


Article

Learning to manipulate a whip with simple primitive actions – A simulation study



Moses C. Nah,
Aleksei Krotov,
Marta Russo,
Dagmar Sternad,
Neville Hogan

mosesnahn@mit.edu

Highlights

Human-like elementary actions enabled complex object manipulation

This approach is more efficient than modern learning approaches

This approach may facilitate control of soft and flexible objects



Article

Learning to manipulate a whip with simple primitive actions – A simulation study

Moses C. Nah,^{1,7,*} Aleksei Krotov,² Marta Russo,^{3,4} Dagmar Sternad,⁵ and Neville Hogan^{1,6}

SUMMARY

This simulation study investigated whether a 4-degrees-of-freedom (DOF) arm could strike a target with a 50-DOF whip using a motion profile similar to discrete human movements. The interactive dynamics of the multi-joint arm was modeled as a constant joint-space mechanical impedance, with values derived from experimental measurement. Targets at various locations could be hit with a single maximally smooth motion in joint-space coordinates. The arm movements that hit the targets were identified with fewer than 250 iterations. The optimal actions were essentially planar arm motions in extrinsic task-space coordinates, predominantly oriented along the most compliant direction of both task-space and joint-space mechanical impedances. Of the optimal movement parameters, striking a target was most sensitive to movement duration. This result suggests that the elementary actions observed in human motor behavior may support efficient motor control in interaction with a dynamically complex object.

INTRODUCTION

Human dexterity is remarkable. Our ability to create and use tools has given us a significant evolutionary advantage. On a daily basis, we manipulate objects with many degrees of freedom (DOF) that exhibit complex behavior — folding a newspaper, tying shoelaces, donning a jacket. A more exotic example is cracking a whip, which confronts us with the daunting complexity of the tools that humans can master. A whip is a deformable structure with non-uniform mechanical properties, that interacts with the complex compressible fluid dynamics of air. If a whip is ‘cracked’, it operates into the supersonic regime.^{1–3} A physical model to describe the fluid-structure interaction requires nonlinear partial differential equations of infinite order. Nevertheless, humans can learn to manipulate a whip, with some ‘whip masters’ reaching impressive levels of expertise — a trained whip master can halve an apple on a person’s head without touching their hair.^{4,5}

How do humans achieve such astonishing dexterity? A recent experimental study examined how humans strike a target with a whip manipulating it in discrete and rhythmic fashion.^{5,6} Analysis of both movement styles showed that the hand motion exhibited a pronounced speed peak during the throwing action. The motion profile was reminiscent of the bell-shaped speed profiles observed in a plethora of other behaviors: discrete human actions such as point-to-point motions exhibit a robustly repeatable pattern with a roughly symmetrical bell-shaped speed profile, both for horizontal-plane motions (which are generally straight) and for vertical-plane motions (which are usually curved).^{7–9}

This bell-shaped speed profile appears to characterize one of the building blocks or ‘primitives’ of human motor control.¹⁰ The earliest movements made by persons recovering after a stroke (cerebral vascular accident) are a sequence of submovements with this profile; and they progressively overlap and blend as recovery proceeds.¹¹ One or more such (sub)movements with this speed profile would be well-suited to propagate a pulse of deformation into and along a whip, thereby exciting a robust traveling-wave behavior. Following the analysis of experimental data in human whip manipulation,⁵ this study used simulation to test whether such discrete movements were sufficient to generate whip motions that hit a target.

A possible advantage of using such movements is that it simplifies controlling a whip. Whichever procedure is used, manipulating an object may be understood as searching the space of parameters that determine the action. Due to the complexity of whip dynamics, the dimension of the space of possible control actions — e.g., trajectories of the whip handle — is prohibitively large. However, a movement with a pre-determined bell-shaped speed profile may be defined by a small number of parameters: start position, end

¹Department of Mechanical Engineering, Massachusetts Institute of Technology, Cambridge, MA 02139, USA

²Department of Bioengineering, Northeastern University, Boston, MA 02115, USA

³Department of Biology, Northeastern University, Boston, MA 02115, USA

⁴Department of Neurology, Policlinico Tor Vergata and the Laboratory of Neuromotor Physiology, IRCCS Santa Lucia Foundation, Rome, Italy

⁵Department of Biology, Department of Electrical and Computer Engineering, Department of Physics, Institute of Experiential Robotics, Northeastern University, Boston, MA 02115, USA

⁶Department of Brain and Cognitive Sciences, Massachusetts Institute of Technology, Cambridge, MA 02139, USA

⁷Lead contact

*Correspondence: mosesnah@mit.edu

<https://doi.org/10.1016/j.isci.2023.107395>



position, and duration. As a result, provided a small number of submovements is competent to make the whip hit a target, the dimension of the search space would be reduced dramatically.

This study extended previous work by Nah et al.^{12,13} and examined whether discrete movements of the upper arm can generate whip movements that aim to strike targets at different distances and directions. In simulation, we developed a simple arm model with motion elements that employed maximally smooth movements with bell-shaped speed profiles. The spring-like property of the human's neuromuscular system was modeled as constant joint-space mechanical impedances, where the values were derived from experimental measurement. We used an optimization algorithm to find the optimal parameters for the joint motions to strike the targets.

The following five specific questions were addressed:

1. Are maximally smooth discrete movements (with a bell-shaped speed profile) sufficient to reach different target locations?
2. How many search iterations are required to find an optimal movement that achieved the targeting task?
3. How sensitive is performance to the optimal movement parameters?
4. How planar is the optimal hand movement in task-space coordinates?
5. How are the limb's mechanical impedance affected by performance?

RESULTS

To simulate the control of a whip striking a target, a multi-joint model of the human upper limb with four controllable DOF was coupled to a discretized model of a whip with 50 DOF. The upper limb movements were controlled by a joint-space impedance controller: constant joint-space impedances were connected to a virtual trajectory planned in joint-space coordinates.¹² The virtual trajectory was a minimum-jerk motion, which modeled the bell-shaped speed profile commonly observed in discrete human movements.⁹ This enabled the movement to be defined by 9 parameters: 1 for movement duration, 4 for the starting (virtual) joint posture, and 4 for the ending (virtual) joint posture. The performance of the movement was evaluated by the closest distance that the whip reached with respect to the target. A robust optimization procedure¹⁴ was used to identify discrete movements that could strike targets at different locations.

In overview, the results showed that a single discrete (maximum-smoothness) movement in joint space, defined by 9 parameters, was able to make the whip strike or closely approach each of 6 different targets (Figure 1). The controller was able to strike 5 out of 6 targets. For the target located at the highest and most distant location, the closest approach of the whip to the target was 10 cm. Remarkably, no more than 250 iterations (and as few as 39 iterations for target 3) were required to identify a movement that achieved the task (Figure 2; Table 1). Systematic similarities of the movements to different targets were observed. For all targets, the optimal hand movements were mostly planar in task-space coordinates. The optimal movements were predominantly in the most compliant direction of both joint-space and task-space mechanical impedances. Of the optimal movement parameters, striking the target was most sensitive to movement duration.

Hitting a target

Figure 1 illustrates the placement of the 6 targets with respect to the upper limb and whip model. Targets 1 to 3 were located well within reach of the extended whip, overlapping by roughly a quarter of the whip's length. Targets 4 to 6 were located at a distance equal to the combined length of the fully extended upper-limb and whip. Hence, targets 4 to 6 could only be hit by striking precisely with the tip of the whip, requiring greater precision.

Despite the high dimensionality of the model's state-space — a 54-DOF system, yielding a 108-dimensional state-space representation — the DIRECT-L algorithm successfully identified a single discrete movement that was able to strike or closely approach the target. Figure 2 shows the optimization results for minimum distance between whip and target. For targets 1 to 5, the optimized movement managed to successfully strike the targets. Remarkably few iterations were required — fewer than 100

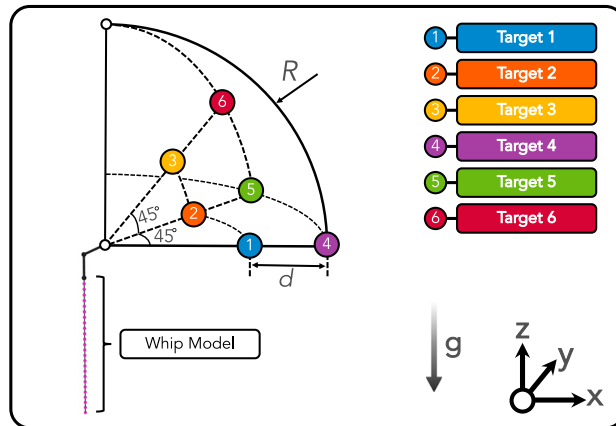


Figure 1. Target locations

Graphical depiction of the six target positions and the task-space coordinate frame of the simulation. Distance R , which is the radius of a sphere centered at the shoulder of the upper limb (depicted as a white marker), is equal to the sum of the lengths of the upper limb and whip (i.e., $R = L_1 + L_2 + l \cdot N = 2.385\text{m}$) (Table 3). The three nearby targets are separated from the three farther-away targets by a constant radial distance d (0.4m). Target locations are shown in a spherical coordinate system (radius-azimuth-elevation): Target 1: $(R - d, 0^\circ, 0^\circ)$, Target 2: $(R - d, 45^\circ, 0^\circ)$, Target 3: $(R - d, 45^\circ, 45^\circ)$, Target 4: $(R, 0^\circ, 0^\circ)$, Target 5: $(R, 45^\circ, 0^\circ)$, Target 6: $(R, 45^\circ, 45^\circ)$.

iterations for targets 1 to 3 that were well within reach; fewer than 250 iterations for all other targets. Even for the most difficult target 6, at the highest and most distant location, the closest distance between the whip and target was 10cm and was identified within 250 iterations.

