Proceedings Paper:

Agudelo-España, D., Álvarez, M.A. orcid.org/0000-0002-8980-4472 and Orozco, Á.A. (2017) Definition and composition of motor primitives using latent force models and hidden Markov models. In: Progress in Pattern Recognition, Image Analysis, Computer Vision, and Applications. CIARP 2016. Iberoamerican Congress on Pattern Recognition, 08/11/2016-11/11/2016, Lima, Peru. Lecture Notes in Computer Science, 10125 . Springer Verlag , pp. 249-256. ISBN 9783319522760

https://doi.org/10.1007/978-3-319-52277-7_31

Definition and Composition of Motor Primitives using Latent Force Models and Hidden Markov Models

Diego Agudelo-España, Mauricio A. Álvarez, Álvaro A. Orozco

Faculty of Engineering, Universidad Tecnológica de Pereira,
La Julita, Pereira, Colombia
{dialagudelo,malvarez,aaog}@utp.edu.co

Abstract. Traditional methods for representing motor primitives have been purely data-driven or strongly mechanistic. In this work a different probabilistic motor primitive parameterization is proposed using latent force models (LFMs). The sequential composition of different motor primitives is also addressed using hidden Markov models (HMMs) which allows to capture the redundancy over dynamics by using a limited set of hidden primitives. The capability of the proposed model to learn and identify motor primitive occurrences over unseen movement realizations is validated using synthetic and motion capture data.

Keywords: Movement Representation, Motor Primitives, Latent Force Models, Hidden Markov Models, Switched Models, Multi-output GPs.

1 Introduction

The movement representation problem is at the core of areas such as robot imitation learning and motion synthesis. In these fields, approaches oriented to the definition of motor primitives as basic building blocks of more complex movements have been extensively used. This is supported by the fact that using a limited set of adjustable primitives tremendously reduces the number of parameters to be estimated for a particular movement and it seems to be the only way to cope with the movement high dimensionality [10].

Particularly, in the context of motor learning in humanoid robotics, motor primitives have been defined based on the theory of dynamical systems. The dynamical motor primitives (DMP) make use of the basic nonlinear dynamical systems behaviors represented by second order differential equations to encode a particular movement [5]. DMPs have been used extensively in the context of imitation learning [5]. However, a relevant limitation of these approaches is that the trajectory represented by DMPs is fixed and non-reactive to variations over the environment they were learnt from [9]. Also, the way DMPs are represented does not allow to take advantage of the potential correlation between the different modeled entities (e.g. robot joints).

In this work, a novel parameterization of motor primitives is proposed relying on the LFM framework [2]. The proposed solution represents a contribution along the same line of the switched dynamical latent force model (SDLFM) introduced in Álvarez et al. [1]. However, the main difference lies in the composition mechanism for different LFMs because in the SDLFM these primitives are articulated via switching points which become hyper-parameters of a Gaussian process covariance matrix. This covariance matrix grows quadratically on the length of the movement time series and as a consequence, the method is unscalable for long realizations of movement. HMM are used for composing different LFMs which allows to have a fixed-length covariance matrix for each primitive, and a simpler representation of the sequential dynamics of motor primitives. HMMs have been used before in the context of motor primitives [4, 11] to combine them either sequentially or simultaneously.

This paper is organized as follows. In Section 2 the probabilistic formulation of the motor primitive representation and the composition mechanism is introduced along with the inference algorithm. In Section 3 we show experimental results over synthetic and real data, with the corresponding discussion. Finally, some conclusions are presented in section 4.

