53]. The synergies determined from NNMF are included in the electronic supplementary material. Similarly to the results reported in d'Avella *et al.* [7,42,54], the first synergy is related to elbow flexion, while synergies 2–4 are related to elbow extension and shoulder flexion.

There are several other modelling simplifications and assumptions that may have impacted our study. First, we did not model spinal circuitry feedback or individual motor units. Second, choosing a different cost function or weights in our cost function might lead to different solutions. However, we chose a cost function and optimization strategy that we have previously shown produces realistic movements with features reported in motor studies [31]. Finally, since our trajectory optimizer is stochastic, it is more likely to return a local minima than the global minimum. In our experiments, the performance of the full-dimensional optimization was not always optimal (e.g. see figure 6). The reason is that it is computationally difficult to find the optimal solution with such a high-dimensional system. We chose to terminate our optimizations when the maximum number of iterations was reached for all cases except the learning experiments (where we stopped iterating when a specified performance threshold was achieved). We tested stricter convergence criteria and did not observe markedly different solutions. Our results in §2.4 indicated that, when employing more exhaustive optimizations, most of the improvements were in the effort term (up to a 16 (20)% improvement in some cases). The standard deviation in the effort term was large, but this is

partly due to the fact that this term can have a relatively small contribution to the overall cost function (the optimization aims to reduce the overall cost function value). The improvements in the final hand positions were less than one centimetre. These results indicate that our optimization parameters were sufficient to achieve good local minima.

Our results in §3.1 indicate that we should expect variability of the order of half a centimetre in the final hand position and of 10% in the movement effort due to the stochastic optimization. This also limits the resolution of our study: it becomes difficult to assess the implications of synergies at this magnitude. However, we have chosen to perform our analysis on tasks and movements (i.e. three-dimensional upper extremity reaching) where this magnitude is typically not significant.

We have shown that modularity speeds motor learning with our evolutionary optimization algorithm. There may exist other (potentially unknown) algorithms where having independent control of each muscle does not significantly impede motor learning. This touches upon the age-old question of whether intelligent behaviour arises because of a prior, possibly innate, structure or because of general learning algorithms. Our results indicate that a modular architecture as a prior structure for reaching movements could improve motor learning, without significantly degrading performance and generalizability. Synergies may be favoured by Darwinian natural selection to speed up both motor development and skill acquisition, a hypothesis that is consistent with results in the lower limb that suggest that synergies are conserved across individuals with significantly different motor experiences and even across species [12,13,19]. How this prior structure—if it exists—is implemented remains an open question.

Data accessibility. The code for the computational simulations and our data are available at <https://simtk.org/projects/ue-synergies>.

Authors' contributions. M.A. conceived and implemented the study, and drafted the manuscript. J.H. helped interpret results and drafted the manuscript. S.D. provided edits to the manuscript and supervised all aspects of this work.

Competing interests. We declare we have no competing interests.

Funding. The authors were supported by the Mobilize Center, a National Institutes of Health Big Data to Knowledge (BD2K) Center of Excellence through grant no. U54EB020405.

