Movement Generation and Drawing in Robotics

Sylvain Calinon
Idiap Research Institute and Ecole Polytechnique Fédérale de Lausanne (EPFL)
sylvain.calinon@idiap.ch

***Abstract*—** We present an overview of our research to help robots acquire manipulation and drawing skills by imitation and self-refinement. We advocate frugal learning in our research, where frugality has two goals: 1) learning movement behavior skills from only few demonstrations or exploration trials; and 2) learning only the components of the skill that really need to be learned. Toward this goal, our work investigates the roles of movement primitives, distance fields, optimal control and ergodic control strategies as inductive biases to facilitate movement generation and drawing in robotics.

***Keywords*—** *robotics, movement primitives, ergodic control, optimal control, frugal learning*

I. Introduction

Despite significant advances in AI, robots still struggle with tasks involving physical interaction. Robots can easily beat humans at board games such as Chess or Go but are incapable of skillfully moving the game pieces by themselves (the part of the task that humans subconsciously succeed in). What makes research in robotics both hard and fascinating is that movement skills are tightly connected to our physical world and to embodied forms of intelligence. In the next sections, an overview of representation and control strategies is presented, with the goal of handling limited training data and uncertainty in robot manipulation and drawing activities.

II. Movement primitives as a superposition of basis functions

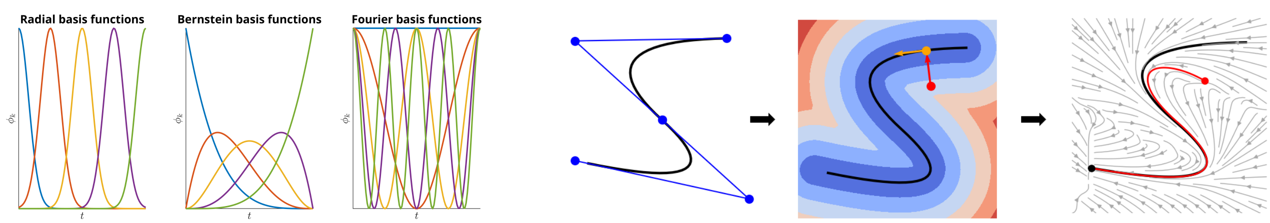

Fig. 1: *Left:* Examples of basis functions to encode movement primitives, see [1] for details. *Right:* Automatic transformation of a piecewise Bézier curve into a distance field and a dynamical system, see [2] for details.

The term *movement primitive* is often used in robotics to refer to the use of basis functions to encode trajectories [1]. It corresponds to an organization of continuous motion signals in the form of a superposition of simpler signals (in parallel and in series), which can be viewed as modular "building blocks" to create more complex movements, see Fig. 1-*left*. This principle, coined in the context of motor control [3], remains valid for a wide range of continuous time signals (for both analysis and synthesis).

The simpler form of *movement primitives* consists of encoding a movement as a weighted superposition of basis functions. This compression/factorization aims at working in a subspace of reduced dimensionality (smaller number of decision variables), while denoising the signal and capturing the essential aspects of a movement (e.g., continuity and smoothness).

Basis functions can for example be used to encode trajectories, whose input is a 1D time variable and whose output can be multidimensional. For basis functions $\phi(t)$ described by a time or phase variable t (see Fig. 1-*left*), the corresponding continuous signal $x(t)$ is encoded with $x(t) = \phi(t)\, w$, with t a continuous time variable and w a vector containing the superposition weights.

When using Bernstein basis functions, instead of considering high-order polynomials to encode long or complicated movement, is is useful in practice to instead consider a concatenation of lower-order polynomials (see Fig. 1-*right* for an example of concatenation with two quadratic Bézier curves). When concatenating curves, constraints on the superposition weights are typically considered. Since these weights can be represented as control points in the case of Bézier curves and splines (see the blue points in Fig. 1-*left*), the last point(s) of a curve and the first point(s) of the next curve will typically be constrained to maintain the continuity of the curve.

This work was supported by the State Secretariat for Education, Research and Innovation in Switzerland and the European Commission's Horizon Europe Program through the INTELLIMAN project (https://intelliman-project.eu/, HORIZON-CL4-Digital-Emerging Grant 101070136) and the SESTOSENSO project (http://sestosenso.eu/, HORIZON-CL4-Digital-Emerging Grant 101070310).

While movement primitives (MPs) provide a compact and modular representation of continuous trajectories, they are essentially time-based or phase-based representations, which can only handle perturbations in a limited way. This motivated the line of research of considering autonomous systems as a time-independent alternative providing control policies based on state-action pairs [4]–[6].