2 Materials and Methods

2.1 Latent Force Models

The definition of motor primitives is done using the latent force model framework which was introduced in Álvarez et al. [2] motivated by the idea that for some phenomena a weak mechanistic assumption underlies a data-driven model. The mechanistic assumptions are incorporated using differential equations. Similarly to the DMP approach, a primitive is defined by a second order dynamical system described by

$$\frac{d^2 y_d(t)}{dt} + C_d \frac{dy_d(t)}{dt} + B_d y_d(t) = \sum_{q=1}^{Q} S_{d,q} u_q(t), \tag{1}$$

where C_d and B_d are known as the damper and spring coefficients respectively, and $\{y_d(t)\}_{d=1}^{D}$ is the set of D outputs of interest. In order to keep the model flexible enough to fit arbitrary trajectories even under circumstances where the mechanistic assumptions are not rigorous fulfilled [3], a forcing term is added and it is shown in the right side of equation (1). This forcing term is governed by a set of Q latent functions $\{u_q(t)\}_{q=1}^{Q}$ whose contribution to the outputs dynamics is regulated by a set of constants $\{S_{d,q}\}$ which are known as the sensitivities. Note that the fact of having multiple outputs governed by the same set of latent functions induces correlations between those outputs which can be of great benefit when motor primitives are used to model the movement of a set of robot joints.

The main difference of the LFM approach in contrast with the classical DMP is that it assumes Gaussian process priors with RBF covariance over the latent forcing functions. As a consequence of this assumption, it turns out that outputs are jointly governed by a Gaussian process as well and the corresponding covariance function can be derived analytically. The cross-covariance between outputs $y_p(t)$ and $y_r(t)$ under the dynamical model of equation (1) can be computed with

$$k_{y_p,y_r}(t,t') = \sum_{q=1}^{Q} \frac{S_{p,q}S_{r,q}l_q\sqrt{\pi}}{8\omega_p\omega_r} k^{(q)}_{y_p,y_r}(t,t'), \tag{2}$$

where $k^{(q)}_{y_p,y_r}(t,t')$ represents the covariance between outputs $y_p(t)$ and $y_r(t)$ associated to the effect of the $u_q(t)$ latent force and its exact form, along with definitions for constants ω_p, ω_r and l_q can be found in [2]. Remarkably, many of the desired properties of a motor primitive representation such as co-activation, modulation, coupling and learnability naturally arise as a consequence of the probabilistic formulation [6] which also allows to quantify the uncertainty over the learned movements.

2.2 Hidden Markov Models

Formally, an HMM models a sequence of observations $\mathbf{Y} = \{\mathbf{y}_1, \mathbf{y}_2, \ldots, \mathbf{y}_n\}$ by assuming that the observation at index i (i.e $\mathbf{y}_i$) was produced by an emission process associated to the k-valued discrete hidden state z_i and that the sequence of hidden states $\mathbf{Z} = \{z_1, z_2, \ldots, z_n\}$ was produced by a first-order Markov process. Therefore, the complete-data likelihood for a sequence of length n can be written as

$$p(\mathbf{Y}, \mathbf{Z}|\mathbf{A}, \boldsymbol{\pi}, \boldsymbol{\theta}) = p(z_1|\boldsymbol{\pi})p(\mathbf{y}_1|z_1, \boldsymbol{\theta}) \prod_{i=2}^{n} p(z_i|z_{i-1}, \mathbf{A})p(\mathbf{y}_i|z_i, \boldsymbol{\theta}), \tag{3}$$

where $\mathbf{A} = \{a_{j,j'}\}$ denotes the hidden state transition matrix, $\boldsymbol{\pi} = \{\pi_j\}$ is the initial hidden state probability mass function and $\boldsymbol{\theta}$ represents the set of emission parameters for each hidden state. The problem of how to estimate the HMM parameters $\zeta = \{\mathbf{A}, \boldsymbol{\pi}, \boldsymbol{\theta}\}$ is well-known and solutions for particular choices of emission processes have been proposed [8]. However, the use of novel probabilistic models as emission processes bring new challenges from the perspective of probabilistic inference and it also broadens the horizon of potential applications.