References

- d'Avella A, Bizzi E. 2005 Shared and specific muscle synergies in natural motor behaviors. *Proc. Natl Acad. Sci. USA* **102**, 3076–3081. (doi:10.1073/pnas.0500199102)
- Bizzi E, Cheung VC. 2013 The neural origin of muscle synergies. *Front. Comput. Neurosci.* **7**, 51. (doi:10.3389/fncom.2013.00051)
- Kargo WJ, Giszter SF. 2008 Individual premotor drive pulses, not time-varying synergies, are the units of adjustment for limb trajectories constructed in spinal cord. *J. Neurosci.* **28**, 2409–2425. (doi:10.1523/JNEUROSCI.3229-07.2008)
- Lafreniere-Roula M, McCrea DA. 2005 Deletions of rhythmic motoneuron activity during fictive locomotion and scratch provide clues to the organization of the mammalian central pattern generator. *J. Neurophysiol.* **94**, 1120–1132. (doi:10.1152/jn.00216.2005)
- Chvatal SA, Ting LH. 2013 Common muscle synergies for balance and walking. *Front. Comput. Neurosci.* **7**, 48. (doi:10.3389/fncom.2013.00048)
- Neptune RR, Clark DJ, Kautz SA. 2009 Modular control of human walking: a simulation study. *J. Biomech.* **42**, 1282–1287. (doi:10.1016/j.jbiomech.2009.03.009)
- d'Avella A, Fernandez L, Portone A, Lacquaniti F. 2008 Modulation of phasic and tonic muscle synergies with reaching direction and speed. *J. Neurophysiol.* **100**, 1433–1454. (doi:10.1152/jn.01377.2007)
- Muceli S, Boye AT, Farina D. 2010 Identifying representative synergy matrices for describing muscular activation patterns during multidirectional reaching in the horizontal plane. *J. Neurophysiol.* **103**, 1532–1542. (doi:10.1152/jn.00559.2009)
- De Groot F, Jonkers I, Duysens J. 2014 Task constraints and minimization of muscle effort result in a small number of muscle synergies during gait. *Front. Comput. Neurosci.* **8**, 115. (doi:10.3389/fncom.2014.00115)
- de Rugy A, Loeb GE, Carroll TJ. 2013 Are muscle synergies useful for neural control? *Front. Comput. Neurosci.* **7**, 19. (doi:10.3389/fncom.2013.00019)
- Nazarpour K, Barnard A, Jackson A. 2012 Flexible cortical control of task-specific muscle synergies. *J. Neurosci.* **32**, 12 349–12 360. (doi:10.1523/JNEUROSCI.5481-11.2012)