We proposed in [2] a simple and flexible approach that gathers the advantages of both representations by transforming MPs into autonomous systems. The key idea is to transform the explicit representation of a trajectory into an implicit shape encoded as a distance field. This conversion from a time-dependent motion to a spatial representation enables the definition of an autonomous dynamical system with modular reactions to perturbation. Asymptotic stability guarantees are provided by using Bernstein basis functions in the MPs, representing trajectories as concatenated quadratic Bézier curves, which provide analytic distance computation. This approach bridges conventional MPs with distance fields, ensuring smooth and precise motion encoding, while maintaining a continuous spatial representation. By simply leveraging the analytic gradients of the curve and its distance field, a stable dynamical system can be computed to reproduce the demonstrated trajectories while handling perturbations, without requiring a model of the dynamical system to be estimated, see Fig. 1-*right*.

III. Ergodic control

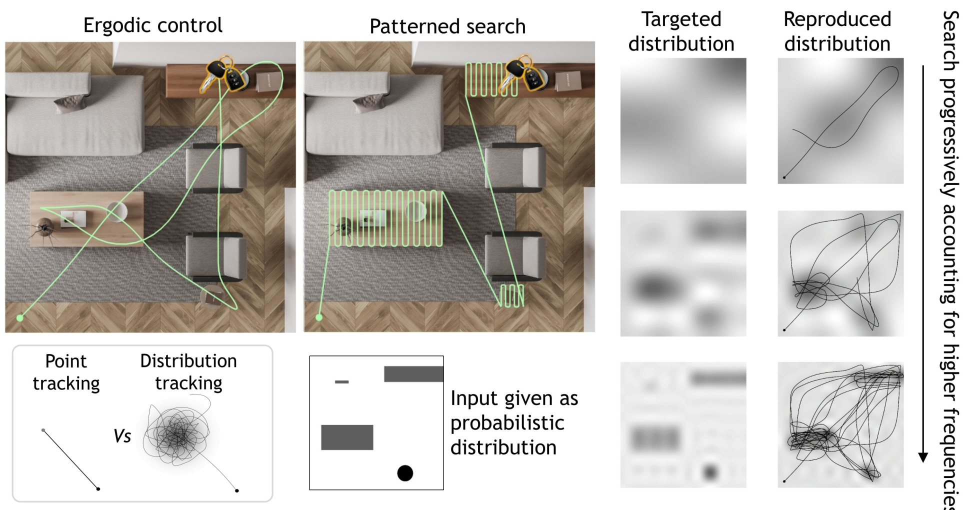

Fig. 2: Ergodic control with a spectral multiscale coverage (SMC) problem formulation.

An accurate reproduction of movements is often not enough to account for all possible sources of uncertainty that can interfere with a robotic task. To handle this issue, ergodic control can be exploited as way to cope with uncertainty of the real world. It can more generally be applied to a wide range of problems requiring the automatic exploration of regions of interest, including active sensing, localization, surveillance, peg-in-hole insertion or construction.

This is particularly helpful when the available sensing information and/or robot actuators are not accurate enough to fulfill the task with a standard controller. This includes situations with limited vision, soft robotic manipulator, or when the relative location of the robot gripper and the grasped tool/object is uncertain (namely, by specifying the task as a distribution instead of a point, see Fig. 2-*bottom*). In all these situations, the robot still has crude information about these locations, which can be exploited for guidance towards most promising areas.

In a collaborative task, ergodic control can also be used when the user's input is not accurate enough to fully reproduce the task, which then requires the robot to explore around the requested input (e.g., a point of interest selected by the operator). For picking and insertion problems, ergodic control can be applied to move in the region of the picking/insertion point, thereby facilitating the prehension/insertion [7]. It can also be employed for active sensing and localization (either detected autonomously, or with help by the operator). Here, the robot can plan movements based on the current information density, and can recompute the commands when new measurements are available (i.e., updating the spatial distribution used as target).

Research in ergodic control addresses fundamental challenges linking machine learning, optimal control, signal processing and information theory. Mathematically, ergodic theory studies the connection between the time-averaged and space-averaged behaviors of a dynamical system, which enables optimal exploration of an information distribution. A conventional tracking problem in robotics is characterized by a target to reach, requiring a controller to be computed to reach this target. In *ergodic control*, instead of providing a single target point, a probability distribution is given to the robot, which must cover the distribution in an efficient way. Ergodic control thus consists of moving within a spatial distribution by spending time in each part of the distribution in proportion to its density. The term ergodicity corresponds here to the difference between the time-averaged spatial statistics of the agent's trajectory and the target distribution to search in. By minimizing this difference, the resulting controller generates natural exploration behaviors.