2.3 HMM and LFM

In this work HMMs are used differently by introducing a hybrid probabilistic model as emission process to represent motor primitives, namely the latent force models. The proposed model is based on the idea that movement time-series can be represented by a sequence of non-overlapping latent force models. This is motivated by the fact that movement realizations have some discrete and non-smooth changes on the forces which govern the movements. These changes can not be modeled by a single LFM because it generates smooth trajectories. Moreover, the use of HMM and LFM emissions enables us to capture the existing redundancy in dynamics over a movement trajectory since the whole trajectory is explained by a limited set of hidden primitives.

Formally, the overall system is modeled as an HMM where the emission distribution for each hidden state is represented by a LFM. Therefore the complete-data likelihood still fulfills the equation in (3) but the emission process is performed as follows

$$p(\mathbf{y}_i|z_i, \boldsymbol{\theta}, \boldsymbol{\chi}) = \mathcal{N}(\mathbf{y}_i|f(\boldsymbol{\chi}), I\sigma^2), \qquad f(t) \sim \mathcal{GP}(0, k_{y_p,y_r}(t, t'; \boldsymbol{\theta}_{z_i})), \tag{4}$$

where $k_{y_p,y_r}(.,.)$ represents the second order LFM kernel already defined in equation (2) with hyper-parameters given by $\boldsymbol{\theta}_{z_i}$ and σ^2 denotes the noise variance. Notice that the HMM framework allows the observable variables to have a different dimensionality with respect to the latent variables, in this case $\mathbf{y}_i$ is a continuous multivariate vector whereas z_i is univariate and discrete. It should also be noticed that there is an additional variable χ conditioning the emission process, this variable denotes the set of sample locations where the LFMs are evaluated and this set is assumed to be independent of the hidden variable values.

2.4 Learning the model

A common approach for estimating HMMs parameters is maximum likelihood via the expectation-maximization (EM) algorithm which is also known as the Baum-Welch algorithm. It can be shown that the E-step and the update equations for the parameters associated to the hidden dynamics $\{\mathbf{A}, \boldsymbol{\pi}\}$ are unchanged by the use of LFMs as emission processes and their exact form can be found in [8]. In order to update the emission process parameters $\boldsymbol{\theta}$, only one term of the $Q(\zeta, \zeta^{old})$ equation of the EM algorithm must be taken into account. This term is basically a weighted sum of Gaussian log-likelihood functions and gradient ascent methods can be used for optimizing and updating the emission parameters in the M-step.

3 Experimental Results

3.1 Synthetic data

A particular instance of the model was used for generating 20 trajectories with 20 segments each. This data-set was divided in two parts. One half was used for learning the motor primitives and the rest was used for validation (i.e. motor primitive identification). Notice that, as the model was formulated in section 2.3, it does not necessarily generate continuous trajectories, thus the synthetic trajectories were generated in such a way that the constant mean of each segment is equal to the last value of the last segment which produces real-looking realizations. To generate the synthetic trajectories, we consider an HMM with three

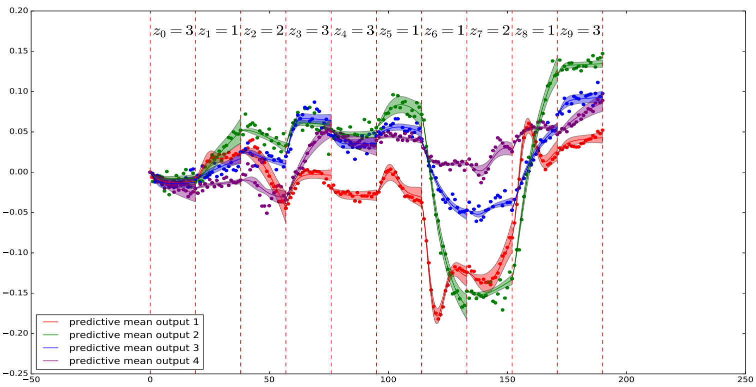

Fig. 1: Primitives identification over a synthetic trajectory. In the top, the most probable hidden state sequence $\{z_0, z_1, \ldots, z_9\}$ given by the inferred model is shown. The predictive mean after conditioning over the hidden state sequence and observations is also depicted with error bars accounting for two standard deviations.

