
Sequence Alignment with Dirichlet Process Mixtures

Ieva Kazlauskaitė
Department of Computer Science
University of Bath
UK
i.kazlauskaitė@bath.ac.uk

Ivan Ustyuzhaninov
Centre for Integrative Neuroscience
University of Tübingen
Germany

Carl Henrik Ek
Faculty of Engineering
University of Bristol
UK

Neill D. F. Campbell
Department of Computer Science
University of Bath
UK

Abstract

We present a probabilistic model for unsupervised alignment of high-dimensional time-warped sequences based on the Dirichlet Process Mixture Model (DPMM). We follow the approach introduced in [7] of simultaneously representing each data sequence as a composition of a true underlying function and a time-warping, both of which are modelled using Gaussian processes (GPs) [10], and aligning the underlying functions using an unsupervised alignment method. In [7] the alignment is performed using the GP latent variable model (GP-LVM) [9] as a model of sequences, while our main contribution is extending this approach to using DPMM, which allows us to align the sequences temporally and cluster them at the same time. We show that the DPMM achieves competitive results in comparison to the GP-LVM on synthetic and real-world data sets, and discuss the different properties of the estimated underlying functions and the time-warps favoured by these models.

1 Introduction

For many machine learning tasks we are faced with the scenario of learning a model from non-stationary sequential data that results from studying several views of the same underlying phenomenon. Consider the following examples: replicated scientific experiments often vary in timing across different trials, different subjects might require different amounts of time to perform the same task, the peaks of the neuronal action potential waveforms do not match each other perfectly, telecommunications suffer from temporal jitter, climate patterns are often cyclic though particular events take place at slightly different times in the year. However, most sample statistics, such as mean and variance, are designed to capture variations in amplitude rather than phase or timing; this results in increased sample variance, blurred fundamental data structures and an exaggerated number of principal components necessary to describe the data. Therefore, the alignment of data is often performed as a non-trivial pre-processing step.

Traditionally, the notion of sequence correspondence is defined using a measure of pairwise similarity integrated across the sequences [2, 8, 4, 6, 12, 11]. We build on [7] where the models of the individual sequences and the alignment across sequences are cast within a single framework. Differently from the commonly used approach of interpolating between observations, the underlying functions and the warps are modelled using GPs which jointly regularise the solution space of the ill-posed alignment problem. The alignment is performed using dimensionality reduction which preserves similarities in the observation space while imposing the preference for dissimilar data points to be placed far apart in the latent space. In particular, the authors propose to use a GP-LVM that places independent GPs over the data features and optimises the corresponding latent variables.

Similarly to [7], we use GPs to model the data and the warps, which allows us to reject the observation noise in a principled manner, and imposes a smoothness constraint on the warping functions. In contrast, however, we consider a Dirichlet process mixture model [5] as a model for alignment. The DPMM performs clustering explicitly and allows us to automatically find the optimal number of underlying functions explaining the observed sequences.

2 Methodology

Let us assume that we are given J noisy observed sequences $\{\mathbf{y}_j\}_{j=1}^J$ where each sequence comprises of N time samples $\mathbf{y}_j := (y_{j,1}, \dots, y_{j,N}) \in \mathbb{R}^N$. We consider each sequence to be generated by sampling a latent function $f_j(x)$ at points $\{g_j(x_n)\}_{n=1}^N$, *i.e.* $y_{j,n} = f_j(g_j(x_n)) + \varepsilon_j$, where $\mathbf{x} := (x_1, \dots, x_N) \in \mathbb{R}^N$ are known evenly-spaced inputs (same for all sequences), $\{g_j(\cdot)\}_{j=1}^J$ are the warping functions (different for different sequences) and $\varepsilon_j \sim \mathcal{N}(0, \beta_j^{-1})$ is i.i.d. Gaussian noise.

We are interested in the case where the number of distinct latent functions $f_j(\cdot)$ is smaller than J , *i.e.* some of the observed sequences are generated from the same underlying functions. In this setting, the observed sequences were *misaligned* by applying different time warpings $g_j(\cdot)$ to the input locations. We treat the *aligned* sequences $\mathbf{s}_j := (f_j(x_1), \dots, f_j(x_N)) \in \mathbb{R}^N$ as if they were observed, and hence we refer to $\{\mathbf{s}_j\}$ as pseudo-observations. The rest of this section is structured as follows: first, we briefly describe the model over individual sequences $\{\mathbf{y}_j\}$, which is the same as in [7], and we refer the reader to this work for more details; next, we introduce an aligning objective based on the DPMM.