Several ergodic control problem formulations exist, including *spectral multiscale coverage* (SMC) [8], diffusion-based [9], energy-based [10] and kernel-based [11] approaches. It can also be applied to surfaces characterized by meshes or point clouds [12]. The objective of SMC corresponds to a tracking problem in the spectral domain (matching of frequency components), see Fig. 2-*bottom*. The advantage of this simple control formulation is that it can be easily combined with other control objectives and constraints. It requires the spatial distribution to be decomposed as Fourier series, with a cost function comparing the spectral decomposition of the robot path with the spectral decomposition of the distribution, in the form of a weighted distance function with a decreasing weight from low frequency to high frequency components. The resulting controller allows the robot to explore the given spatial distribution in a natural manner, by starting from a crude exploration and by refining the search progressively (i.e., matching the Fourier coefficients with an increasing importance from low to high frequency components), see Fig. 2-*right*.

Fig. 2-*left* illustrates the use of ergodic control to search for a key in a living room, where regions of interest are specified in the form of a distribution over the spatial domain (see *bottom-center* inset). These regions of interest correspond to the different tables in the room where it is likely to find the key. The left and right columns respectively depict an ergodic search and a patterned search. For a coverage problem in which the duration is given in advance, a patterned search is well suited because the agent will optimize the path to reduce the number of transitions from one area to the other. In contrast, an ergodic control search will typically result in multiple transitions between the different regions of interest. When searching for a key, since we are interested in finding the key as early as possible, we have a problem in which the total search duration cannot be specified in advance. In this problem setting, an ergodic search is better suited, generating a natural way of exploring the environment by first proceeding with a crude search in the whole space and by progressively fine-tuning the search with additional details until the key is finally found. Indeed, to search for a key, a patterned search would be an inefficient and unnatural way of scanning the whole environment regularly by looking at all details regions by regions, instead of treating the problem at different scales (i.e., at different spatial frequencies as in SMC). Insertion tasks can similarly be achieved robustly with ergodic control, yielding a control strategy capable of coping with limited or inaccurate sensing information, see [7] for details.

IV. Application: Drawing robots

The modular movement representations and the control strategies presented in the previous sections are also exploited in our work to provide robots with drawing and writing capabilities.

In Fig. 3-*(a)*, movement primitives were used to generate the upper-body movement of a mini-humanoid robot drawing portraits (incl. bimanual gestures with natural-looking oscillations) [13]. Different shades of gray were generated through concentric arc circles to emulate the movement of a wrist during hatching.

In Fig. 3-*(b)*, a torque-controlled full-size humanoid was used to generate rapid and compliant movements to produce calligraphic graffiti art [14], [15]. In this experiment in collaboration with researchers from Goldsmiths University of London, an interface for artists was developed to let them edit trajectory paths by also considering variations in the underlying dynamics of the traces. To do this, the notion of control points in splines was extended to the use of Gaussian distributions to provide coordination and variation information. Optimal control was then used as the underlying mechanism to generate dynamic motions.

Fig. 3-*(c)* exploits a diffusion-based ergodic control method called HEDAC to draw portraits with a series of small strokes [16]. The underlying optimal control problem consists of covering a given distribution (the grayscale portrait image), so that the robot will pass more frequently on the dark regions of the image, effectively producing different shades of gray with a single ballpoint pen. When observing the robot at work, the doodles appear meaningless at first sight, with the global objective progressively revealed as time passes. Interestingly, while the result might appear to be composed of aleatory strokes, the whole process is purely deterministic and does not introduce any stochasticity in the generative process (namely, the code does not involve any random number generation). These doodles appear to be random when they are observed from close distance, but they are in reality the result of the optimal control problem, with the doodles revealing the global objective only when the drawing is finished.

References

[1] S. Calinon, "Mixture models for the analysis, edition, and synthesis of continuous time series," in *Mixture Models and Applications* (N. Bouguila and W. Fan, eds.), pp. 39–57, Springer, 2019.

[2] Y. Li and S. Calinon, "From movement primitives to distance fields to dynamical systems," *arXiv:2504.09705*, 2025.

[3] F. A. Mussa-Ivaldi, S. F. Giszter, and E. Bizzi, "Linear combinations of primitives in vertebrate motor control," *Proc. National Academy of Sciences*, vol. 91, pp. 7534–7538, August 1994.

[4] M. Hersch, F. Guenter, S. Calinon, and A. G. Billard, "Dynamical system modulation for robot learning via kinesthetic demonstrations," *IEEE Trans. on Robotics*, vol. 24, no. 6, pp. 1463–1467, 2008.