hidden states with transition matrix $\mathbf{A}$ and initial state probability mass function $\boldsymbol{\pi}$. The same parameters inferred from training data are shown alongside (i.e. $\mathbf{A}^*$ and $\boldsymbol{\pi}^*$)

$$\mathbf{A} = \begin{bmatrix} 0.8 & 0.1 & 0.1 \\ 0.6 & 0.3 & 0.1 \\ 0.3 & 0.2 & 0.5 \end{bmatrix}, \boldsymbol{\pi} = \begin{bmatrix} 0.1 \\ 0.3 \\ 0.6 \end{bmatrix}, \qquad \mathbf{A}^* = \begin{bmatrix} 0.83 & 0.08 & 0.09 \\ 0.63 & 0.27 & 0.1 \\ 0.27 & 0.25 & 0.48 \end{bmatrix}, \boldsymbol{\pi}^* = \begin{bmatrix} 0.0 \\ 0.4 \\ 0.6 \end{bmatrix}.$$

Regarding the emission process, a LFM with 4 outputs ($D = 4$), a single latent force ($Q = 1$) and sample locations set $\boldsymbol{\chi} = \{0.1, \ldots, 5.1\}$ with $|\boldsymbol{\chi}| = 20$ was chosen. The sensitivities were fixed to one and they were not estimated. The actual emission parameters and the corresponding inferred values are depicted in table 1, where the spring and damper constants are indexed by the output index.

It can be argued that there is a similar pattern with respect to the original emission values, particularly for hidden states 1 and 3. The hidden state 2 exhibits a more diverse pattern in comparison to its original parameters. Nevertheless, when the inferred LFM covariances are plotted (see figure 2) for each hidden state, it is easy to see that the underlying correlations between outputs were successfully captured by the inferred emission processes. The poor length-scale estimation can be explained by the fact that the range covered by the sample locations set $\boldsymbol{\chi}$ (i.e. from 0.1 to 5.1) is reduced for this property to be noticeable and inferable.

A sample synthetic trajectory can be seen in figure 1. This is a trajectory used to validate the model's capability to detect motor primitives (i.e hidden

Table 1: LFM emission parameters

	Hidden State 1	Hidden State 2	Hidden State 3
Spring const.	{3., 1, 2.5, 10.}	{1., 3.5, 9.0, 5.0}	{5., 8., 4.5, 1.}
Damper const.	{1., 3., 7.5, 10.}	{3., 10., 0.5, 0.1}	{6., 5., 4., 9.}
Lengthscale	10	2	5
Spring const.*	{3.21, 1.05, 2.67, 11.09}	{0.67, 1.72, 3.39, 2.26}	{6.92, 10.66, 6.32, 2.3}
Damper const.*	{1.2, 3.36, 8.39, 9.61}	{0.5, 2.62, 1.07, 0.27}	{6.09, 5.54, 4.47, 9.76}
Lengthscale*	85.77	180.48	159.96

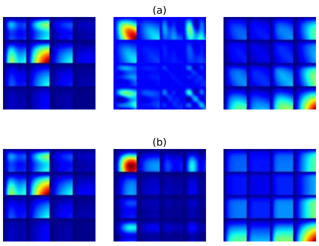

Fig. 2: Actual and inferred toy experiment covariance matrices for the 3 hidden states. Top row **(a)**: LFM covariance matrices used for generating the synthetic trajectories. Bottom row**(b)**: Estimated LFM covariance matrices.

states) given the inferred model parameters, which is achieved using the Viterbi algorithm[8]. The resulting hidden state sequence is shown on the top of figure 1, and it turned out to be exactly the same sequence used for generation, which was assumed to be unknown.

