Learning from Elephants: Trajectory Planning via Kinematic Invariants

Francesca Sparnacci[1,2], Enrico Donato[1], Elisa Setti[1], Camilla Agabiti[1], Manolo Garabini[2], Lucia Beccai[3] and Egidio Falotico[1]
[1] *The BioRobotics Institute & the Dept. of Excellence in Robotics and AI, Sant'Anna School of Advanced Studies, Italy*
[2] *Department of Information Engineering & Research Center E. Piaggio, University of Pisa, Italy*
[3] *Soft BioRobotic Perception lab, Istituto Italiano di Tecnologia, Italy*

***Abstract*—Adaptive trajectory planning remains a central challenge in robotics, requiring solutions that integrate precision and generalizability. Derived from elephant trunk movements, this study introduces a novel trajectory planner for point reaching that builds upon biological demonstrations and satisfies kinematic invariants governing these movements. A Learning from Demonstration framework is designed to enable the extraction of reference trajectories from multiple demonstrations for different targets and their generalization to novel positions. In addition, we prove that the elephant trunk movements are governed by a generalized power law that relates the tangential velocity along a trajectory to a combination of its local curvature and torsion. Thus, this kinematic invariant is integrated into the planner design to ensure that generated trajectories are biologically faithful as well as capable of adapting to diverse goal positions.**

***Keywords*—Trajectory Planning, Soft Continuum Robotics, Elephant Trunk, Learning from Demonstration**

I. Introduction

The ability of robots to perform reaching tasks depends on effective trajectory planning, which ensures smooth, accurate, and adaptable movements [1], [2]. Addressing this challenge requires solutions that combine workspace generalizability, target precision, and motion efficiency [3], [4], qualities often identified in biological systems. Indeed, biological motion planning is driven by energy optimization principles to generate goal-directed, smooth and efficient paths [5], [6]. These principles give rise to governing kinematic invariants, such as power laws relating path geometry and travelling velocity, resulting in constant equi-affine velocity movements [7]. These insights have inspired human-like motion planning strategies for anthropomorphic robots, integrating biological principles to generate efficient reaching trajectories [8], [9].

Similarly, the elephant trunk provides a unique model for trajectory generation in continuum soft robots due to its compliant yet precise and coordinated movements [10]. Although soft robots already perform precise reaching [11] and agile motion planning [12], biological principles can improve their performance in terms of efficiency and adaptability [13].

This paper presents an elephant-inspired trajectory planner that learns reaching skills from trunk tip demonstrations.

Building on the work of *Zhao et al.* [14], we first generate reference trajectories from demonstrations and later generalize to novel goal positions. Our main contribution focused on preserving kinematic invariants to retain biologically relevant characteristics in both learned and generalized trajectories.

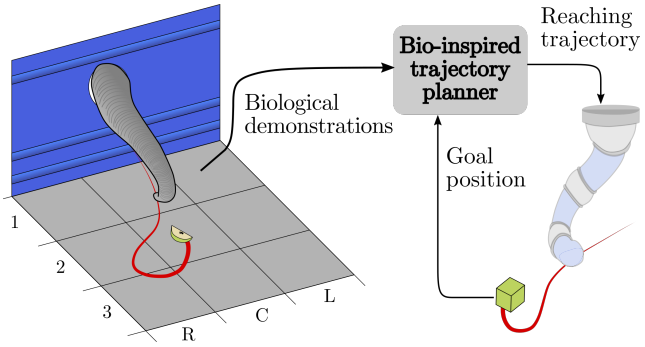

Fig. 1: This work aims to design a bio-inspired trajectory planner that learns goal-directed reaching motions from elephant trunk demonstrations and generalizes on the observed workspace.

II. Elephant Trunk Reaching Trajectory Analysis

Elephant trunk movements are investigated during goal-directed reaching. We provide insights on data collection and kinematic invariants evaluation.