Characteristics of the optimal movements

The optimal movements were qualitatively and quantitatively similar for targets in the same direction, i.e., targets 1 and 4, or targets 2 and 5, or targets 3 and 6 (Figure 1). For targets 1 and 4, which were located in the frontal direction, the optimal trajectory was a large nearly planar ‘sweeping’ movement of handle and whip aligned with a sagittal plane (Figures 3A and 4A). The measure of average deviation from a planar surface, planar surface deviation (PSD) (Table 2), was smaller than for other targets. For targets 1 and 4, the average deviation from the best-fit plane was approximately 1cm. The normal vectors of the best-fit planes almost co-aligned with the normal vector of the xz-plane, i.e., a sagittal plane (Figures 5A and 6A). For




Figure 2. Result of the optimization

Minimum distance between the whip and target L^* over iterations for targets 1 to 3 (top panel) and 4 to 6 (bottom panel). For the optimization of target 6, the occasional spikes are due to the DIRECT-L algorithm’s procedure.^{12,13}

Table 1. Optimal movement parameters and L^* , the closest distance between the whip and target

Target	Optimal movement parameters								D	L^* (m)
	$q_{0,i,J1}$	$q_{0,i,J2}$	$q_{0,i,J3}$	$q_{0,i,J4}$	$q_{0,f,J1}$	$q_{0,f,J2}$	$q_{0,f,J3}$	$q_{0,f,J4}$		
1	- 1.361	0.000	- 0.349	1.414	1.728	0.000	0.000	1.414	0.950	0.000
2	- 0.943	0.000	- 1.047	1.414	1.728	- 1.047	0.000	0.471	0.950	0.000
3	- 0.943	1.047	0.000	1.414	1.728	- 1.047	0.000	1.414	0.583	0.000
4	- 1.548	0.000	- 0.349	1.414	1.728	0.000	0.000	0.367	0.950	0.000
5	- 1.082	1.047	1.047	0.820	1.728	- 1.047	1.396	0.157	0.950	0.000
6	- 0.943	1.060	0.349	0.995	1.740	- 0.970	- 0.272	0.948	0.525	0.102

$L^* = 0$ denotes that the whip hit the target. Subscript J1~J4 correspond to the joints of the upper-limb model (Figure 11). Units of q_0 and D are radian and second, respectively.

targets 1 and 4, the angles between the normal vectors of the xz-plane and the best-fit planes were approximately 10° (Table 2).

Targets 2 and 5 were rotated 45° about a vertical axis through the shoulder from targets 1 and 4, and the optimal movements were qualitatively and quantitatively different (Figures 3B and 4B). The PSD values for targets 2 and 5 were approximately 6cm, which was larger than for other targets. The best-fit planes were strongly tilted in the roll direction (Figures 5B and 6B). For targets 2 and 5, the angles between the normal vectors of the xz-plane and the best-fit planes were approximately 45° .

Targets 3 and 6 were elevated 45° from targets 2 and 5. The optimal hand movements for targets 3 and 6 were more planar than for targets 2 and 5 and closer to a large planar sweeping motion (Figures 3C and 4C). For targets 3 and 6, the PSD values were 3.9cm and 2.6cm, respectively. These values were between those of targets 1 and 4 and targets 2 and 5 (Table 2). While the movements for targets 3 and 6 were largely planar, the best-fit planes were strongly tilted in the roll direction (Figures 5C and 6C). As with targets 2 and 5, for targets 3 and 6, the angles between the normal vectors of the xz-plane and the best-fit planes were approximately 45° .

Despite these quantitative and qualitative differences, the average deviation from the planar movement was relatively small — maximum 6.5cm, which is 10% of the length of the upper limb (Table 2). Hence, for all targets, the resulting hand movements that achieved the targeting task were mostly planar in task-space coordinates.

Sensitivity of the optimal movement parameters

The sensitivity of successful hits to deviations from the 9 optimal movement parameters was assessed. The analysis was conducted for targets 1 to 5, since only for these targets were optimal parameters found to hit the target. Of the 9 movement parameters for these five targets, striking the target was most sensitive to movement duration (Figure 7). Even an error as small as 10ms (targets 4, 5) or 40ms (target 3) from the optimal movement duration failed to hit the target. This result emphasizes the highly nonlinear and complex dynamics of the whip, where even a negligibly small error from the optimal movement results in a failure to hit a target.

Relation of the optimal movements to limb mechanical impedance

The qualitative differences in the movements toward different targets might be related to the joint-space mechanical impedance of the upper-limb model. To explore that possibility, the directions corresponding to the eigenvectors of the 4-dimensional joint-space stiffness were computed and ranked in ascending order of the corresponding eigenvalues. A consistent pattern emerged: regardless of target location, the contribution of the first joint-space eigenvector (i.e., eigenvector corresponding to the smallest eigenvalue) was larger than the contributions of all other eigenvectors (Figures 5A, 6A, and 8). Qualitatively, the optimal movements were mostly oriented along the most compliant direction of the joint-space impedance matrices. This pattern was also observed in task-space mechanical impedances, where the impedances were naturally induced by the joint-space impedances (Figures 8, 9, and 10). This result indicates a strong relation between mechanical impedances and task performance, where the optimal movement was predominantly oriented along the most compliant direction of both task-space and joint-space mechanical impedances.

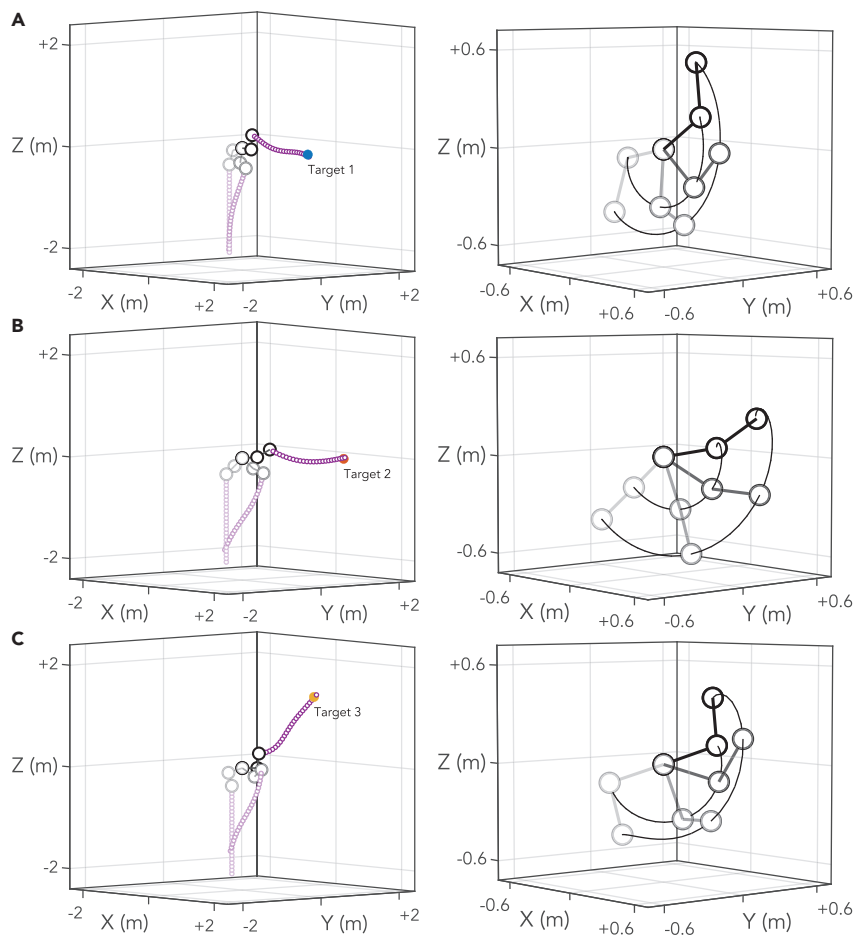


Figure 3. Simulation results for targets 1, 2, and 3, shown in panels A, B, and C, respectively (Left column) Time-sequence of upper-limb and whip model. Three frames of the simulation were taken: at the start of the movement, at the moment when the whip hit the target, and a moment between the two. Black circles depict the shoulder, elbow and hand of the upper-limb model. Purple circles depict the point-masses of the whip model. Opacity of the color increases from the start to the end of the movement. (Right Column) Time-sequence of the corresponding optimal trajectories of elbow and end-effector. Four frames are plotted from the start to the end of the movement, i.e., for time $0 \leq t \leq D$, where D is the duration of the virtual (zero-torque) trajectory (Equation 3).

DISCUSSION

Manipulating soft, flexible, deformable objects poses significant challenges. Unlike the control of rigid objects, their infinite-dimensional structure exacerbates the complexity of the control problem.^{15,16} These challenges are highlighted by the shortcomings of popular optimization-based approaches — they scale poorly with system dimension. Known as the “curse of dimensionality”,¹⁷ the complexity of optimization grows exponentially with the number of DOF and with the number of time-samples included in the computation. Even with modern computational resources, real-time optimization of Lagrangian dynamic systems with as few as tens of DOF is still profoundly challenging.¹⁸ Despite these challenges, humans excel at manipulating soft, flexible, deformable objects as demonstrated in a recent experimental study by our group.⁵ Complementary to that experimental study that analyzed data of the human hand and whip, the current study simulated, i.e., synthesized, a similar action.

This simulation study examined whether targets at various locations could be reached with a simple model of the arm controlled by simple control components, or ‘primitives’.¹⁹ Specifically, we tested whether a whip could be manipulated with a minimum-jerk movement whose parameters were learned through optimization. This required coordination of four actively controlled joints to control a system with 108 state-space dimensions. Assuming a coarse sampling rate of 4Hz (once every 250ms), the corresponding Nyquist

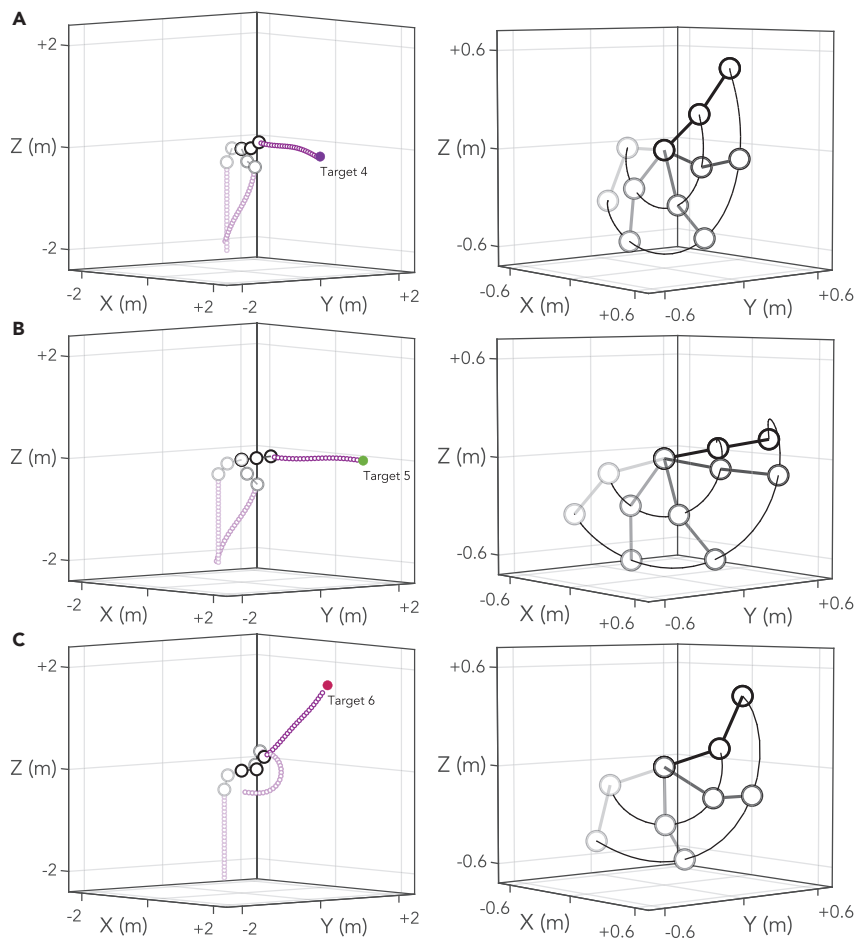


Figure 4. Simulation results for targets 4, 5, and 6, shown in panels A, B, and C, respectively