The Viterbi algorithm was executed for each validation trajectory and the resulting hidden state sequences were compared against the actual values used during the data-set generation. The correct hidden state was recovered with a success rate of 95% failing only in 10 out of 200 validation segments.

3.2 Real data

The real-data experiment consists in inferring a set of motor primitives from a group of realizations of the same particular behavior and assessing whether or not, a similar pattern is recovered over the inferred activated motor primitives for unseen realizations of the same behavior. To achieve this, the CMU motion capture database (CMU-MOCAP) is used. The chosen behavior is the walking action because it exhibits a rich redundancy over dynamics as a consequence of the cyclic nature of gait. Specifically, the subject No. 7 was used with trials $\{01, 02, 03, 06, 07, 08, 09\}$ for training and trials $\{10, 11\}$ for validation. In order to take advantage of the multiple-output nature of LFMs a subset of four joints was selected for the validation. The chosen joints are both elbows and both knees since they are relevant for the walking behavior and their potential correlations might be exploited by the model.

The model formulation given in section 2 implies fixing some model parameters a priori such as the number of hidden primitives, the number of latent forces and the sample locations set $\boldsymbol{\chi}$ used in the emission process in equation 4. Using a number of hidden states equals to three was enough for capturing the dynamics of the chosen behavior. Experiments with a higher number of hidden states were carried out with equally good results but the local variations associated to each observation started to be captured by the model thanks to the higher number of available hidden states. Similarly, the use of a high number of latent forces (e.g equals to the number of outputs) makes the model more flexible from the point of view of regression at the expense of favoring overfitting. By using three latent forces a good trade-off is obtained between the regression capability and the intended behavior generalization at the motor primitive level. Finally, the sample locations set $\boldsymbol{\chi}$ was defined to cover the interval $[0.1, 5.1]$ with 20 sample locations equally spaced (i.e. $|\boldsymbol{\chi}| = 20$). This choice was motivated by the functional division of a gait cycle into eight phases [7]: initial contact, loading response, mid stance, terminal stance, pre-swing, initial swing, mid swing, and terminal swing. Having $|\boldsymbol{\chi}| = 20$ a complete gait cycle over the selected observations is made up of roughly seven segments which is in accordance with the functional subdivision given that the initial contact can be considered an instantaneous event.

In figure 3 the identified motor primitives are shown along the resulting fit. The resulting Viterbi sequences over training and validation observations are

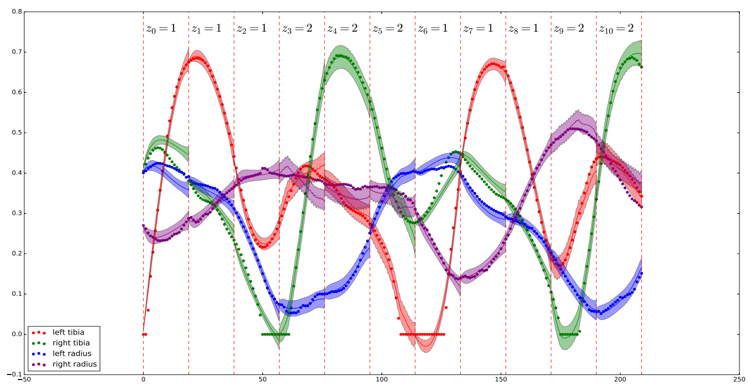

Fig. 3: Primitives identification over a real walking trajectory. In the top the most probable hidden state sequence $\{z_0, z_1, \ldots, z_{10}\}$ given by the inferred model is shown. The predictive mean with error bars is depicted for the four joints.

shown in table 2. From this table a cyclic sequential pattern over the identified motor primitives can be observed with a period of seven segments as expected. The first three segments of a gait cycle correspond to activations of the hidden state one, and the remaining gait cycle segments are generally explained by the hidden state two, although the last cycle segment exhibits high variability. Remarkably, the discussed pattern was also obtained over the unseen trajectories (i.e. No. 10, 11) which suggests that the sequential dynamics of the motor primitives associated to the walking behavior of subject seven were learned by the proposed model.