A. Experimental protocol and data collection

An 18-year-old male African elephant named Itzek, living at the ZooSafari Park in Fasano (Italy), was involved in the experimentation, and all possible measures were taken to avoid any risk related to the experiments and to guarantee the animal's welfare. The experimental setup, shown in Figure 1, consisted of a closed structure within which the elephant reached the food by inserting his trunk through a $30cm$-diameter hole, as in a previous study [15]. Therefore, vision was impaired by the wall dividing the experimental area from the position of the elephant body, and the movements of the trunk were guided mainly by olfactory and tactile senses. Food targets were alternately arranged in a cross pattern: three in front of the hole (close: C1, mid: C2, and far: C3) and two lateral (right: R2, left: L2). Over targets, the elephant first explored the workspace to locate the food and learned the target location and then performed direct reaching.

Demonstrations were collected through RGB-D cameras, whose recordings were post-processed through a point identification algorithm to extract tip trajectories. In this study, only direct-reaching trials are considered, resulting in 41 trials (i.e., target: trials $\rightarrow$ C1: 3, C2: 20, C3: 12, R2: 3, L2: 3).

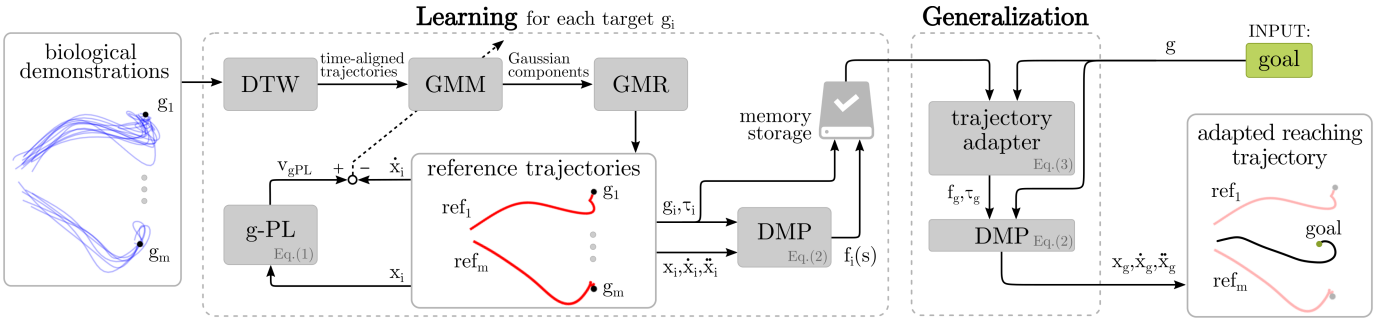

Fig. 2: Adaptive Learning from Demonstrations architecture. For each target, biological demonstrations are time-aligned via Dynamic Time Warping (DTW), and a reference trajectory is extracted through Gaussian Mixture Models (GMM) and Regression (GMR). The number of Gaussian components is optimized according to the generalized Power Law (g-PL) fulfillment. Each reference is separately modelled through a Dynamic Movement Primitive (DMP) and his forcing term is stored along with the respective target and duration. Finally, references are combined through a weighted sum adapter to generate an adapted reaching trajectory towards a new goal in the workspace.

B. Reaching kinematic invariants: the power law case

Demonstrations showcase the presence of a power law dependence between tip tangential velocity v and local path curvature κ_c [7]. Evidence comes from state-of-the-art models, which evaluate either planar or spatial trajectories.

- The 1/3-Power Law (PL) holds for planar movements [5], suggesting that the velocity increases under high curvature and decreases in straighter curves.

  $$\textbf{1/3-PL:}\quad v = \alpha\kappa_c^{-\beta}$$

  This relationship is captured by the invariant exponent $\beta_{planar} = 1/3$ and the trial-dependent scale factor α.
- In the case of spatial movements, we must consider the presence of curve torsion, which measures how much a path deviates from planarity. The 1/6-PL [16] provides a relationship between the tangential velocity v and a combination of path curvature κ_c and torsion τ_c. In particular, $\beta_{spatial} = 1/6$.

  $$\textbf{1/6-PL:}\quad v = \alpha\kappa_c^{-\frac{1}{3}}|\tau_c|^{-\frac{1}{6}} = \alpha(\kappa_c^2|\tau_c|)^{-\beta}$$

Nevertheless, both laws exhibit limitations: the planar law predicts unrealistic velocity peaks when curvature approaches zero, while the spatial law struggles when torsion is negligible.