(Left column) Time-sequence of upper-limb and whip model. Three frames of the simulation were taken: at the start of the movement, at the moment when the whip hit the target (for target 6, at the moment when the whip and the target were closest), and a moment between the two. Black circles depict the shoulder, elbow and hand of the upper-limb model. Purple circles depict the point-masses of the whip model. Opacity of the color increases from the start to the end of the movement. (Right Column) Time-sequence of the corresponding optimal trajectories of elbow and end-effector. Several frames are plotted from the start to the end of the movement, i.e., for time $0 \leq t \leq D$, where D is the duration of the virtual (zero-torque) trajectory (Equation 3).

frequency of 2Hz is just greater than 1.79Hz, which is the fastest oscillatory mode of the whip model (details of the modal analysis of the linearized whip model are presented in the Supplementary Materials "Modal Analysis of the Linearized Whip Model"). With such sampling rate, a 1-s-long simulation of a 108-dimensional system lies in a space with $4^{108} \approx 10^{65}$ sample points. A complete search of this space would be profoundly impractical. However, encoding upper-limb action using the parameters of discrete smooth movements simplified the optimization. Our results showed that the optimization successfully managed the complexity of this extremely high-dimensional problem, and thereby escaped the curse of dimensionality.

With reference to the five questions raised in the introduction, the results to the corresponding questions can be summarized as follows:

1. A single discrete movement planned in joint-space coordinates was sufficient to reach various target locations.
2. The proposed method was able to obtain the optimal solution in 250 iterations or fewer, with a maximum of 249 iterations for target 5 and a minimum of 39 iterations for target 3 (Figure 2). This method did not require sophisticated computational resources, nor any collection of human data for demonstration.

Table 2. Planar surface deviation (PSD), the normal vector n and centroid c of the best-fit plane of the optimal hand movements

Target	PSD (m)	Normal Vector n (m)	Centroid c (m)
1	0.012	[- 0.002, - 0.976, 0.220]	[0.076, - 0.050, - 0.102]
2	0.059	[- 0.426, 0.711, - 0.560]	[0.082, 0.065, - 0.197]
3	0.039	[- 0.361, 0.744, - 0.562]	[0.181, - 0.058, - 0.154]
4	0.014	[0.085, - 0.984, 0.158]	[0.090, - 0.059, - 0.183]
5	0.065	[0.686, - 0.689, 0.235]	[0.011, 0.062, - 0.244]
6	0.026	[0.566, - 0.748, - 0.347]	[0.156, - 0.097, - 0.247]

The movements were generated from the movement parameters of Table 1.

3. Within the 9 optimal movement parameters, striking the target was most sensitive to movement duration (Figure 7). Even a small deviation (10ms–40ms) from the optimal duration resulted in failure to hit the target.
4. For all targets, the optimal hand movement was essentially planar.

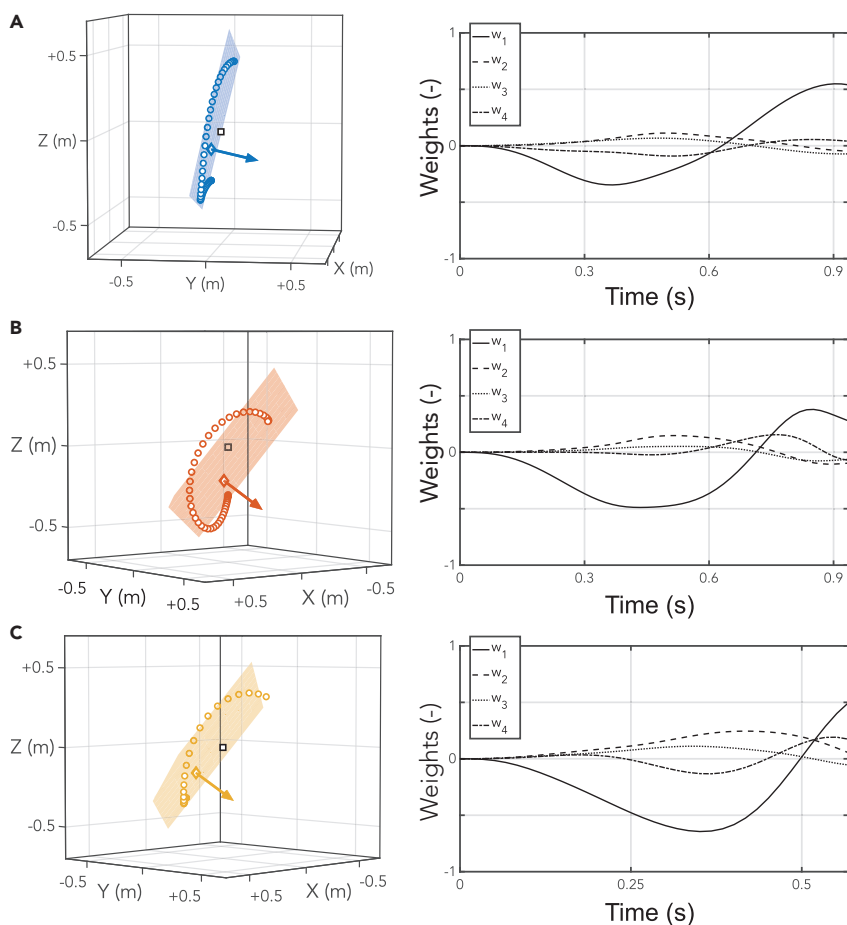


Figure 5. Optimal end-effector movement for targets 1, 2, and 3, shown in panels A, B, and C, respectively (Left Column) The end-effector trajectory (Figure 3, right column) and its best-fit plane. The centroid c and the normal vector n of the best-fit plane are depicted respectively as a diamond symbol and an arrow originating from it. The square depicts the shoulder's location. (Right Column) Corresponding weights of the eigenmovements, $w_i, i = 1, \dots, 4$ of the optimal upper-limb movement vs. time (Equation 6). Values between the start and end of the movement are plotted, i.e., the width of the x-axis is D , the duration of the virtual (zero-torque) trajectory (Equation 3).

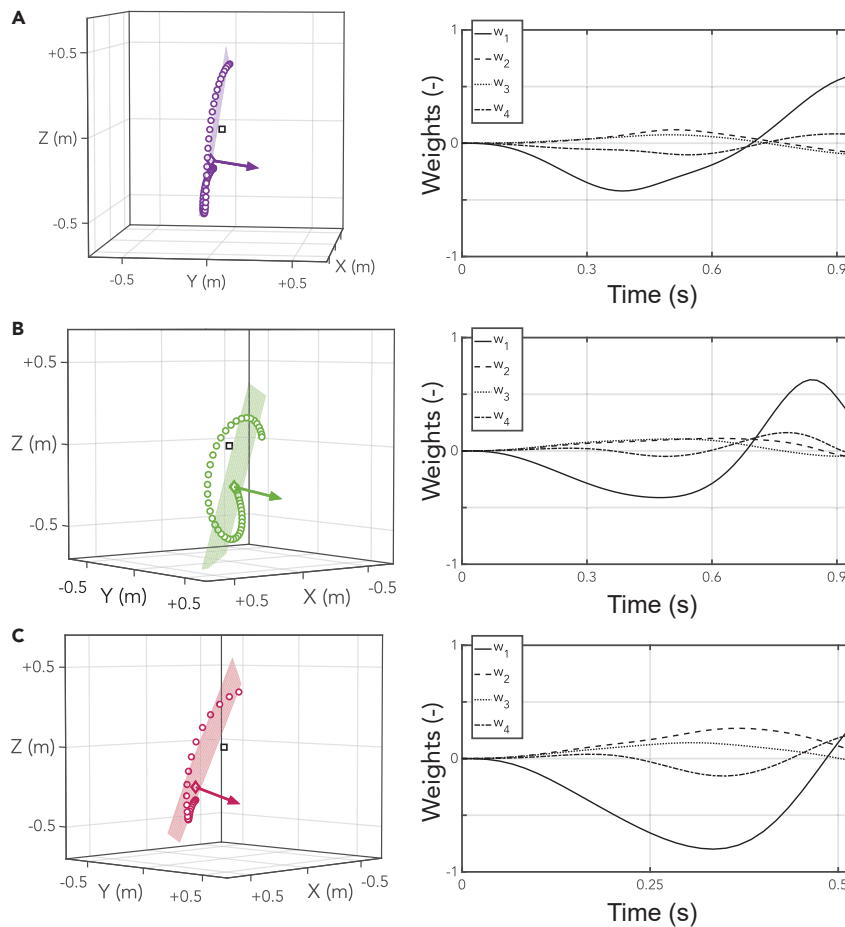


Figure 6. Optimal end-effector movement for targets 4, 5, and 6, shown in panels A, B, and C, respectively (Left Column) The end-effector trajectory (Figure 4, right column) and its best-fit plane. The centroid c and the normal vector n of the best-fit plane are depicted respectively as a diamond symbol and an arrow originating from it. The square depicts the shoulder's location. (Right Column) Corresponding weights of the eigenmovements, $w_i, i = 1, \dots, 4$ of the optimal upper-limb movement vs. time (Equation 6). Values between the start and end of the movement are plotted, i.e., the width of the x-axis is D , the duration of the virtual (zero-torque) trajectory (Equation 3).