12. Dominici N *et al.* 2011 Locomotor primitives in newborn babies and their development. *Science* **334**, 997–999. (doi:10.1126/science.1210617)
13. Hart CB, Giszter SF. 2010 A neural basis for motor primitives in the spinal cord. *J. Neurosci.* **30**, 1322–1336. (doi:10.1523/JNEUROSCI.5894-08.2010)
14. Alessandro C, Delis I, Nori F, Panzeri S, Berret B. 2013 Muscle synergies in neuroscience and robotics: from input-space to task-space perspectives. *Front. Comput. Neurosci.* **7**, 43. (doi:10.3389/fncom.2013.00043)
15. Inouye JM, Valero-Cuevas FJ. 2016 Muscle synergies heavily influence the neural control of arm endpoint stiffness and energy consumption. *PLoS Comput. Biol.* **12**, e1004737. (doi:10.1371/journal.pcbi.1004737)
16. Berniker M, Jarc A, Bizzi E, Tresch MC. 2009 Simplified and effective motor control based on muscle synergies to exploit musculoskeletal dynamics. *Proc. Natl Acad. Sci. USA* **106**, 7601–7606. (doi:10.1073/pnas.0901512106)
17. Kargo WJ, Ramakrishnan A, Hart CB, Rome LC, Giszter SF. 2010 A simple experimentally based model using proprioceptive regulation of motor primitives captures adjusted trajectory formation in spinal frogs. *J. Neurophysiol.* **103**, 573–590. (doi:10.1152/jn.01054.2007)
18. Meyer AJ, Eskinazi I, Jackson JN, Rao AV, Patten C, Fregly BJ. 2016 Muscle synergies facilitate computational prediction of subject-specific walking motions. *Front. Bioeng. Biotechnol.* **4**, 77. (doi:10.3389/fbioe.2016.00077)
19. Yang Q, Logan D, Giszter SF. 2019 Motor primitives are determined in early development and are then robustly conserved into adulthood. *Proc. Natl Acad. Sci. USA* **116**, 12 025–12 034. (doi:10.1073/pnas.1821455116)
20. Thomas PS, Barto AG. 2012 Motor primitive discovery. In *2012 IEEE Int. Conf. on Development and Learning and Epigenetic Robotics, San Diego, CA, USA, 7–9 November 2012*. (doi:10.1109/DevLrn.2012.6400845)
21. Loeb GE. 2012 Optimal isn't good enough. *Biol. Cybern* **106**, 757–765. (doi:10.1007/s00422-012-0514-6)
22. Berger DJ, Gentner R, Edmunds T, Pai DK. 2013 Differences in adaptation rates after virtual surgeries provide direct evidence for modularity. *J. Neurosci.* **33**, 12 384–12 394. (doi:10.1523/JNEUROSCI.0122-13.2013)
23. Torres-Oviedo G, Macpherson JM, Ting LH. 2006 Muscle synergy organization is robust across a variety of postural perturbations. *J. Neurophysiol.* **96**, 1530–1546. (doi:10.1152/jn.00810.2005)
24. Chvatal SA, Torres-Oviedo G, Savafynia SA, Ting LH. 2011 Common muscle synergies for control of center of mass and force in nonstepping and stepping postural behaviors. *J. Neurophysiol.* **106**, 999–1015. (doi:10.1152/jn.00549.2010)
25. Torres-Oviedo G, Ting LH. 2010 Subject-specific muscle synergies in human balance control are consistent across different biomechanical contexts. *J. Neurophysiol.* **103**, 3084–3098. (doi:10.1152/jn.00960.2009)
26. Roh J, Rymer WZ, Beer RF. 2012 Robustness of muscle synergies underlying three-dimensional force generation at the hand in healthy humans. *J. Neurophysiol.* **107**, 2123–2142. (doi:10.1152/jn.00173.2011)
27. Gentner R, Edmunds T, Pai DK. 2013 Robustness of muscle synergies during visuomotor adaptation. *Front. Comput. Neurosci.* **7**, 120. (doi:10.3389/fncom.2013.00120)
28. McGowan CP, Neptune RR, Clark DJ, Kautz SA. 2010 Modular control of human walking: adaptations to altered mechanical demands. *J. Biomech.* **43**, 412–419. (doi:10.1016/j.jbiomech.2009.10.009)
29. Giszter SF. 2015 Motor primitives—new data and future questions. *Curr. Opin. Neurobiol.* **33**, 156–165. (doi:10.1016/j.conb.2015.04.004)
30. Saul KR, Hu X, Goehler CM, Vidt ME, Daly M, Velisar A, Murray WM. 2015 Benchmarking of dynamic simulation predictions in two software platforms using an upper limb musculoskeletal model. *Comput. Methods Biomech. Biomed. Eng.* **18**, 1445–1458. (doi:10.1080/10255842.2014.916698)
31. Al Borno M, Vyas S, Shenoy KV, Delp SL. 2019 High-fidelity musculoskeletal modeling reveals a motor planning contribution to the speed-accuracy tradeoff. *bioRxiv* 804088. (doi:10.1101/804088)
32. DeWolf T, Stewart TC, Slotine J-J, Eliasmith C. 2016 A spiking neural model of adaptive arm control. *Proc. R. Soc. B* **283**, 20162134. (doi:10.1098/rspb.2016.2134)