Model over time We model latent functions $f_j(\cdot)$ as stationary GPs, $f_j \sim \mathcal{GP}(0, k_{\theta_n}(\cdot, \cdot))$ with a squared exponential kernel. Denoting the finite-dimensional evaluation of the warping function as $G_j := (g_j(x_1), \dots, g_j(x_N))$, the joint likelihood of observed and aligned sequences is

$$p\left(\begin{bmatrix} \mathbf{s}_j \\ \mathbf{y}_j \end{bmatrix} \middle| G_j, \mathbf{x}, \theta_j\right) \sim \mathcal{N}\left(0, \begin{bmatrix} k_{\theta_j}(\mathbf{x}, \mathbf{x}) & k_{\theta_j}(\mathbf{x}, G_j) \\ k_{\theta_j}(G_j, \mathbf{x}) & k_{\theta_j}(G_j, G_j) \end{bmatrix} + \beta_j^{-1} I\right). \quad (1)$$

We encode our preference for smooth warping functions by also modelling $g_j(\cdot)$ as a GP with a smooth kernel function. The $\{G_j\}$ are constrained to be monotonic in the range $[-1, 1]$ using a cumulative sum of the set of auxiliary variables $\mathbf{u}_j \in \mathbb{R}^N$ with a softmax renormalisation. The distribution over G_j is $p(G_j | \mathbf{x}, \omega) \sim \mathcal{N}(0, k_{\omega_j}(X, X))$ with hyperparameters ω_j .

Alignment model So far we have modelled each sequence \mathbf{y}_j in isolation, and there is nothing in the model that encourages it to use the smallest possible number of distinct latent functions $\{f_j(\cdot)\}$ to explain the data. One way to introduce such a constraint is to add a regularisation term that encourages clustering of aligned sequences $\{\mathbf{s}_j\}$ into a small number of clusters. Indeed, the subset $\{\mathbf{s}_{c_j}\}$ of the aligned sequences that belong to the same cluster would be similar to each other, meaning that the corresponding latent functions $\{f_{c_j}(\cdot)\}$ are also similar when evaluated at \mathbf{x} , while the GP prior on $\{f_{c_j}(\cdot)\}$ enforces that the functions are smooth and nearby values are correlated.

We propose clustering $\{\mathbf{s}_j\}$ using the DPMM. As a non-parametric mixture model, DPMM allows us to automatically infer the number of clusters (*i.e.* distinct aligned sequences) from the data. The locations of the mixture components and the mixing weights are distributed according to a sample from a Dirichlet process, and, depending on the parameters of the process, such a prior seeks to explain the data using only a few mixture components, corresponding to aligning the underlying latent functions. Formally, the DPMM can be defined as follows:

$$G \sim \mathcal{DP}(\alpha, G_0), \quad \eta_j \sim G, \quad \mathbf{s}_j \sim p(\mathbf{s}_j | \eta_j), \quad (2)$$

where G , a sample from the Dirichlet process, can be thought of as an infinite weighted sum of delta functions at locations sampled from the base distribution G_0 (*e.g.* a Gaussian). The scaling parameter α controls the entropy of those weights; for small values of α only a few weights are significantly above 0, meaning that the data \mathbf{s}_j are sampled from one of a few mixture distributions with parameters η_j . An explicit construction of the DPMM is based on a stick breaking representation:

$$\begin{aligned} v_i &\sim \text{Beta}(1, \alpha), \quad i = \{1, 2, 3, \dots\}, \\ \pi_i &= v_i(1 - v_1) \dots (1 - v_{i-1}), \\ \eta_i &\sim G_0, \quad z_i \sim \text{Mult}(\pi), \quad \mathbf{s}_j \sim p(\mathbf{s}_j | \eta_{z_j}). \end{aligned} \quad (3)$$

We use DPMM to regularise the GPs that are fitted to the data, and hence we optimise the data likelihood of the DPMM jointly with the GP models. Since the data likelihood is not available in closed form, we approximate it with a variational lower bound [3]. Specifically, we approximate posterior distributions over v_i, η_i and z_j with factorised Beta, Gaussian and Multinomial distributions respectively:

$$q(\mathbf{v}, \boldsymbol{\eta}, \mathbf{z}) = \prod_{t=1}^{T-1} q_{\gamma_t}(v_t) \prod_{t=1}^T q_{\tau_t}(\eta_t) \prod_{j=1}^J q_{\phi_j}(z_j), \quad (4)$$