5. The optimal movements were predominantly in the most compliant direction, both in intrinsic (joint-space) and extrinsic (task-space) coordinates.

Compared to state-of-the-art research,^{20–22} the two primary contributions of this work are:

- *Computational and data efficiency of the approach.* In contrast to modern learning approaches,^{20,21} the dynamic manipulation tasks were achieved without requiring a large amount of pre-collected data. An optimal solution was found within 250 iterations, which is significantly fewer than these state-of-the-art methods.
- *Insensitivity to the dimensionality or complexity of the manipulated object.* In contrast to the trajectory optimization methods, which are influenced by the dimensionality of the object,²² our approach was unaffected by such considerations. Given a well-defined objective function (e.g., the minimum distance between the tip of the whip and target), our approach can successfully manipulate objects with high dimensionality and complexity.




Figure 7. Sensitivity of task performance to the movement parameters for targets 1 to 5
 Gray bars depict the upper and lower bounds of the search space of the DIRECT-L optimization algorithm. The upper (respectively lower) error bar depicts the upper (respectively lower) value of the error to miss the target. Dots within the error bar depict the optimal movement value (Table 1). T1 to T5 denote target 1 to target 5, respectively.

Using primitive actions to learn complex object manipulation

Recent state-of-the-art machine learning algorithms proposed ‘structured policy methods’ to tackle high-dimensional optimization problems.^{23–28} Compared to model-free reinforcement learning methods that are notoriously expensive in terms of their sampling complexity,^{29,30} structured policy methods alleviate computational complexity by first embedding predefined structures in the controller and then training the control policy to learn the optimal parameters of those structures. This approach has been successful in reproducing a rich repertoire of motions, e.g., bicycle stunts³¹ and various locomotion skills for quadrupeds.³² These methods could also generate soft, flexible and deformable object manipulation where a small set of motion primitives enabled robots to manipulate cloths²⁷ and disentangle multi-cable knots.²⁸ Nevertheless, the cloth-manipulation study²⁷ required over 100,000 iterations for their optimization to converge, while the knot-disentangling study²⁸ required 7,000 images to train their controller. Importantly, the choice of structure, or representation, impacted the learning, robustness, and motion quality of the controller, and the best form for the embedded structure has remained unresolved.²⁶

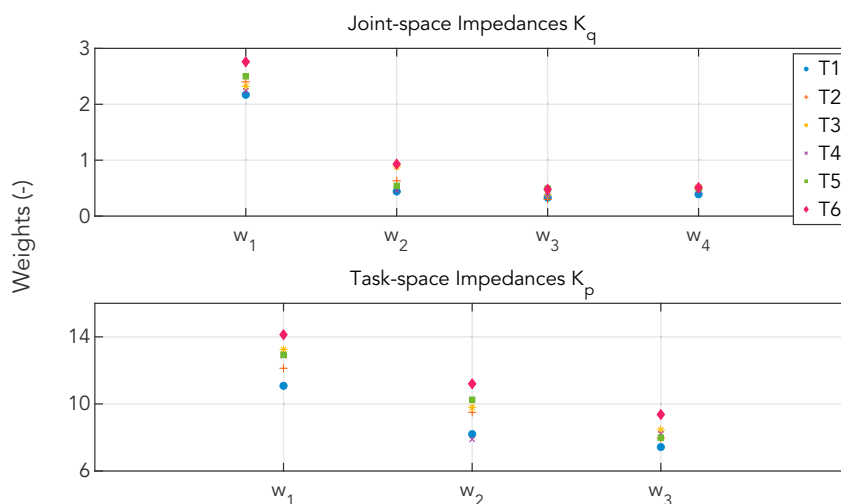


Figure 8. Square root of the sum of the squared weight values of the joint-space and task-space impedances
 T1 to T6 denote target 1 to target 6, respectively.

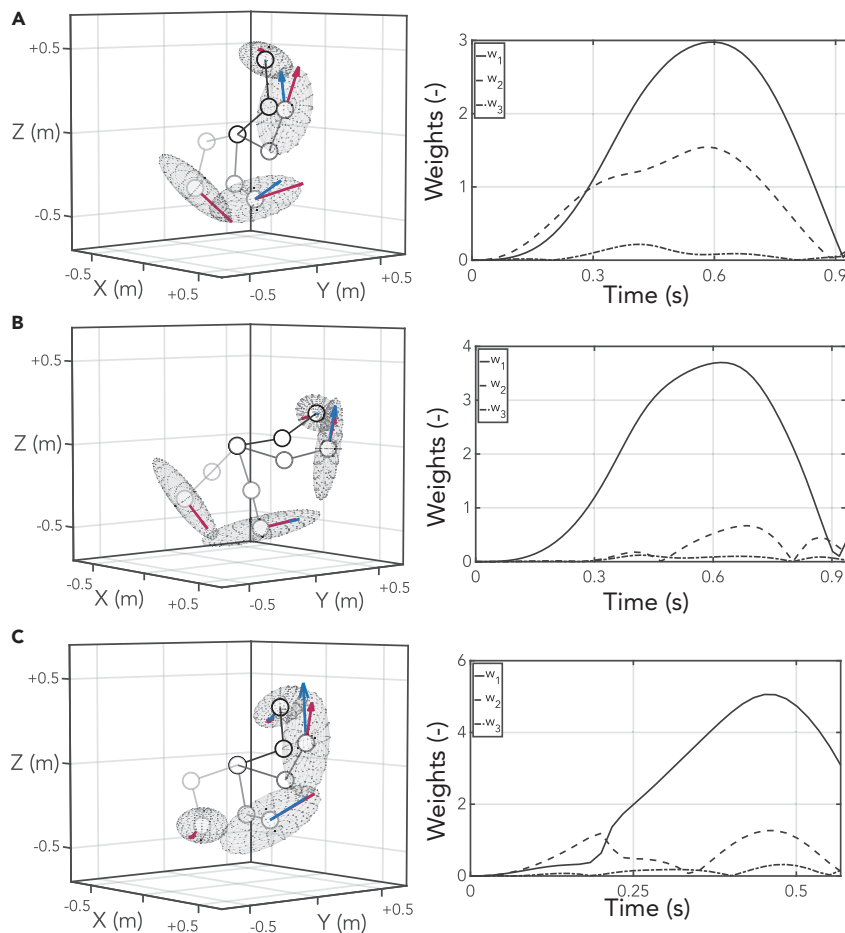


Figure 9. Task-space stiffness ellipsoids for targets 1, 2, and 3, shown in panels A, B, and C, respectively

(Left Column) Time-sequence of upper-limb (black) and its stiffness ellipsoids. Four frames of the simulation were taken: at the start of the movement, at the end of the movement, and frames between the two. The iso-potential energy surface of the task-space impedance K_p is plotted as an ellipsoid, i.e., the width of the ellipsoid is inversely proportional to the eigenvalue of K_p . Red arrow depicts the eigenvector of K_p with the smallest eigenvalue. Blue arrow depicts the velocity of the hand. (Right Column) Corresponding weights of the eigenvectors, w_i , $i = 1, 2, 3$ of the stiffness ellipsoid vs. time (Equation 8). Values between the start and end of the movement are plotted, i.e., the width of the x-axis is D , the duration of the virtual (zero-torque) trajectory (Equation 3).

The work reported here resembles a structured policy method, but the structures were formulated in terms of biologically inspired motor primitives — robust patterns found in prior human neuroscience research.^{7–10,33,34} A feedforward open-loop motion command using a minimum-jerk trajectory with constant joint-space impedance terms, derived from experimental measurements, was sufficient to achieve the task. This approach enabled convergence to the optimal set of movement parameters in less than 250 iterations — a maximum of 249 iterations for target 5, a minimum of 39 iterations for target 3 (Figure 2). This is a dramatically smaller number than state-of-the-art machine learning algorithms which typically require many thousands of iterations.²⁷ For target 6, which was located at the highest and most distant location, the closest distance between the whip and target was 10cm. This error suggests that a single primitive movement may not be sufficient to reach all targets. Whether two or more submovements might suffice is a topic for future studies. It is also important to note that the optimal movement parameters were well within the ranges used by the search algorithm (Tables 1 and 3). Hence, it is unlikely that movements of shorter (or longer) duration and/or range would achieve results better than those reported here.

It is worth emphasizing that this approach was judiciously biomimetic. It was based on the hypothesis that humans achieve their remarkable dexterity by taking advantage of motor primitives. Without any detailed

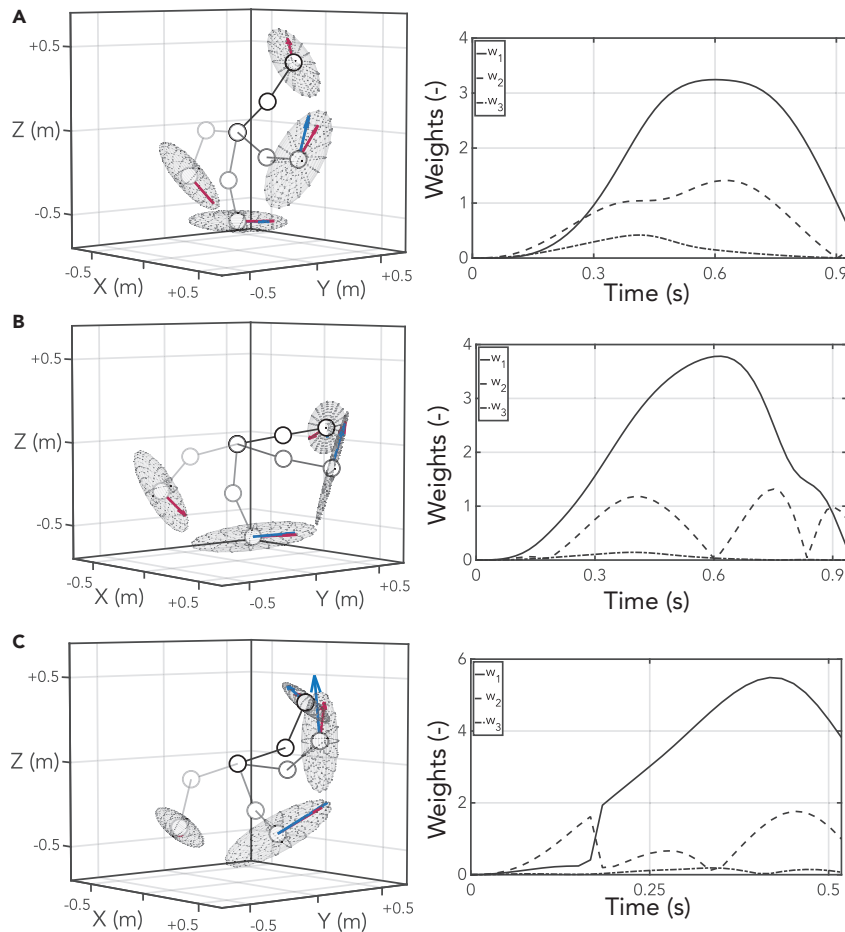


Figure 10. Task-space stiffness ellipsoids for targets 4, 5, and 6, shown in panels A, B, and C, respectively (Left Column) Time-sequence of upper-limb (black) and its stiffness ellipsoids. Four frames of the simulation were taken: at the start of the movement, at the end of the movement, and frames between the two. The iso-potential energy surface of the task-space impedance K_p is plotted as an ellipsoid, i.e., the width of the ellipsoid is inversely proportional to the eigenvalue of K_p . Red arrow depicts the eigenvector of K_p with the smallest eigenvalue. Blue arrow depicts the velocity of the hand. (Right Column) Corresponding weights of the eigenvectors, w_i , $i = 1, 2, 3$ of the stiffness ellipsoid vs. time (Equation 8). Values between the start and end of the movement are plotted, i.e., the width of the x-axis is D , the duration of the virtual (zero-torque) trajectory (Equation 3).