Table 2: Motor primitives identified over walking realizations.

Trial Number	**Motor primitives identified**
01	1, 1, 1, 2, 2, 2, 1, 1, 1, 1, 2, 2, 2, 1, 1, 1
02	1, 1, 1, 2, 2, 2, 1, 1, 1, 1, 2, 2, 2, 1, 1, 1
03	1, 1, 1, 2, 2, 2, 2, 1, 1, 1, 2, 2, 2, 2, 2, 1, 1, 1, 2, 2, 2
06	1, 1, 1, 2, 2, 2, 2, 1, 1, 1, 2, 2, 2, 2, 1, 1
07	1, 1, 1, 2, 2, 2, 2, 1, 1, 1, 2, 2, 2, 2, 2, 1, 1, 1, 2
08	1, 1, 1, 2, 2, 2, 2, 2, 1, 1, 1, 1, 2
09	1, 1, 1, 2, 1, 2, 1, 1, 1, 2, 2
10*	1, 1, 1, 2, 2, 2, 1, 1, 1, 2, 2
11*	1, 1, 1, 2, 2, 2, 2, 1, 1, 1, 2, 2, 2, 2, 2, 1, 1, 1, 2

4 Conclusions

In this work, a novel probabilistic parameterization of motor primitives and their sequential composition is proposed relying on LFMs and HMMs. We showed how to estimate the model's parameters using the EM algorithm and, through synthetic and real data experiments, the model's capability to identify the occurrence of motor primitives after a training stage was successfully validated.

For future work, alternative formulations which explicitly include the switching points as parameters are suggested to increase the model's flexibility. A further step in the validation might involve using it as a probabilistic generative model in a classification task. Thereby, potential applications can be related with distinguishing between different walking styles or identifying pathological walking.

Acknowledgments. This work has been supported by the project *Human-motion Synthesis through Physically-inspired Machine Learning Models* funded by Universidad Tecnológica de Pereira with code 6-15-3.

References

1. Alvarez, M., Peters, J.R., Lawrence, N.D., Schölkopf, B.: Switched latent force models for movement segmentation. In: Advances in neural information processing systems. pp. 55–63 (2010)
2. Alvarez, M.A., Luengo, D., Lawrence, N.D.: Latent force models. In: International Conference on Artificial Intelligence and Statistics. pp. 9–16 (2009)
3. Álvarez, M.A., Luengo, D., Lawrence, N.D.: Linear latent force models using gaussian processes. Pattern Analysis and Machine Intelligence, IEEE Transactions on 35(11), 2693–2705 (2013)
4. Chiappa, S., Peters, J.R.: Movement extraction by detecting dynamics switches and repetitions. In: Advances in neural information processing systems. pp. 388–396 (2010)
5. Ijspeert, A.J., Nakanishi, J., Schaal, S.: Learning attractor landscapes for learning motor primitives. Tech. rep. (2002)
6. Paraschos, A., Daniel, C., Peters, J.R., Neumann, G.: Probabilistic movement primitives. In: Advances in neural information processing systems. pp. 2616–2624 (2013)
7. Perry, J., Burnfield, J.M., Cabico, L.M.: Gait analysis: normal and pathological function (1992)
8. Rabiner, L.R.: A tutorial on hidden markov models and selected applications in speech recognition. Proceedings of the IEEE 77(2), 257–286 (1989)
9. Rückert, E.A., Neumann, G., Toussaint, M., Maass, W.: Learned graphical models for probabilistic planning provide a new class of movement primitives (2013)
10. Schaal, S.: Is imitation learning the route to humanoid robots? Trends in cognitive sciences 3(6), 233–242 (1999)
11. Williams, B., Toussaint, M., Storkey, A.J.: Modelling motion primitives and their timing in biologically executed movements. In: Advances in neural information processing systems. pp. 1609–1616 (2008)