where T is the maximal number of clusters in the mixture (the infinite mixture model is truncated for the variational approximation). This approximation allows us to obtain a lower bound on the data likelihood:

$$\log p(\mathbf{s}_j) \geq \mathcal{L}_q(\mathbf{s}_j, \mathbf{v}, \boldsymbol{\eta}, \mathbf{z}) := \mathbb{E}_q[\log p(\mathbf{s}_j, \mathbf{v}, \boldsymbol{\eta}, \mathbf{z})] - \mathbb{E}_q[\log q(\mathbf{v}, \boldsymbol{\eta}, \mathbf{z})], \quad (5)$$

where both terms are analytically tractable, and we refer the reader to [3] for their exact expressions.

Learning We want the observed data $\{\mathbf{y}_j\}$ and aligned sequences $\{\mathbf{s}_j\}$ to be modelled by the GPs (by $f(g(\mathbf{x}))$ and by $f(\mathbf{x})$ respectively), and the aligned sequences $\{\mathbf{s}_j\}$ to be clustered into groups, therefore, we simultaneously maximise the GP data likelihood (1) and the lower bound on the DPMM likelihood (4). The GP likelihood (1) includes the evaluation of the warping GP, G_j , which we cannot integrate out analytically, therefore, we use a point estimate by including the likelihood of G_j in the objective, and directly optimising $\{\mathbf{u}_j\}$, which parametrise G_j . Overall, the optimisation objective is

$$\mathcal{J}(\mathbf{s}, G, \mathbf{v}, \boldsymbol{\eta}, \mathbf{z}) = \sum_{j=1}^J \left[p \left(\begin{bmatrix} \mathbf{s}_j \\ \mathbf{y}_j \end{bmatrix} \middle| G_j, \mathbf{x}, \theta_j \right) + \mathcal{L}_q(\mathbf{s}_j, \mathbf{v}, \boldsymbol{\eta}, \mathbf{z}) + p(G_j | \mathbf{x}, \omega) \right], \quad (6)$$

which we maximise w.r.t. the pseudo-observations of the aligned sequences $\{\mathbf{s}_j\}$, the DPMM variational parameters $\mathbf{v}, \boldsymbol{\eta}, \mathbf{z}$, the auxiliary variables of the warping functions $\{\mathbf{u}_j\}$, and the hyper-parameters of the GPs and the DPMM. Among the hyper-parameters of the DPMM are the scaling parameter α , variance of the base distribution G_0 (we assume it is a zero-mean Gaussian), and the variance of mixture components (we assume diagonal-covariance Gaussians with the same variance for all components). The scaling parameter α and the variance of G_0 are directly optimised (we include a Gamma prior on them), while the variance of mixture components is set to $1/\beta$, *i.e.* the estimated noise in the GP fits, which we assume to be the same for all sequences.

3 Experiments

Synthetic dataset We consider a set of 10 sequences generated using $\text{sinc}(\mathbf{x})$ and \mathbf{x}^3 , where \mathbf{x} is a linearly spaced vector of values in $[-1, 1]$, and warped using randomly generated monotonically increasing warping functions. We define (1) the mean (median) alignment error as the sum of means (medians) of pairwise distances between observations within each group of sequences in the N -dimensional space, (2) the data fit as the standard deviation of the estimated observational noise ($\sqrt{1/\beta}$), and (3) the warping complexity as the sum of the absolute values of differences between components of \mathbf{u}_j , which corresponds to the total variation of variables that define the warps. We provide a quantitative comparison between our method and the method that uses GP-LVM [7] in terms of the three criteria in Fig. A1 (see Appendix) as a function of a warping parameter ω where 0 corresponds to no warping and the warps get progressively larger as ω increases. Our model achieves lower alignment error in situations where ω is small or intermediate while the GP-LVM achieves smaller mean alignment errors for very large warps. However, the median errors of our model stay low even for high values of ω . An example of this behaviour is provided in Fig. A2 (see Appendix). If one of the sequences is an outlier due to a large warp, the DPMM tends to create a new cluster for it, while the GP-LVM favours the solution that recovers the two groups of sequences but is unable to align the sequences within the two groups accurately.

Heartbeats data We consider a data set of heart beat sounds, which have a clear "lub dub, lub dub" pattern that varies temporally depending on the age, health, and state of the subject [1]. Similarly to [7], our model is able to automatically align and cluster the heart sounds into two distinct patterns. In both approaches a Matérn 3/2 kernel is used, that takes into account the rapid variations in the recordings while also limiting the effect of the uninformative high frequency noise. Fig. 1 gives a comparison of the alignment and the clustering of a set of heartbeats. Our model achieves a smaller alignment error and estimates simpler warps (closer to an identity function).