representation of the object being manipulated, the approach was highly effective to manage the complexity of an extremely high-dimensional system. This approach can be classified as ‘model-free’ learning, since by using parameterized motor primitives, it operated without relying on a detailed model of the underlying system dynamics.³⁵ Moreover, our approach is distinct from “adaptive control methods” which pertain to tracking controllers.³⁵ Instead, the minimum-jerk trajectory planned in joint-space coordinates served as a “virtual trajectory” connected to joint stiffness and damping, rather than a “desired trajectory” aimed at achieving perfect tracking. These results suggest that motor primitives are key structures humans use to achieve their remarkable dexterity manipulating high-dimensional or even continuum dynamic objects.

The role of mechanical impedance

In this work, mechanical impedance was included to account for physical interaction between the hand and whip.^{19,36–38} A dynamic operator which determines the force evoked by an imposed displacement, mechanical impedance ‘shapes’ the dynamics of interaction with the object being manipulated. Versatile interaction with objects may be achieved by choosing a specific impedance, especially for tasks in which simply mimicking a kinematic trajectory from a demonstrator may not be sufficient.^{39–42} A striking relation between limb impedance and task performance was discovered: the optimization algorithm converged to

Table 3. Model parameters

	Description	Notations	Values	Unit
Geometric and Inertial Parameters of Limb Segments	Mass of limb segment	M_1, M_2	1.595, 0.869	(kg)
	Length of limb segment	L_1, L_2	0.294, 0.291	(m)
	Length from proximal joint to COM	L_{1c}, L_{2c}	0.129, 0.112	(m)
	Principal moment of inertia, x' - axis	$I_{1,xx}, I_{2,xx}$	0.0119, 0.0048	(kg · m ²)
	Principal moment of inertia, y' - axis	$I_{1,yy}, I_{2,yy}$	0.0119, 0.0049	(kg · m ²)
	Principal moment of inertia, z' - axis	$I_{1,zz}, I_{2,zz}$	0.0013, 0.00047	(kg · m ²)
Parameters of the Whip Model	Number of sub-models (i.e., node number)	N	25	(–)
	Value of point-mass	m	0.012	(kg)
	Length of massless cylinder	l	0.072	(m)
	Coefficient of rotational spring	k	0.242	(N · m / rad)
	Coefficient of rotational damper	b	0.092	(N · m · s / rad)

Subscripts 1 and 2 denote the upper arm and forearm, respectively. Principal moments of inertia of limb segments were calculated with respect to the center of mass (COM) (Figure 11).

movements predominantly in the most compliant direction. This result is reminiscent of previous experimental observations, in which mechanical impedance decreased with practice.^{43,44}

These results indicate that limb mechanical impedance may be an important factor in determining the optimal movement. For all targets, the optimal hand movement was essentially a planar movement. If the dimensionality of the movement is considered as a measure of its complexity, a planar motion (which recruits shoulder and elbow flexion/extension) is arguably the ‘simplest’ motion that can accomplish the task. Moreover, a predominantly planar hand action evoked a predominantly planar whip motion. As a result, the targeting task could be achieved within a 27-dimensional subset of the 54-dimensional configuration space of the whip (25 whip DOF plus 2 limb DOF). Reducing the relevant dimensions of configuration space by half is arguably a substantial simplification.

While the objective of striking a target with a whip was successfully achieved using constant joint-space impedances, the possibility of variable joint-space impedances for control purposes is not precluded by this result. Prior research showed that impedance modulation was necessary to achieve complex manipulation tasks, e.g., managing physical interaction with the environment.^{45,46} In fact, incorporating time-varying joint-space impedances into the controller could potentially broaden the range of targets that can be successfully hit by the whip, e.g., target 6. Investigating the potential advantages of variable joint-space impedances for complex object manipulation is a topic for future studies.

A human-inspired approach

Human sensory-motor capabilities far exceed those of contemporary robots, even though our actuation, communication, and computation are orders of magnitude slower than those of robotic systems. The fastest neural transmission speed in humans is about 100 m/s,^{47,48} about a million times slower than its robotic counterparts which can reach up to about 10⁸ m/s. The bandwidth of skeletal muscle is considerably less than 10Hz,⁴⁷ whereas electromechanical technology routinely achieves bandwidths in excess of hundreds of Hz.⁴⁹

Since the slow neuromuscular system impairs reactive feedback control, human motor control relies heavily on predictive feedforward action. However, a detailed prediction of a whip’s behavior would be quite unworkable for real-time control, even with modern computational resources. A practical alternative is to simplify control by taking advantage of the whip’s natural behavior. Experimental results by our group reported a propagating wave in the whip, with a speed profile strikingly similar to the speed profile of discrete human movements.⁵ Hence, a bell-shaped speed profile as the basic movement shape may be sufficient to manipulate a whip, and the results in this paper have demonstrated that the approach successfully achieved the targeting task.

A human operates the whip via its handle, but the upper limb has more DOF than the handle, so that a specified hand motion does not uniquely determine corresponding joint motions. How the biological controller




Figure 11. Model used for the simulation

(A) Simulation model in its virtual configuration, rendered with the MuJoCo simulator.

(B) Model of the upper limb and its model parameters. Rotational joints of shoulder (J1-J3), elbow (J4) and their axes of rotation are visualized as bullet shapes. Joints J1-J4 were equipped with torque actuators.

(C) Sub-model of the whip and its model parameters. Each rotational joint is equipped with a linear rotational spring k and rotational damper b , visualized as bullet shapes.

manages this ill-posed inverse kinematic map from hand motions to joint motions is presently unknown.³⁶ To focus on the feasibility of managing the whip with pulse-like 'building-block' motions, this problem was circumvented by assuming minimum-jerk motion profiles in joint space. The resulting hand paths are mostly vertically oriented and curved, not unlike the hand paths reported by Atkeson and Hollerbach.⁸

Previous work has argued that human neuromotor performance appears to be a composition of motor primitives.^{50,51} These are defined as patterns of behavior that manifest as stable attractors of the (nonlinear) neuromechanical system. They are conceived as building blocks that can be parametrized and combined to produce complex behavior.⁵²⁻⁵⁴ At least three distinct classes of primitives have been identified — sub-movements (for non-repetitive point-to-point movements), oscillations (for rhythmic, repetitive movements), and mechanical impedance (to manage physical interaction with the environment). Related work in robotics has demonstrated that dynamic movement primitives are a powerful method for planning robot trajectories based on demonstration.⁵⁵⁻⁶⁰ Their effectiveness has been proven in a variety of domains, including biped locomotion,⁶¹ placing and pouring⁶² and dart throwing.⁶³ Nevertheless, to the best of our knowledge, the main focus of this body of research was on unconstrained movements or on the manipulation of rigid objects with comparatively few DOF. Tasks as dynamically complex as manipulating a whip, as reported in this paper, have not yet been fully explored.

Task performance and motor precision

This study simulated a complex manipulation task where the goal was to hit targets at several locations with a 50-DOF whip model. Considering the nonlinear dynamics of such a high-dimensional system, deviations from the optimal movement were expected to result in large errors, and they did. Assuming control based on dynamic motor primitives, learning complex motor skills would consist of finding the optimal parameters of the 'building blocks'. Once these optimal parameters were learned, initiating the optimal action

would consist of retrieving those parameters. However, the motor variability of humans — natural variations in posture, movement and muscle activity of human motor control^{64,65} — prevents perfect repetition of the optimal action. Even skilled movements of experts showed variability across repetitions, discussed by Bernstein in his famous blacksmith example and known as “repetition without repetition”.⁶⁶ Our sensitivity analysis considered the variability of human motor control by quantifying the effect of deviations from the optimal movement parameters (Table 3). The results emphasized the challenge of accurately striking a target. In fact, a whip expert who was able to hit a distant target with a success rate of only about 90% had more than 15 years of experience.⁴

Limitations of the study

The current study focused on complex object manipulation in simulation without actual robot implementation. Future studies are necessary to extend the study from simulation to actual robot experiments.

Conclusions

Despite the significant limitations of our neuromuscular system, humans manipulate flexible, deformable objects of prodigious dynamic complexity with apparent ease. The simulations presented here showed that encoding action based on the parameters of motor primitives enabled optimization to control a flexible object with extremely complex dynamics — a whip. Robotic manipulation of flexible materials continues to pose a significant challenge. Control based on primitives may provide a novel and effective way to control soft, flexible, dynamically complex objects.

STAR★METHODS

Detailed methods are provided in the online version of this paper and include the following:

- RESOURCE AVAILABILITY
 - Lead contact
 - Materials availability
 - Data and code availability
- METHOD DETAILS
 - Modeling
 - Upper-limb controller
 - Evaluation metrics of the upper-limb movements
 - Optimization of the whip task
 - Sensitivity analysis of the optimal movement parameters

SUPPLEMENTAL INFORMATION

Supplemental information can be found online at <https://doi.org/10.1016/j.isci.2023.107395>.