To address these issues, we propose a generalised PL which is consistent with the geometric definition of curvature and torsion and applies the 1/3-PL when torsion is low (i.e., lower than a threshold $\bar{\tau}_c \to 0$) and the 1/6-PL otherwise.

$$\textbf{g-PL:}\quad v_{gPL} = \alpha \cdot \begin{cases} (\kappa_c^2|\tau_c|)^{-1/6} & \text{if } |\tau_c| > \bar{\tau}_c \\ \kappa_c^{-1/3} & \text{if } |\tau_c| \leq \bar{\tau}_c \end{cases} \tag{1}$$

III. Bio-inspired Reaching Trajectory Planner

We design the reaching trajectory planner shown in Figure 2, which generalizes over the workspace starting from biological demonstrations and preserves the power law invariant.

A. Reference trajectories extraction

For each target position, we extract a reference trajectory from a set of reaching demonstrations, while maintaining their geometrical and kinematic properties. First, the Dynamic Time Warping (DTW) algorithm [17] temporally aligns demonstrations so that trajectories can be consistently compared. Subsequently, we integrate Gaussian Mixture Models (GMM) and Regression (GMR) to generalise over multiple demonstrations, capturing their inherent variability. These models help in identifying the underlying probabilistic structure of target-oriented trajectories.

The number of Gaussian components plays a key role. Indeed, as demonstrations differ in terms of quantity and variability, the optimal number of Gaussian components must be determined to avoid overfitting and ensure the model preserves biological characteristics. Thus, a balance between preservation of trajectory variations and generalized power law fulfilment is pursued by RMSE minimization between planner-generated velocities and power law estimation.

B. Adaptive planner design

Once the representative trajectories are extracted, a planner that generalizes over the whole workspace is designed. To this aim, each trajectory is modelled as a Dynamic Movement Primitive (DMP) [18], a dynamical system where the forcing term $f(s)$ encodes the movement characteristics:

$$\begin{cases} \tau\dot{v} = K(g-x) - Dv + f(s)[g - x(0)] \\ \tau\dot{x} = v \end{cases} \tag{2}$$

Here $s \in [0,1]$ is a phase variable evolving according to a canonical system of the type $\tau\dot{s} = -\gamma s$, while $x(0)$ represents the initial condition of the trajectory $x(t)$ - now on simply referred as x - and g the desired goal position. The forcing function $f(s)$ is defined as a weighted sum of Gaussian basis functions and is learned from the reference trajectory. The temporal scaling factor τ, the elastic constant K, the damping gain D and the number of basis functions are constant over the reference trajectories. The DMP guarantees convergence to g, as $f(s)$ decays exponentially with bounded weights and it can reproduce any smooth trajectory while being invariant to spatial and temporal scaling.

Building upon this framework, we integrate an adaptive strategy for generating synthetic trajectories towards novel goals. A weighted sum of forcing functions of each reference trajectory is performed with respect to the distance between

the new target and the original m targets (where $m = 5$ in our case). For a new goal g, the updated forcing term is:

$$f_g(s) = \sum_{k=1}^{m} \mu_k f_k(s) \quad \text{where} \quad \mu_k = \frac{{d_k}^{-1}}{\sum_{k=1}^{m} {d_k}^{-1}} \tag{3}$$

with $d_k = ||g - g_k||$, where $(\cdot)_k$ denotes any state $(\cdot)$ referred to the k-th trajectory. Similarly, the time for the new reaching trajectory τ_g is computed as a weighted sum of reference trajectory durations, reusing the same weights μ_k of Equation 3. The parameters $f_g(s), \tau_g, g$ are then fed to a new DMP which elaborates an adapted reaching trajectory to g.