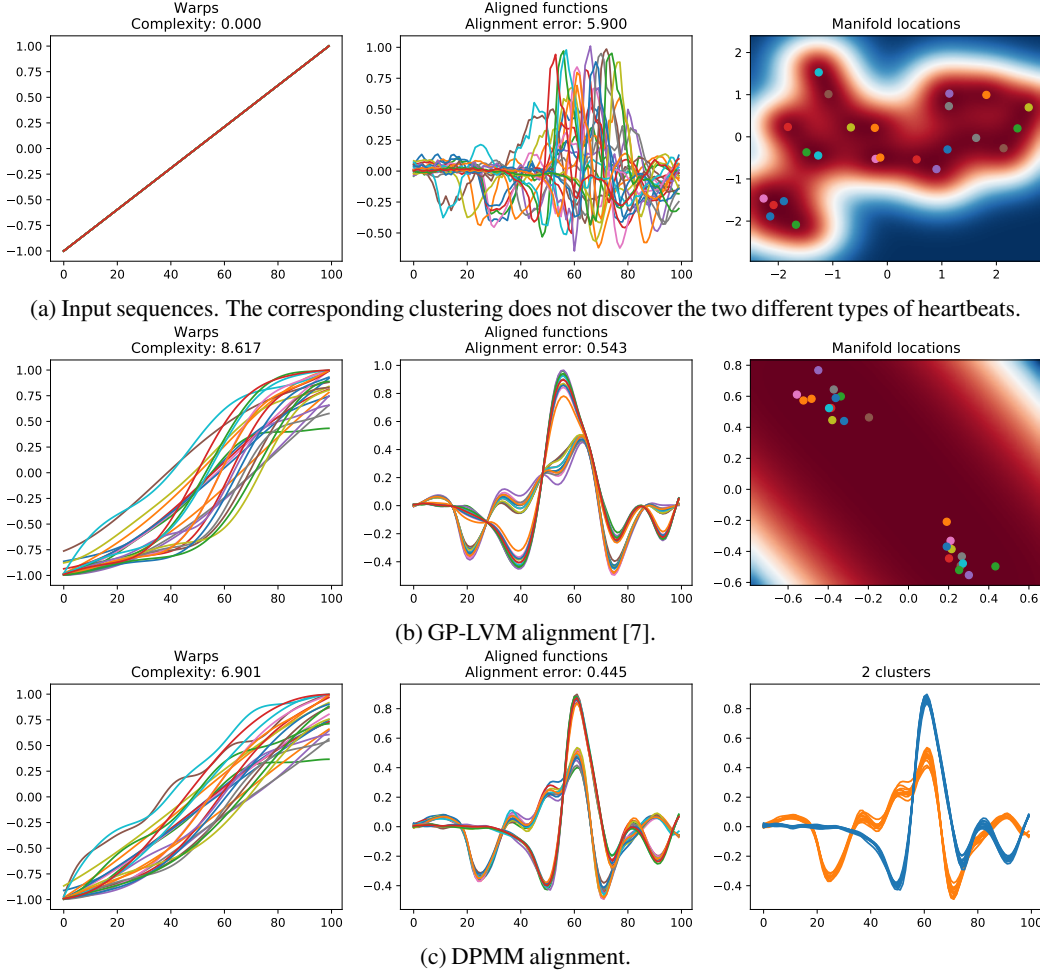


Figure 1: Alignment of heartbeats data [1].

4 Discussion

We have presented a probabilistic model that is able to implicitly align inputs that contain temporal variations. Our approach builds on [7], where we replace the latent variable model with a DPMM. Our model performs explicit clustering in the data space (and automatically estimates the number of clusters), while the GP-LVM performs dimensionality reduction and looks for a set of sequences that exhibit a simple structure in the latent low dimensional space (and allows to sample new sequences). While the experimental results suggest that both models perform well on real and synthetic data sets, they display different behaviour. The GP-LVM aligns the sequences constrained to them staying on a low-dimensional structure; in comparison, the DPMM does not have this global constraint and aligns sequences in each estimated cluster independently from the other clusters. In future work we propose developing a model which makes use of the global low-dimensional structure and the unconstrained alignment within clusters. The main limitation of our method is related to the sensitivity to the choice of hyper-parameters in the variational approximation of the lower bound on the DPMM likelihood, and one possible direction to overcome this limitation could be to perform variational inference over the hyper-parameters instead of optimising them directly. A successful application of our method on new data sets relies on the choice of priors for alignment, and we propose to explore the effect of using alternative prior processes, such as the Pitman-Yor process.