ACKNOWLEDGMENTS

The research was supported by NSF M3X-1826097 (Neville Hogan), NSF M3X-1825942 (Dagmar Sternad), NIH-R37-HD087089 (Dagmar Sternad), and NIH-R01-HD087089.

AUTHOR CONTRIBUTIONS

M.N. performed simulations, software implementation, and contributed to figures and writing. A.K. collected the data to identify the model parameters of the whip and contributed to figures and writing. M.R. collected the data to identify the model parameters of the whip and contributed to figures and writing. D.S. conceived the study, contributed to writing, funding acquisition, and supervision of the project. N.H. conceived the study, contributed to writing, funding acquisition, and supervision of the project. All authors read and approved the final manuscript.

DECLARATION OF INTERESTS

The authors declare no competing interests.

INCLUSION AND DIVERSITY

We support inclusive, diverse, and equitable conduct of research.

Received: December 19, 2022

Revised: May 26, 2023

Accepted: July 11, 2023

Published: July 14, 2023

REFERENCES

- Bernstein, B., Hall, D.A., and Trent, H.M. (1958). On the dynamics of a bull whip. *J. Acoust. Soc. Am.* 30, 691–1115.
- Goriely, A., and McMillen, T. (2002). Shape of a cracking whip. *Phys. Rev. Lett.* 88, 244301.
- McMillen, T., and Goriely, A. (2003). Whip waves. *Phys. Nonlinear Phenom.* 184, 192–225.
- Henrot, C.C.I. (2016). Characterization of Whip Targeting Kinematics in Discrete and Rhythmic Tasks (Massachusetts Institute of Technology). Bachelor's Thesis.
- Krotov, A., Russo, M., Nah, M., Hogan, N., and Sternad, D. (2022). Motor control beyond reach—how humans hit a target with a whip. *R. Soc. Open Sci.* 9, 220581.
- Russo, M., Ozeri-Engelhard, N., Hupfeld, K., Nettekoven, C., Thibault, S., Sedaghat-Nejad, E., Buchwald, D., Xing, D., Zobeiri, O., Kiltner, K., et al. (2021). Highlights from the 30th annual meeting of the society for the neural control of movement. *J. Neurophysiol.* 126, 967–975.
- Gordon, J., Ghilardi, M.F., and Ghez, C. (1995). Impairments of reaching movements in patients without proprioception: I. Spatial errors. *J. Neurophysiol.* 73, 347–360.
- Atkeson, C.G., and Hollerbach, J.M. (1985). Kinematic features of unrestrained vertical arm movements. *J. Neurosci.* 5, 2318–2330.
- Flash, T., and Hogan, N. (1985). The coordination of arm movements: an experimentally confirmed mathematical model. *J. Neurosci.* 5, 1688–1703.
- Hogan, N., and Sternad, D. (2012). Dynamic primitives of motor behavior. *Biol. Cybern.* 106, 727–739.
- Krebs, H.I., Aisen, M.L., Volpe, B.T., and Hogan, N. (1999). Quantization of continuous arm movements in humans with brain injury. *Proc. Natl. Acad. Sci. USA* 96, 4645–4649.
- Nah, M.C., Krotov, A., Russo, M., Sternad, D., and Hogan, N. (2021). Manipulating a whip in 3D via dynamic primitives. In 2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), pp. 2803–2808.
- Nah, M.C., Krotov, A., Russo, M., Sternad, D., and Hogan, N. (2020). Dynamic Primitives Facilitate Manipulating a Whip. 2020 8th IEEE RAS/EMBS International Conference for Biomedical Robotics and Biomechanics (BioRob), pp. 685–691.
- Gablonsky, J.M., and Kelley, C.T. (2001). A locally-biased form of the DIRECT algorithm. *J. Global Optim.* 21, 27–37.
- Sadati, S.M.H., Maiolino, P., Iida, F., Nanayakkara, T., and Hauser, H. (2020). Current advances in soft robotics: best papers from RoboSoft 2018. *Front. Robot. AI* 7, 56.
- Rus, D., and Tolley, M.T. (2015). Design, fabrication and control of soft robots. *Nature* 521, 467–475.
- Bellman, R. (1958). Dynamic programming and stochastic control processes. *Inf. Control* 1, 228–239.
- Kuindersma, S., Deits, R., Fallon, M., Valenzuela, A., Dai, H., Permenter, F., Koolen, T., Marion, P., and Tedrake, R. (2016). Optimization-based locomotion planning, estimation, and control design for the Atlas humanoid robot. *Aut. Robots* 40, 429–455.
- Hogan, N., and Sternad, D. (2013). Dynamic primitives in the control of locomotion. *Front. Comput. Neurosci.* 7, 71.
- Zhang, H., Ichnowski, J., Seita, D., Wang, J., Huang, H., and Goldberg, K. (2021). Robots of the lost arc: Self-supervised learning to dynamically manipulate fixed-endpoint cables. In 2021 IEEE International Conference on Robotics and Automation (ICRA), pp. 4560–4567.
- Chi, C., Burchfiel, B., Cousineau, E., Feng, S., and Song, S. (2022). Iterative residual policy: for goal-conditioned dynamic manipulation of deformable objects. *Proceedings of Robotics: Science and Systems (RSS)*.
- Zimmermann, S., Poranne, R., and Coros, S. (2021). Dynamic manipulation of deformable objects with implicit integration. *IEEE Rob. Autom. Lett.* 6, 4209–4216.
- Peng, X.B., Berseth, G., and Van de Panne, M. (2016). Terrain-adaptive locomotion skills using deep reinforcement learning. *ACM Trans. Graph.* 35, 1–12.
- Peng, X.B., Berseth, G., Yin, K., and Van De Panne, M. (2017). Deeploco: dynamic locomotion skills using hierarchical deep reinforcement learning. *ACM Trans. Graph.* 36, 1–13.
- Schaal, S. (1999). Is imitation learning the route to humanoid robots? *Trends Cognit. Sci.* 3, 233–242.
- Peng, X.B., and van de Panne, M. (2017). Learning locomotion skills using deeprl: does the choice of action space matter? In *Proceedings of the ACM SIGGRAPH/Eurographics Symposium on Computer Animation*, pp. 1–13.
- Ha, H., and Song, S. (2022). Flingbot: The unreasonable effectiveness of dynamic manipulation for cloth unfolding. In *Conference on Robot Learning (PMLR)*, pp. 24–33.
- Viswanath, V., Grannen, J., Sundaresan, P., Thananjeyan, B., Balakrishna, A., Novoseller, E., Ichnowski, J., Laskey, M., Gonzalez, J.E., and Goldberg, K. (2021). Disentangling dense multi-cable knots. In 2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS).
- Haarnoja, T., Zhou, A., Abbeel, P., and Levine, S. (2018). Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. In *International Conference on Machine Learning (PMLR)*, pp. 1861–1870.
- Lin, X., Wang, Y., Olkin, J., and Held, D. (2020). SoftGym: Benchmarking deep reinforcement learning for deformable object manipulation. In *Conference on Robot Learning (PMLR)*, pp. 432–448.
- Tan, J., Gu, Y., Liu, C.K., and Turk, G. (2014). Learning bicycle stunts. *ACM Trans. Graph.* 33, 1–12.
- Coros, S., Karpathy, A., Jones, B., Reveret, L., and Van De Panne, M. (2011). Locomotion skills for simulated quadrupeds. *ACM Trans. Graph.* 30, 1–12.
- Morasso, P. (1981). Spatial control of arm movements. *Exp. Brain Res.* 42, 223–227.
- Park, S.-W., Marino, H., Charles, S.K., Sternad, D., and Hogan, N. (2017). Moving slowly is hard for humans: limitations of dynamic primitives. *J. Neurophysiol.* 118, 69–83.
- Schaal, S., and Atkeson, C.G. (2010). Learning control in robotics. *IEEE Robot. Autom. Mag.* 17, 20–29.
- Hogan, N. (2017). Physical interaction via dynamic primitives. In *Geometric and Numerical Foundations of Movements (Springer)*, pp. 269–299.
- Razavian, R.S., Bazzi, S., Nayeem, R., Sadeghi, M., and Sternad, D. (2021). Dynamic primitives and optimal feedback control for the manipulation of complex objects. In *IEEE International Conference on Robotics and Automation (ICRA)*.
- Razavian, R.S., Sadeghi, M., Bazzi, S., Nayeem, R., and Sternad, D. (2023). Body mechanics, optimality, and sensory feedback in the human control of complex objects. *Neural Comput.* 35, 853–895.
- Atkeson, C.G., and Schaal, S. (1997). Robot learning from demonstration. In *International Conference on Machine Learning (ICML (Citeseer))*, pp. 12–20.