IV. Results

A. Biological reaching trajectories

By analysing each reaching trajectory in the dataset, it becomes evident that they cannot be considered either solely planar or spatial. Indeed, they exhibit a variable torsional component, as clear in Figure 3. Contrary to human arm movements, in which motions are typically (or can be approximated as) planar in sections, elephant trunk trajectories extend across multiple dimensions, suggesting the influence of underlying biomechanical principles related to its peculiar structure.

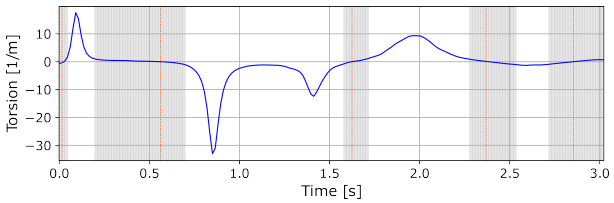

Fig. 3: Evolution of torsion for a single trajectory. When torsion is low (e.g. $|\tau_c| < 1$, in gray), it can be approximated to be planar.

B. Power laws evaluation

The power laws reported in Section II-B are evaluated on the reaching trajectories to assess their validity on the movements of the elephant trunk. We first perform a linear regression between velocity and curvature (and torsion, for the spatial case) in log-log planes for each experiment, to get the regression line that best approximates the data and whose slope corresponds to β_{planar} (or $\beta_{spatial}$). The regression is benchmarked by the standard error $SE(\beta)$ and the coefficient of determination R^2. Finally, the least square optimization is performed to identify the α that best fits the dataset.

For the 1/3-PL, the coefficient is estimated over the entire dataset to be $\bar{\beta}_{planar} = 0.35 \pm 0.08$, confirming a tight closeness to its theoretical value: $\bar{\beta}_{planar} \sim \beta_{planar}$. However, the $R^2 = 0.69$ reflects the discrepancies between the model and the biological data, suggesting that the 1/3-PL does not fully capture the velocity along the trajectory.

In contrast, the linear regression for the 1/6-PL estimates $\bar{\beta}_{spatial} = 0.18 \pm 0.04$, closely matches the respective theoretical value: $\bar{\beta}_{spatial} \sim \beta_{spatial}$. This power law provides a better fit, with a higher $R^2 = 0.76$, which indicates that the spatial power law is generally more effective in predicting velocity along the curve, but it is less suitable to low torsion.

These observations support the hypothesis of combining the two power laws with the g-PL formulation proposed in Equation (1). We select $\bar{\tau}_c = 1$ since it identifies the intervals in which the spatial model exhibits velocity peaks due to low torsion and allows for a continuous piecewise function.

The validity of this approach is additionally confirmed by the RMSE evaluation between the velocities estimated by the three power laws and the actual velocities in Table I, since the g-PL achieves the best fit for all the targets. The tangential velocity ranges across targets vary between 1.01 m/s (C1) and 2.02 m/s (C2). Thus, to provide a more intuitive understanding of the RMSE values, the error of our g-PL model remains within a range of 13–17% of observed biological velocities.

	C1	C2	C3	L2	R2
1/3-PL	0.21±0.03	0.37±0.12	0.34±0.09	0.34±0.07	0.32±0.06
1/6-PL	0.34±0.06	0.35±0.20	0.30±0.05	0.29±0.08	0.36±0.15
g-PL	**0.14±0.02**	**0.29±0.11**	**0.25±0.06**	**0.18±0.05**	**0.25±0.08**

TABLE I: Statistics of the RMSE for the evaluated power laws across different targets. Results are reported as: mean ± std [m/s].