33. Flash T, Hogan N. 1985 The coordination of arm movements: an experimentally confirmed mathematical model. *J. Neurosci.* **5**, 1688–1703. (doi:10.1523/JNEUROSCI.05-07-01688.1985)
34. Shadmehr R. 2017 Distinct neural circuits for control of movement vs. holding still. *J. Neurophysiol.* **117**, 1431–1460. (doi:10.1152/jn.00840.2016)
35. Zajac FE. 1989 Muscle and tendon: properties, models, scaling, and application to biomechanics and motor control. *Crit. Rev. Biomed. Eng.* **17**, 359–411.
36. Delp SL, Anderson FC, Arnold AS, Loan P, Habib A, John CT, Guendelman E, Thelen DG. 2007 Opensim: open-source software to create and analyze dynamic simulations of movement. *IEEE Trans. Biomed. Eng.* **54**, 1940–1950. (doi:10.1109/TBME.2007.901024)
37. Al Borno M, De Lasa M, Hertzmann A. 2013 Trajectory optimization for full-body movements with complex contacts. *IEEE Trans. Vis. Comput. Graph* **19**, 1405–1414. (doi:10.1109/TVCG.2012.325)
38. Hansen N. 2006 The CMA evolution strategy: a comparing review. In *Towards a new evolutionary computation* (eds JA Lozano, P Larrañaga, I Inza, E Bengoetxea). Studies in Fuzziness and Soft Computing, vol. 192, pp. 75–102. Berlin, Germany: Springer. (doi:10.1007/3-540-32494-1_4)
39. Tresch MC, Cheung VC. 2006 Matrix factorization algorithms for the identification of muscle synergies: evaluation on simulated and experimental data sets. *J. Neurophysiol.* **95**, 2199–2212. (doi:10.1152/jn.00222.2005)
40. d'Avella A, Lacquaniti F. 2013 Control of reaching movements by muscle synergy combinations. *Front. Comput. Neurosci.* **7**, 42. (doi:10.3389/fncom.2013.00042)
41. Overduin SA, Roh J, Carmena JM, Bizzi E. 2015 Representation of muscle synergies in the primate brain. *J. Neurosci.* **35**, 12 615–12 624. (doi:10.1523/JNEUROSCI.4302-14.2015)
42. d'Avella A, Portone A, Fernandez L, Lacquaniti F. 2006 Control of fast-reaching movements by muscle synergy combinations. *J. Neurosci.* **26**, 7791–7810. (doi:10.1523/JNEUROSCI.0830-06.2006)
43. Bernstein N. 1967 *The co-ordination and regulation of movements*. New York, NY: Pergamon.
44. Hagio S, Kouzaki M. 2018 Modularity speeds up motor learning by overcoming mechanical bias in musculoskeletal geometry. *J. R. Soc. Interface* **15**, 20180249. (doi:10.1098/rsif.2018.0249)
45. Haith AM, Krakauer JW. 2013 Model-based and model-free mechanisms of human motor learning. In *Progress in motor control* (eds M Richardson, M Riley, K Shockley). Advances in Experimental Medicine and Biology, vol. 782, pp. 1–21. New York, NY: Springer. (doi:10.1007/978-1-4614-5465-6_1)
46. Salimans T, Ho J, Chen X, Sidor S, Sutskever I. 2017 Evolution strategies as a scalable alternative to reinforcement learning. (<http://arxiv.org/abs/1703.03864>)
47. Lillicrap TP, Hunt JJ, Pritzel A, Heess N, Erez T, Tassa Y, Silver D, Wierstra D. 2015 Continuous control with deep reinforcement learning. (<http://arxiv.org/abs/1509.02971>)
48. Jiang Y, Van Wouwe T, De Groot F, Liu CK. 2019 Synthesis of biologically realistic human motion using joint torque actuation. *ACM Trans. Graph.* **38**, 72. (doi:10.1145/3306346.3322966)
49. Sohn MH, Ting LH. 2016 Suboptimal muscle synergy activation patterns generalize their motor function across postures. *Front. Comput. Neurosci.* **10**, 7. (doi:10.3389/fncom.2016.00007)
50. d'Avella A, Saltiel P, Bizzi E. 2003 Combinations of muscle synergies in the construction of a natural motor behavior. *Nat. Neurosci.* **6**, 300–308. (doi:10.1038/nn1010)
51. Gallego JA, Perich MG, Naufel SN, Ethier C, Solla SA, Miller LE. 2018 Cortical population activity within a preserved neural manifold underlies multiple motor behaviors. *Nat. Commun.* **9**, 4233. (doi:10.1038/s41467-018-06560-z)
52. Kim J, Park H. 2011 Fast nonnegative matrix factorization: an active-set-like method and comparisons. *SIAM J. Sci. Comput.* **33**, 3261–3281. (doi:10.1137/110821172)
53. Ting LH, Chvatal SA. 2010 Decomposing muscle activity in motor tasks. In *Motor control: theories, experiments, and applications* (eds F Danion, ML Latash), pp. 102–138. New York, NY: Oxford University Press.
54. Scano A, Chiavenna A, Malosio M, Molinari Tosatti L, Molteni F. 2017 Muscle synergies-based characterization and clustering of poststroke patients in reaching movements. *Front. Bioeng. Biotechnol.* **5**, 62. (doi:10.3389/fbioe.2017.00062)