40. Atkeson, C.G., and Schaal, S. (1997). Learning tasks from a single demonstration. In *Proceedings of International Conference on Robotics and Automation (IEEE)*, pp. 1706–1712.
41. Li, M., Yin, H., Tahara, K., and Billard, A. (2014). Learning object-level impedance control for robust grasping and dexterous manipulation. In *2014 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 6784–6791.
42. West, A.M., Hermus, J., Huber, M.E., Maurice, P., Sternad, D., and Hogan, N. (2022). Dynamic primitives limit human force regulation during motion. *IEEE Rob. Autom. Lett.* 7, 2391–2398.
43. Milner, T.E., and Franklin, D.W. (2005). Impedance control and internal model use during the initial stage of adaptation to novel dynamics in humans. *J. Physiol.* 567, 651–664.
44. Thoroughman, K.A., and Shadmehr, R. (1999). Electromyographic correlates of learning an internal model of reaching movements. *J. Neurosci.* 19, 8573–8588.
45. Lachner, J., Allmendinger, F., Stramigioli, S., and Hogan, N. (2022). Shaping Impedances to Comply with Constrained Task Dynamics. *IEEE Trans. Robot.* 2750–2767.
46. Buchli, J., Stulp, F., Theodorou, E., and Schaal, S. (2011). Learning variable impedance control. *Int. J. Robot Res.* 30, 820–833.
47. Kandel, E.R., Schwartz, J.H., Jessell, T.M., Siegelbaum, S., and Hudspeth, A.J. (2000). *Principles of Neural Science* (McGraw-hill).
48. Sperelakis, N. (2012). *Cell Physiology Source Book: Essentials of Membrane Biophysics* (Elsevier).
49. Paine, N., Oh, S., and Sentis, L. (2014). Design and control considerations for high-performance series elastic actuators. *IEEE ASME Trans. Mechatron.* 19, 1080–1091.
50. Hogan, N. (2014). A general actuator model based on nonlinear equivalent networks. *IEEE ASME Trans. Mechatron.* 19, 1929–1939.
51. Schaal, S., and Sternad, D. (1998). Programmable pattern generators. In *3rd International Conference on Computational Intelligence in Neuroscience*, pp. 48–51.
52. de Rugy, A., and Sternad, D. (2003). Interaction between discrete and rhythmic movements: reaction time and phase of discrete movement initiation during oscillatory movements. *Brain Res.* 994, 160–174.
53. Sternad, D., Dean, W.J., and Schaal, S. (2000). Interaction of rhythmic and discrete pattern generators in single-joint movements. *Hum. Mov. Sci.* 19, 627–664.
54. Sternad, D. (2008). Towards a unified theory of rhythmic and discrete movements—behavioral, modeling and imaging results. In *Coordination: Neural, Behavioral and Social Dynamics* (Springer), pp. 105–133.
55. Schaal, S., Peters, J., Nakanishi, J., and Ijspeert, A. (2003). Control, planning, learning, and imitation with dynamic movement primitives. In *Workshop on Bilateral Paradigms on Humans and Humanoids: IEEE International Conference on Intelligent Robots and Systems (IROS)*, pp. 1–21.
56. Schaal, S., Peters, J., Nakanishi, J., and Ijspeert, A. (2005). Learning movement primitives. In *Robotics Research, the 11th International Symposium* (Springer), pp. 561–572.
57. Schaal, S. (2006). Dynamic movement primitives—a framework for motor control in humans and humanoid robotics. In *Adaptive Motion of Animals and Machines* (Springer), pp. 261–280.
58. Ijspeert, A.J., Nakanishi, J., Hoffmann, H., Pastor, P., and Schaal, S. (2013). Dynamical movement primitives: learning attractor models for motor behaviors. *Neural Comput.* 25, 328–373.
59. Colomé, A., and Torras, C. (2014). Dimensionality Reduction and Motion Coordination in Learning Trajectories with Dynamic Movement Primitives (2014 IEEE/RSJ International Conference on Intelligent Robots and Systems), pp. 1414–1420.
60. Schaal, S., Kotosaka, S., and Sternad, D. (2000). Nonlinear dynamical systems as movement primitives. *IEEE International Conference on Humanoid Robotics*, 1–11.
61. Nakanishi, J., Morimoto, J., Endo, G., Cheng, G., Schaal, S., and Kawato, M. (2004). Learning from demonstration and adaptation of biped locomotion. *Robot. Autonom. Syst.* 47, 79–91.
62. Pastor, P., Hoffmann, H., Asfour, T., and Schaal, S. (2009). Learning and generalization of motor skills by learning from demonstration. *2009 IEEE International Conference on Robotics and Automation*, pp. 763–768.
63. Kober, J., Oztop, E., and Peters, J. (2011). Reinforcement learning to adjust robot movements to new situations. In *22th International Joint Conference on Artificial Intelligence*.
64. Srinivasan, D., and Mathiassen, S.E. (2012). Motor variability—an important issue in occupational life. *Work* 41, 2527–2534.
65. Sternad, D. (2018). It's not (only) the mean that matters: variability, noise and exploration in skill learning. *Curr. Opin. Behav. Sci.* 20, 183–195.
66. Bernstein, N. (1967). *The Co-ordination and Regulation of Movements* (Pergamon Press).
67. Todorov, E., Erez, T., and Tassa, Y. (2012). Mujoco: A Physics Engine for Model-Based Control. *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 5026–5033.
68. Hatze, H. (1980). A mathematical model for the computational determination of parameter values of anthropomorphic segments. *J. Biomech.* 13, 833–843.
69. Duprey, S., Naaim, A., Moissenet, F., Begon, M., and Chèze, L. (2017). Kinematic models of the upper limb joints for multibody kinematics optimisation: An overview. *J. Biomech.* 62, 87–94.
70. Rocke, R.D. (1966). *Transmission Matrices and Lumped Parameter Models for Continuous Systems* (California Institute of Technology), Doctoral dissertation.
71. Krotov, A.D. (2020). *Human Control of a Flexible Object: Hitting a Target with a Bull-Whip* (Northeastern University), Master's Thesis.
72. Hogan, N. (1985). Impedance control—An approach to manipulation. I—Theory. II—Implementation. III—Applications. *J. Dyn. Syst. Meas. Control* 107, 1–7.
73. Lipps, D.B., Baillargeon, E.M., Ludvig, D., and Perreault, E.J. (2020). Quantifying the multidimensional impedance of the shoulder during volitional contractions. *Ann. Biomed. Eng.* 48, 2354–2369.
74. Hermus, J., Doeringer, J., Sternad, D., and Hogan, N. (2020). Separating neural influences from peripheral mechanics: the speed-curvature relation in mechanically constrained actions. *J. Neurophysiol.* 123, 1870–1885.
75. Flash, T. (1987). The control of hand equilibrium trajectories in multi-joint arm movements. *Biol. Cybern.* 57, 257–274.
76. Rusu, R.B. (2010). Semantic 3d object maps for everyday manipulation in human living environments. *Künstl. Intell.* 24, 345–348.
77. Hoppe, H., DeRose, T., Duchamp, T., McDonald, J., and Stuetzle, W. (1992). Surface reconstruction from unorganized points. In *Proceedings of the 19th Annual Conference on Computer Graphics and Interactive Techniques*, pp. 71–78.
78. Sciacivco, L., and Siciliano, B. (2001). *Modelling and Control of Robot Manipulators* (Springer Science & Business Media).
79. Mussa-Ivaldi, F.A., and Hogan, N. (1991). Integrable solutions of kinematic redundancy via impedance control. *Int. J. Robot Res.* 10, 481–491.
80. Buchli, J., Theodorou, E., Stulp, F., and Schaal, S. (2011). Variable impedance control a reinforcement learning approach. *Robot. Sci. Syst.* VI 153.

STAR★METHODS

RESOURCE AVAILABILITY

Lead contact

Further information and requests for resources should be directed to and will be fulfilled by the lead contact, Moses C. Nah (mosesnah@mit.edu).

Materials availability

The study did not generate any new unique reagents, nor are there any restrictions on their availability.

Data and code availability

All data, code, images are available in <https://github.com/mosesnah-shared/whip-project-targeting> and the [Data S2](#).

METHOD DETAILS

Simulations were performed using the MuJoCo simulator.⁶⁷ Numerical integration used the semi-implicit Euler method with a time step of 0.1ms (10,000Hz).

Modeling

Upper-limb model - The manipulator

The human upper limb was modeled as a two-bar open-chain linkage. Everything distal to the wrist (i.e., hand, fingers etc.) was omitted from this model. The two limb segments — the upper arm (which extends from the shoulder to the elbow), and the forearm (which extends from the elbow to the wrist) — were treated as non-uniform cylinders. The geometric and inertial parameters of the limb segments were obtained from a computational model by Hatze ([Table 3](#)).⁶⁸

The upper-limb model had four actively-controlled DOF — three DOF at the shoulder and one DOF at the elbow. The glenohumeral joint of the shoulder was modeled as a 3-DOF spherical joint that was fixed in space, i.e., translational movements of the shoulder were omitted from the model.⁶⁹ The rotational movements of the 3-DOF shoulder joint were modeled as a sequence of three rotational joints whose axes of rotation were initially orthogonal — denoted as J1, J2 and J3 ([Figures 11A](#) and [11B](#)). The three rotational joints in order corresponded to flexion/extension (J1), adduction/abduction (J2) and lateral/medial rotation (J3). The movement of the elbow was modeled as a single-joint elbow flexion/extension (J4). Supination/pronation of the forearm was omitted from the model. At all four joints, independently controlled torque actuators were mounted co-axially.

Whip model - The manipulated object

The continuous dynamic behavior of a whip was modeled as an equivalent lumped-parameter model, in which the continuum was approximated and replaced by a finite-DOF system composed of lumped elements ([Figures 11A](#) and [11C](#)).⁷⁰ Each sub-model of a whip consisted of a uniform sphere mass, a massless cylinder, a linear rotational spring and a linear rotational damper. The sphere mass m was suspended from a massless cylinder with length l ([Figure 11C](#)). The radius of the sphere was 2cm. The other end of the massless cylinder was equipped with a 2-DOF universal joint, which consisted of two rotational joints whose axes of rotation were orthogonal. Each rotational joint was equipped with a linear rotational spring and a linear rotational damper, with coefficients k and b , respectively ([Figure 11C](#)). Values of the whip parameters were borrowed from an “experimentally-fitted” whip model, where the values were derived from experimental observations of an actual bullwhip ([Table 3](#)).^{13,71} 25 identical sub-models were serially connected in an end-to-end chain-like arrangement, which comprised a 50-DOF whip model.

The connection between the two models

To introduce no torque for the connection between the whip and upper-limb model, the k and b of the whip’s sub-model which directly attached to the end-effector of the upper-limb model were set to zero. Summarizing, the whole system resulted in a 54-DOF open-chain linkage.

Upper-limb controller

Impedance controller with gravity compensation

To account for physical interaction between the upper limb and the whip, the upper-limb model included a first-order joint-space impedance controller with gravity compensation:^{12,13,72}

$$\boldsymbol{\tau} = \mathbf{K}_q(\mathbf{q}_0 - \mathbf{q}) + \mathbf{B}_q(\dot{\mathbf{q}}_0 - \dot{\mathbf{q}}) + \boldsymbol{\tau}_G \quad (\text{Equation 1})$$

In this equation, $\boldsymbol{\tau}(t) \in \mathbb{R}^4$ denotes the net torque input commanded on the joint actuators; $\boldsymbol{\tau}_G(t) \in \mathbb{R}^4$ denotes gravity compensation torque; $\mathbf{q}(t) \in \mathbb{R}^4$ denotes the actual joint angle displacement measured from the virtual configuration (Figure 11B); $\mathbf{q}_0(t) \in \mathbb{R}^4$ represents the virtual trajectory of the impedance controller, which is referred to as “zero-torque trajectory”, i.e., neglecting gravitation effects, if the actual joint angle trajectory \mathbf{q} exactly matches the zero-torque trajectory \mathbf{q}_0 , no torque is exerted by the actuators. Subscripts 1 to 4 correspond to joints J1 to J4, respectively (Figure 11B); the zero-torque trajectory $\mathbf{q}_0(t)$ is the feedforward motion command of the controller which generates the upper-limb movement; $\mathbf{K}_q, \mathbf{B}_q \in \mathbb{R}^{4 \times 4}$ are constant symmetric joint stiffness and damping matrices, respectively.