C. Learning reaching trajectories from demonstrations

The reference trajectory extraction methodology is evaluated by varying the number of GMM clusters between 5 and 25 and computing the RMSE to the generalized power law output. The outcome from GMR of optimal solutions per each target is reported in Figure 4. In addition, we have estimated the jerk cost - i.e., cumulative jerk over the trajectory - of the optimal reference trajectories and others obtained simply as having minimal distance from biological demonstrations. Table II reports - for both approaches - the number of clusters that provides the lowest RMSE for each target together with the respective jerk cost. Our approach always reduces the jerk cost - with a decrease from -13% up to -95% for the same targets - ensuring lower energy consumption under equal robot control.

	C1	C2	C3	L2	R2
# clusters *	22	16	9	14	8
RMSE [m/s] *	0.22	0.24	0.15	0.21	0.18
Jerk cost [$\times 10^3$] *	**1300**	**4.8**	**0.84**	**2.3**	**9.1**
# clusters	23	22	19	20	23
Jerk cost [$\times 10^3$]	1500	11	16	14	180

TABLE II: Identification of the optimal clustering for target-specific GMMs. The * indicates results under g-PL optimization.

D. Adaptive reaching trajectory planning

After reference trajectory extraction, each is individually modelled through a DMP. Later, they are combined to generalize over new goals - using the approach described in Section III. Thus, the planner is able to generate a trajectory that not only precisely reaches the goal but also shares the kinematic characteristics of biological movements. To demonstrate the results, a grid of workspace points is selected as new goals for evaluating the planner, as shown in Figure 5.

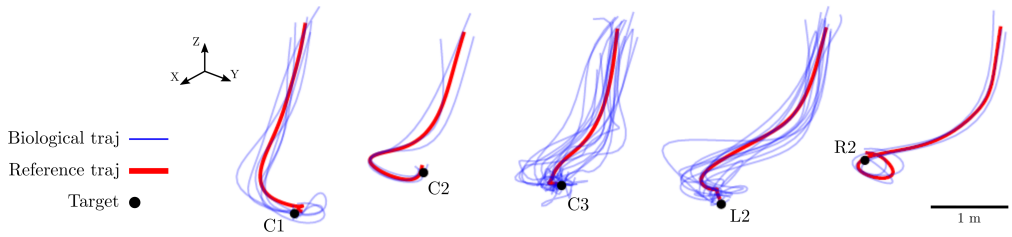

Fig. 4: Comparison of biological (in blue) and reference trajectories (in red) for each of the five targets, illustrating the alignment between observed and modeled reaching movements of the elephant trunk tip.

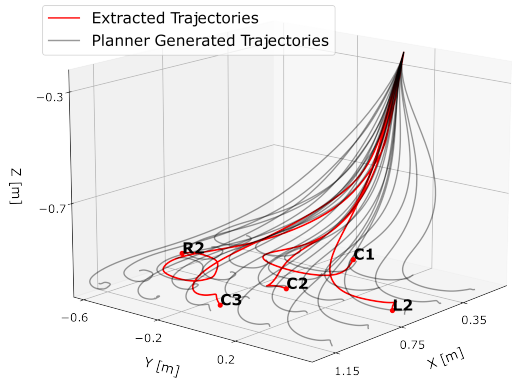

Fig. 5: Planning new reaching trajectories, extending the motion toward new goals and demonstrating the ability to generalize.

We evaluate the validity of the power law model on planner-generated trajectories to determine whether they capture the kinematic properties observed in the elephant trunk movement. The RMSE between the velocities of each trajectory and those predicted by the power law is evaluated and found to be 0.24 ± 0.11 m/s, in line with results in Table I.

V. Conclusion and Future Works

Results prove that the planner successfully generates trajectories that closely resemble the biological observations and respect their kinematic invariants. Furthermore, the proposed strategy allows more types of movement to be encoded in the same system, enabling its application to arbitrary goal positions. Our bio-inspired planner shows promise in reducing jerk costs, suggesting potential improvements in movement efficiency and energy consumption for soft robotic systems [19]. More empirical evidence will be needed to support these claims. In addition, future investigations will expand the study to the entire backbone to move towards elephant-inspired shape control strategies, essential to actively command deformations and perform dynamic tasks.