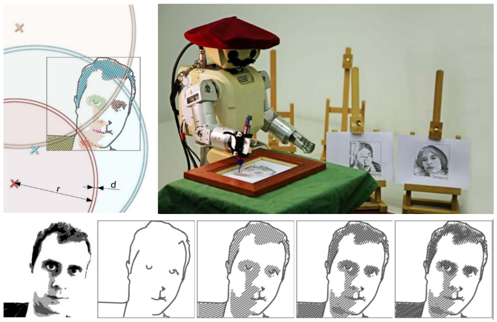

(a) Salvador DaBot [13]

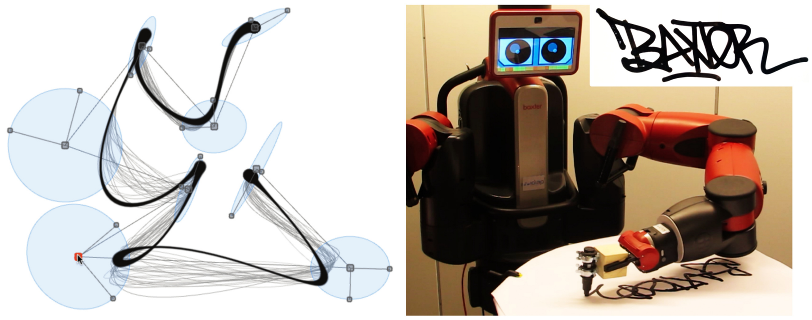

(b) Dynamic tags writing with Baxter [14], [15]

(c) DrozBot [16]

Fig. 3: Experiments with various forms of drawing robots.

[5] S. M. Khansari-Zadeh and A. Billard, “Learning control Lyapunov function to ensure stability of dynamical system-based robot reaching motions,” *Robotics and Autonomous Systems*, vol. 62, no. 6, pp. 752–765, 2014.

[6] M. Saveriano, “An energy-based approach to ensure the stability of learned dynamical systems,” in *Proc. IEEE Intl Conf. on Robotics and Automation (ICRA)*, pp. 4407–4413, 2020.

[7] S. Shetty, J. Silvério, and S. Calinon, “Ergodic exploration using tensor train: Applications in insertion tasks,” *IEEE Trans. on Robotics*, vol. 38, no. 2, pp. 906–921, 2022.

[8] G. Mathew and I. Mezić, “Spectral multiscale coverage: A uniform coverage algorithm for mobile sensor networks,” in *Proc. IEEE Conf. on Decision and Control*, pp. 7872–7877, Dec 2009.

[9] S. Ivić, B. Crnković, and I. Mezić, “Ergodicity-based cooperative multiagent area coverage via a potential field,” *IEEE Trans. on Cybernetics*, vol. 47, no. 8, pp. 1983–1993, 2017.

[10] A. Seewald, C. J. Lerch, M. Chancán, A. M. Dollar, and I. Abraham, “Energy-aware ergodic search: Continuous exploration for multi-agent systems with battery constraints,” in *Proc. IEEE Intl Conf. on Robotics and Automation (ICRA)*, pp. 7048–7054, 2024.

[11] M. M. Sun, A. Gaggar, P. Trautman, and T. Murphey, “Fast ergodic search with kernel functions,” *IEEE Transactions on Robotics*, vol. 41, pp. 1841–1860, 2025.

[12] C. Bilaloglu, T. Löw, and S. Calinon, “Tactile ergodic coverage on curved surfaces,” *IEEE Transactions on Robotics (T-RO)*, vol. 41, pp. 1421–1435, 2025.

[13] S. Calinon, J. Epiney, and A. G. Billard, “A humanoid robot drawing human portraits,” in *Proc. IEEE-RAS Intl Conf. on Humanoid Robots (Humanoids)*, pp. 161–166, 2005.

[14] D. Berio, S. Calinon, and F. Fol Leymarie, “Learning dynamic graffiti strokes with a compliant robot,” in *Proc. IEEE/RSJ Intl Conf. on Intelligent Robots and Systems (IROS)*, pp. 3981–3986, 2016.

[15] D. Berio, F. Fol Leymarie, and S. Calinon, “Interactive generation of calligraphic trajectories from Gaussian mixtures,” in *Mixture Models and Applications* (N. Bouguila and W. Fan, eds.), pp. 23–38, Springer, 2019.

[16] T. Löw, J. Maceiras, and S. Calinon, “drozbot: Using ergodic control to draw portraits,” *IEEE Robotics and Automation Letters (RA-L)*, vol. 7, no. 4, pp. 11728–34, 2022.