Acknowledgements

This work has been supported by EPSRC CDE (EP/L016540/1) and CAMERA (EP/M023281/1) grants.

References

- [1] P. Bentley, G. Nordehn, M. Coimbra, and S. Mannor. The PASCAL Classifying Heart Sounds Challenge 2011 (CHSC2011) Results. <http://www.peterjbentley.com/heartchallenge>.
- [2] D. J. Berndt and J. Clifford. Using Dynamic Time Warping to Find Patterns in Time Series. In *International Conference on Knowledge Discovery and Data Mining (KDD)*, 1994.
- [3] David M. Blei and Michael I. Jordan. Variational inference for dirichlet process mixtures. *Bayesian Analysis*, 1:121–144, 2005.
- [4] I. L. Dryden and K. V. Mardia. *Statistical Shape Analysis, with Applications in R. Second Edition*. 2016.
- [5] Thomas S. Ferguson. A bayesian analysis of some nonparametric problems. *Ann. Statist.*, 1(2):209–230, 03 1973.
- [6] Eugene Hsu, Kari Pulli, and Jovan Popović. Style translation for human motion. *ACM Trans. Graph.*, 24(3):1082–1089, 2005.
- [7] I. Kazlauskaitė, C. H. Ek, and N. D. F. Campbell. Gaussian Process Latent Variable Alignment Learning. *ArXiv e-prints*, March 2018.
- [8] E. J. Keogh and M. J. Pazzani. Derivative Dynamic Time Warping. In *SIAM International Conference on Data Mining*, 2001.
- [9] N. D. Lawrence. Probabilistic Non-Linear Principal Component Analysis with Gaussian Process Latent Variable Models. *Journal of Machine Learning Research (JMLR)*, 6:1783–1816, 2005.
- [10] C. E. Rasmussen and C. K. I. Williams. *Gaussian Processes for Machine Learning (Adaptive Computation and Machine Learning)*. 2005.
- [11] F. Zhou. Generalized Time Warping for Multi-modal Alignment of Human Motion. In *International Conference on Computer Vision and Pattern Recognition (CVPR)*, 2012.
- [12] F. Zhou and F. de la Torre. Canonical Time Warping for Alignment of Human Behavior. In *Advances in Neural Information Processing Systems (NIPS)*, 2009.

Appendix: Experiments on the synthetic data

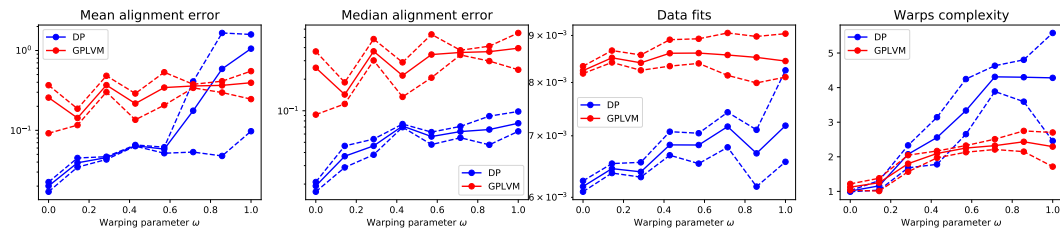


Figure A1: Comparison of alignment error, data fit and warp complexity for our model and [7]. Solid lines show mean values across 5 trials, dashed lines show min and max values.

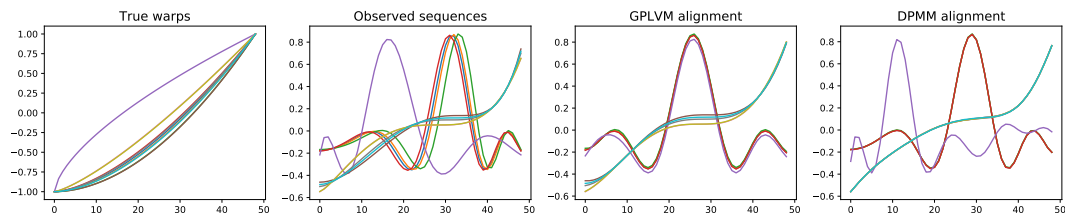


Figure A2: An example of behaviour for [7] and our model for the warping parameter ω of 1.0 in Fig. A1.