Acknowledgement

This work received funding from the European Union's Horizon 2020 research and innovation program under grant agreement No. 863212 (PROBOSCIS project) and by the Italian Ministry of Foreign Affairs and International Cooperation, project DESTRO grant n. PGR02061. We thank the staff of ZooSafari (Fasano, Italy) for their help in defining a safe and ethical experimental protocol and their support during the experiments.

References

[1] T. Zhang *et al.*, "Time-optimal and smooth trajectory planning for robot manipulators," *International Journal of Control, Automation and Systems*, vol. 19, pp. 521–531, 2021.

[2] F. Iori *et al.*, "Dmp-based reactive robot-to-human handover in perturbed scenarios," *International Journal of Social Robotics*, vol. 15, pp. 233–248, 2023.

[3] A. Ata, "Optimal trajectory planning of manipulators: A review," *Journal of Engineering Science and Technology*, vol. 2, no. 1, pp. 32–54, 2007.

[4] E. Donato *et al.*, "Plant-inspired behavior-based controller to enable reaching in redundant continuum robot arms," in *2023 IEEE International Conference on Soft Robotics (RoboSoft)*, 2023.

[5] F. Lacquaniti *et al.*, "The law relating the kinematic and figural aspects of drawing movements," *Acta psychologica*, vol. 54, no. 1-3, pp. 115–130, 1983.

[6] T. Flash *et al.*, "Models of human movement: Trajectory planning and inverse kinematics studies," *Robotics and Autonomous Systems*, vol. 61, pp. 330–339, 2013.

[7] P. Viviani *et al.*, "Minimum-jerk, two-thirds power law, and isochrony: Converging approaches to movement planning," *J Exp Psychol Hum Percept Perform*, vol. 21, no. 1, pp. 32–53, 1995.

[8] G. Averta *et al.*, "Exploiting upper-limb functional principal components for human-like motion generation of anthropomorphic robots," *Journ of NeurEng and Rehab*, vol. 17, pp. 1–15, 2020.

[9] G. Gulletta *et al.*, "A human-like upper-limb motion planner: Generating naturalistic movements for humanoid robots," *Intern Journ of Adv Rob Sys*, vol. 18, no. 2, p. 1 729 881 421 998 585, 2021.

[10] P. Dagenais *et al.*, "Elephants evolved strategies reducing the biomechanical complexity of their trunk," *Current Biology*, vol. 31, no. 21, pp. 4727–4737, 2021.

[11] M. Nazeer *et al.*, "Rl-based adaptive controller for high precision reaching in a soft robot arm," *IEEE Transactions on Robotics*, vol. 40, pp. 2498–2512, 2024.

[12] R. Jitosho *et al.*, "Reinforcement learning enables real-time planning and control of agile maneuvers for soft robot arms," in *Proceedings of 7th Conference on Robot Learning*, 2023.

[13] M. Hammond *et al.*, "Bioinspired soft robotics: State of the art, challenges, and future directions," *Current Robotics Reports*, vol. 4, no. 3, pp. 65–80, 2023.

[14] Y. Zhao *et al.*, "Generating a style-adaptive trajectory from multiple demonstrations," *International Journal of Advanced Robotic Systems*, vol. 11, no. 7, p. 103, 2014.

[15] M. Lo Preti *et al.*, "Sensorized objects used to quantitatively study distal grasping in the african elephant," *iScience*, vol. 26, no. 9, 2023.

[16] F. E. Pollick *et al.*, "Three-dimensional arm movements at constant equi-affine speed," *Cortex*, vol. 45, no. 3, pp. 325–339, 2009.

[17] S. Salvador *et al.*, "Toward accurate dynamic time warping in linear time and space," *Intelligent Data Analysis*, vol. 11, pp. 561–580, 2007.

[18] M. Ginesi *et al.*, "Overcoming some drawbacks of dynamic movement primitives," *Robot Auton Syst*, vol. 144, p. 103 844, 2021.

[19] J. D. Wong *et al.*, "The energetic basis for smooth human arm movements," *Elife*, vol. 10, e68013, 2021.