Gravitational effects were compensated with $\boldsymbol{\tau}_G$, such that the actual upper-limb posture \mathbf{q} could exactly match with the zero-torque posture \mathbf{q}_0 when the whole model was at rest:¹²

$$\boldsymbol{\tau}_G = \mathbf{J}_{01}^T \mathbf{f}_{1,G} + \mathbf{J}_{02}^T \mathbf{f}_{2,G} + \mathbf{J}_{03}^T \mathbf{f}_{3,G} \quad (\text{Equation 2})$$

where $\mathbf{J}_{ij} \in \mathbb{R}^{3 \times 4}$ is a Jacobian matrix of frame j relative to frame i ; $\mathbf{f}_{i,G} \in \mathbb{R}^3$ denotes the gravitational force applied to frames $i = 1, 2$ and 3 , which are attached to the center of mass of the upper-limb, center of mass of the lower-limb and the hand, respectively.¹²

Motion planning – Zero-torque trajectory

The zero-torque trajectory $\mathbf{q}_0(t)$ (Equation 1) of the upper-limb model followed a discrete rest-to-rest minimum-jerk profile in joint-space coordinates¹²:

$$\mathbf{q}_0(t) = \mathbf{q}_{0,i} + (\mathbf{q}_{0,f} - \mathbf{q}_{0,i}) \cdot \left\{ 10 \left(\frac{t}{D} \right)^3 - 15 \left(\frac{t}{D} \right)^4 + 6 \left(\frac{t}{D} \right)^5 \right\} \quad (\text{Equation 3})$$

where subscripts i and f denote the initial and final (zero-torque) postures, respectively. For times greater than duration D (i.e., $t > D$), the zero-torque trajectory of the upper limb remained at the final posture $\mathbf{q}_{0,f}$. The zero-torque trajectory $\mathbf{q}_0(t)$ was determined by 9 movement parameters: 4 for initial posture $\mathbf{q}_{0,i}$, 4 for final posture $\mathbf{q}_{0,f}$, and 1 for movement duration D .

Stiffness and damping matrices for the joint-space impedance controller

The values of the stiffness matrix \mathbf{K}_q and damping matrix \mathbf{B}_q used for the joint-space impedance controller (Equation 1) were based on measurements of human subjects:^{73–75}

$$\mathbf{K}_q = \begin{bmatrix} 17.4 & 6.85 & -7.75 & 8.40 \\ 6.85 & 33.0 & 3.70 & 0.00 \\ -7.75 & 3.70 & 27.7 & 0.00 \\ 8.40 & 0.00 & 0.00 & 23.2 \end{bmatrix}, \mathbf{B}_q = 0.05\mathbf{K}_q \quad (\text{Equation 4})$$

Details of deriving the \mathbf{K}_q and \mathbf{B}_q matrices are presented in the Supplementary Materials, “Stiffness and Damping Matrices of the Impedance Controller”.

Evaluation metrics of the upper-limb movements

Data to evaluate the upper-limb movements were extracted with a sampling rate of 60Hz.

Best-fit plane and planar surface deviation (PSD)

To qualitatively assess the optimal upper-limb movements, the planarity of the upper limb’s end-effector (or hand) trajectory was calculated. The planarity of the movement was analyzed to check whether the movement in three-dimensional space can be reduced to a motion in two-dimensional space. To calculate the planarity of the hand trajectory, the best-fit plane and the corresponding planar surface deviation (PSD) of the end-effector’s trajectory were calculated.⁷⁶ If the value of planar surface deviation is small, then the corresponding optimal movement in 3D space is regarded to be ‘planar’ and the movement is approximated as a planar movement.

Given a set of sample points of the trajectory, the best-fit plane minimized the squared sum of orthogonal distances between the plane and the points:⁷⁷

$$\mathbf{n} \cdot (\mathbf{x} - \mathbf{c}) = 0$$

In this equation, $\mathbf{x} \in \mathbb{R}^3$ denotes the Cartesian position of an arbitrary point which lies on the best-fit plane; $\mathbf{n}, \mathbf{c} \in \mathbb{R}^3$ are the unit normal vector and the centroid of the best-fit plane, respectively.⁷⁷ Centroid \mathbf{c} was derived by:

$$\mathbf{c} = \frac{1}{N_s} \sum_{i=1}^{N_s} \mathbf{p}_i$$

In this equation, N_s denotes the number of sample points and $\mathbf{p}_i \in \mathbb{R}^3$ denotes the i^{th} sample point of the Cartesian position of the hand. \mathbf{n} is derived by taking the normalized eigenvector corresponding to the smallest eigenvalue of the 3D covariance matrix of N_s sample points.⁷⁷

The corresponding PSD of the best-fit plane was calculated by taking the root mean square of the orthogonal distances between the N_s sample points and the best-fit plane:

$$\text{PSD} = \sqrt{\frac{1}{N_s} \sum_{i=1}^{N_s} [\mathbf{n} \cdot (\mathbf{p}_i - \mathbf{c})]^2} \quad (\text{Equation 5})$$

Eigenstructure analysis of joint-space impedances

With the controller assumed in Equation 1, mechanical impedance plays a key role in the production of torque to move the limb. To assess the relation between the optimal joint-space action and the choice of mechanical impedance, the optimal joint-space movement was projected onto the eigenstructure common to the \mathbf{K}_q and \mathbf{B}_q matrices (Equation 4). A movement of the 4-DOF upper-limb model was decomposed as a unique weighted linear combination of the eigenvectors of matrix \mathbf{K}_q . Given their key role in the production of movement, these eigenvectors were termed the “eigenmovements” of the upper-limb model.

$\mathbf{q}_0 - \mathbf{q}$ (Equation 1) was described as a linear combination of the eigenmovements:

$$\mathbf{q}_0 - \mathbf{q} = w_1 \mathbf{v}_1 + w_2 \mathbf{v}_2 + w_3 \mathbf{v}_3 + w_4 \mathbf{v}_4 \quad (\text{Equation 6})$$

In this equation, $\mathbf{v}_1, \dots, \mathbf{v}_4 \in \mathbb{R}^4$ are the orthonormal eigenvectors of the stiffness matrix \mathbf{K}_q and coefficients $w_1, \dots, w_4 \in \mathbb{R}$ are the corresponding weights; subscripts are numbered in ascending size order of the corresponding eigenvalues. The eigenvectors and corresponding eigenvalues are presented in the Supplementary Materials, “Eigenmovements and the Corresponding Eigenvalues”.

To quantify the total “contribution” of the i^{th} eigenmovement to the upper-limb movement, the square root of the sum of the squared w_i -values from the onset to the end of the movement with duration D (Equation 3) was calculated. The weight with the highest value was identified, and the corresponding eigenmovement was regarded as making the highest contribution.

Eigenstructure analysis of task-space impedances

To assess the relation between the optimal end-effector movement and the choice of mechanical impedance, the optimal end-effector movement was also projected onto the eigenstructure of the task-space impedances. The relation between task-space and joint-space impedances is:^{78,79}

$$\mathbf{K}_p = \left(\mathbf{J}_{03} \mathbf{K}_q^{-1} \mathbf{J}_{03}^T \right)^{-1}, \mathbf{B}_p = \left(\mathbf{J}_{03} \mathbf{B}_q^{-1} \mathbf{J}_{03}^T \right)^{-1} \quad (\text{Equation 7})$$

In these equations, $\mathbf{K}_p, \mathbf{B}_p \in \mathbb{R}^{3 \times 3}$ are symmetric task-space stiffness and damping matrices, respectively. The kinematic stiffness term which emerges from the change of Jacobian matrix was neglected for the calculation.⁸⁰

The optimal end-effector velocity, $\mathbf{v}_{EE} \in \mathbb{R}^3$ was decomposed as a unique weighted linear combination of the eigenvectors of matrix \mathbf{K}_p and quantified how much the instantaneous hand movement was aligned along the eigenvectors of \mathbf{K}_p :

$$\mathbf{v}_{EE} = w_1 \mathbf{v}_1 + w_2 \mathbf{v}_2 + w_3 \mathbf{v}_3 \quad (\text{Equation 8})$$

In this equation, $\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3 \in \mathbb{R}^3$ are the orthonormal eigenvectors of the symmetric stiffness matrix \mathbf{K}_p and coefficients $w_1, w_2, w_3 \in \mathbb{R}$ are the corresponding weights; subscripts are numbered in ascending size order of the corresponding eigenvalues.

To quantify the alignment of the instantaneous hand movement along the eigenvectors of \mathbf{K}_p , the square root of the sum of the squared w_i -values from the onset to end of the movement with duration D (Equation 3) was calculated. The weight with the highest value was identified, and the corresponding eigenvector was regarded as the direction with which the hand movement was most aligned.

Optimization of the whip task

The task of reaching a spatial target with a whip was quantified as minimizing the distance between the whip and the target with a single discrete upper-limb movement. The target was modelled as a massless sphere with radius of 3cm. The distances between 25 nodes of the whip and the center of the target were tracked, and within these 25 nodes, the node that achieved the minimum distance with the target was used for the evaluation. This minimum distance L^* reached with a single discrete upper-limb movement served to assess performance. Formulating the targeting task as an optimization problem, the objective was to find the optimal 9 movement parameters ($\mathbf{q}_{0,i}, \mathbf{q}_{0,f}, D$) which minimized L^* . Optimization was performed using the DIRECT-L algorithm in the NLOpt Python toolbox.

Six targets with different distances and directions were defined for the targeting task. Targets 1, 2 and 3 were located closer to the shoulder to be well within reach, and targets 4, 5 and 6 located further from the shoulder to be just outside the range of reach (Figure 1). The optimization halted when the whip hit the target, or when the maximum iteration of 600 was conducted. When one of the 25 nodes of the whip hit the target, L^* was set to be 0m.

Sensitivity analysis of the optimal movement parameters

Sensitivity of task performance to the movement parameters was quantified by calculating the amount of error required to miss the target. Sensitivity analysis was conducted for targets 1 to 5, since for these targets the optimal parameters exist to hit the target. To identify the sensitivity of one movement parameter on the task performance, each of the 9 movement parameters were varied, while all other parameters were fixed to their optimal values (Table 1). For each movement parameter, the upper and lower value of the error to miss the target was identified. If the interval of the error was small, then the corresponding movement parameter was regarded to be sensitive for the task performance.