

DmpBbo: A versatile Python/C++ library for Function Approximation, Dynamical Movement Primitives, and Black-Box Optimization

Frek Stulp¹ and Gennaro Raiola^{1, 2}

¹ ENSTA-ParisTech (at the time of this work) ² Department of Advanced Robotics, Istituto Italiano di Tecnologia (IIT)

DOI: [10.21105/joss.01225](https://doi.org/10.21105/joss.01225)

Software

- [Review](#) ↗
- [Repository](#) ↗
- [Archive](#) ↗

Submitted: 26 January 2019

Published: 07 May 2019

License

Authors of papers retain copyright and release the work under a Creative Commons Attribution 4.0 International License ([CC-BY](#)).

General overview

Dynamical movement primitives (DMPs) (A. J. Ijspeert, Nakanishi, & Schaal, 2002, A. Ijspeert, Nakanishi, Pastor, Hoffmann, & Schaal (2013)) are one of the most popular representations for goal-directed motion primitives in robotics. They are also often used as the policy representation for policy improvement in robotics, a particular form of reinforcement learning. `dmpbbo` provides five software modules for the representation and optimization of dynamical movement primitives. These five modules are:

- `dynamicalsystems/`, various dynamical systems representing for instance exponential decay or spring-damper systems (standalone module).
- `functionapproximators/`, various function approximators such as Gaussian process regression, radial basis function networks, and Gaussian mixture regression (standalone module).
- `dmp/`, implementation of dynamical movement primitives, where various dynamical systems and function approximators in the first modules can be easily exchanged to get DMPs with different properties.
- `bbo/`, implementations of several stochastic optimization algorithms for the optimization of black-box cost functions (standalone module)
- `dmp_bbo/`, applies black-box optimization to the parameters of a DMP (depends on all other modules)

`dmpbbo` provides both a real-time C++ implementation, as well as an implementation in Python for non-roboticists.

`dmpbbo` is accompanied by an extensive tutorial on the motivation for dynamical movement primitives, and their mathematical derivation.

Advanced features

Several more advanced features implemented in `dmpbbo` are:

- Contextual dynamical movement primitives, which can adapt to variations of tasks (Stulp, Raiola, Hoarau, Ivaldi, & Sigaud, 2013)
- Dynamical movement primitives with gain schedules (Buchli, Stulp, Theodorou, & Schaal, 2011)




Figure 1: Overview

- Unified models for function approximators (Stulp & Sigaud, 2015)
- Covariance matrix adaptation in black-box optimization, which enables automatic exploration tuning (Stulp, 2012)

Applications

This library and its predecessors were used in the following scientific publications (Stulp, 2012, Stulp et al. (2013), Stulp, Herlant, Hoarau, & Raiola (2014), Stulp & Sigaud (2015)). The images below are snapshots of robotic applications where ‘dmpbbo’ was used. And here a list of videos:

- <https://www.youtube.com/watch?v=R7LWkh1UMII>
- <https://www.youtube.com/watch?v=MAiw3Ke7bh8>
- https://www.youtube.com/watch?v=jkaRO8J_1XI
- https://www.youtube.com/watch?v=i_JBRojCqcc

Robot names and credits in order of appearance: iCub (Photo by ISIR), MEKA (Photo by ENSTA ParisTech), Pepper (Photo by SoftBank)

References

- Buchli, J., Stulp, F., Theodorou, E., & Schaal, S. (2011). Learning variable impedance control. *International Journal of Robotics Research*, 30(7), 820–833. doi:[10.1177/02783364911402527](https://doi.org/10.1177/02783364911402527)
- Ijspeert, A. J., Nakanishi, J., & Schaal, S. (2002). Movement imitation with nonlinear dynamical systems in humanoid robots. In *Proceedings of the IEEE international conference on robotics and automation (ICRA)*. doi:[10.1109/ROBOT.2002.1014739](https://doi.org/10.1109/ROBOT.2002.1014739)
- Ijspeert, A., Nakanishi, J., Pastor, P., Hoffmann, H., & Schaal, S. (2013). Dynamical Movement Primitives: Learning attractor models for motor behaviors. *Neural Computation*, 25(2), 328–373. doi:[10.1162/neco_a_00393](https://doi.org/10.1162/neco_a_00393)
- Stulp, F. (2012). Adaptive exploration for continual reinforcement learning. In *International conference on intelligent robots and systems (IROS)* (pp. 1631–1636). doi:[10.1109/IROS.2012.6385818](https://doi.org/10.1109/IROS.2012.6385818)
- Stulp, F., & Sigaud, O. (2015). Many regression algorithms, one unified model – a review. *Neural Networks*. doi:[10.1016/j.neunet.2015.05.005](https://doi.org/10.1016/j.neunet.2015.05.005)

Stulp, F., Herlant, L., Hoarau, A., & Raiola, G. (2014). Simultaneous on-line discovery and improvement of robotic skill options. In *International conference on intelligent robots and systems (IROS)*. doi:[10.1109/IROS.2014.6942741](https://doi.org/10.1109/IROS.2014.6942741)

Stulp, F., Raiola, G., Hoarau, A., Ivaldi, S., & Sigaud, O. (2013). Learning compact parameterized skills with a single regression. In *IEEE-RAS international conference on humanoid robots*. doi:[10.1109/HUMANOIDS.2013.7030008](https://doi.org/10.1109/HUMANOIDS.2013.7030008)