Preprocessing Data of Varying Trial Duration with Linear Time Warping to Extend on the Applicability of SNP-GPFA

Arjan Dhesi
School of Informatics
The University of Edinburgh
Informatics Forum EH8 9AB
A.Dhesi-1@sms.ed.ac.uk

Arno Onken
School of Informatics
The University of Edinburgh
Informatics Forum EH8 9AB
aonken@inf.ed.ac.uk

Abstract

Signal-noise Poisson-spiking Gaussian Process Factor Analysis (SNP-GPFA) is a popular model for analyzing neural data. However, a limitation exists, in that it cannot be applied to data of varying trial duration, limiting the range of experiments that can be performed. This work proposes data preprocessing techniques to feature align uneven length spike data, as well as findings from the application of SNP-GPFA to transformed rodent V1 data. We find that stretching followed by linear time warping is sufficient to align rodent V1 data in time and with respect to a paired visual stimulus and reward feature for successful application of SNP-GPFA.

1 Introduction

When analysing neuronal spiking data, dimensionality and noise reduction proves useful to present a clear understanding of typical cognitive activity. To this end, Gaussian Process Factor Analysis can be utilised [8]. One issue with typical readings is that there is a notable amount of trial varying noise that can make results more difficult to clearly examine and analyze. The noise in this case is defined as the variations about the determined "true signal" that is fully dependent on the presented stimulus. Additionally, it has been found this noise resides in a subspace that is mostly orthogonal to that of the true signals [6, 4], and there is still much unknown about the possible purpose of noise in the brain [3]. Taking this into account, it is useful to be able to visualize this noise and its subspace orientation alongside the true signals. The signal-noise Poisson-spiking Gaussian Process Factor Analysis (SNP-GPFA) model was designed for this purpose [4]. Another issue with typical behavioral data is that the duration of the trials often vary. This highlights an important limitation to SNP-GPFA, that it is inapplicable to data with varying trial duration, which constrains the freedom researchers have for the experiments they can conduct if they wish to apply this model to their collected data.

The current SNP-GPFA model being limited to data with fixed trial length limits experiments to those where the timing of the stimulus presentation is fully controlled by the researchers, or where a time limit is introduced. However, for experiments where the subject has control over the environment and the timings when certain stimuli are received, data with varying trial duration must be handled. As a general purpose solution, data can be transformed by time warping models to be of fixed length and feature aligned. Here, we utilize the time warping models presented by Williams et al. [7] of shift, linear, and piecewise linear warping for their proposed benefits over non-linear time warping techniques like dynamic time warping, which risk overfitting, and as they were designed explicitly for use on spike data. Hyperparameter values of the time warping and SNP-GPFA models were chosen through optimisation with Weights and Biases [1]. Application of the SNP-GPFA model to this transformed data results in signal latents that are as expected, clearly showcasing reaction to the reward zone, as well as findings related to the noise subspace orientation that agree with prior work.

2022 NeurIPS Workshop on Gaussian Processes, Spatiotemporal Modeling, and Decision-making Systems.

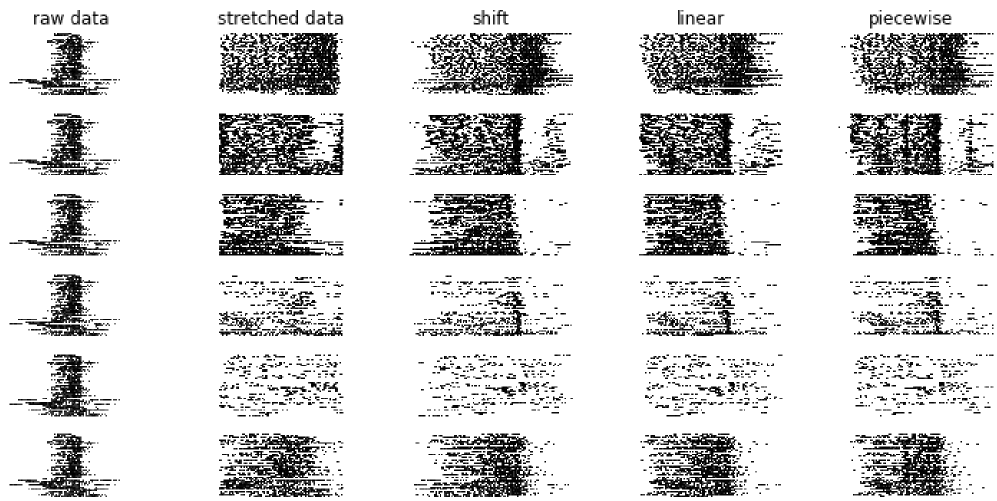

Figure 1: From left to right: the centre aligned raw data and its stretched, shift, linear and 1 knot piecewise linear warped outcomes. Each row is a different neuron, and within these rows are 69 smaller rows for the spike trains of that neuron over every trial.

2 Neuroscience dataset

The neuroscience dataset for this project contains the activity of many neurons in mice V1 as the subjects walk down a 160 cm length corridor in a virtual reality environment towards a reward zone beginning at the 120 cm mark, defined by walls of a different visual pattern, where a water reward is supplied to the mice. This data was recorded using two-photon calcium imaging recorded from 4 mice over 4 days. The recorded calcium traces were then deconvolved using MLSpike [2]. Full experimental details were described before [5]. We analyzed a sub-dataset that contained the recordings of 23 neurons over 69 trials. No neurons or trials were omitted.

3 Time warping

We propose to preprocess spike trains of different lengths for SNP-GPFA by applying time warping. Initially, the data must be stretched up to a fixed length before the time warping models can be applied. The results of the stretched data being warped by the shift, linear and piecewise linear models [7] can be seen in Figure 1, where the pre and post-warped spike trains for the first 6 neurons over every trial can be seen. It can be clearly seen in Figure 1 that a significant feature worth aligning is where the walls of the virtual corridor change appearance and the mouse receives a water reward. This is indicated by the sharp cut off point in spiking activity that arises around three quarters through the warped data, where the reward zone would be. When tested on the larger sub-datasets available, no noticeable change in warping quality was observed.

Cross validation was performed on the shift, linear and piecewise linear warping models with a 3/1/1 train/test/validation split. The results from this can be seen in Figure 2. Linear warping proves to be significantly advantageous over shift warping due to the difference in their validation and test loss results. Additionally, it can be seen that linear warping outperforms the piecewise warping models with increasing number of knots. This is likely due to the piecewise linear models overfitting to the data, as with an increasing number of knots the piecewise linear warping becomes more and more like a nonlinear warping model. Williams et al. mention that in most cases, shift and linear warping are sufficient to align unaligned spike data, but are insufficient to capture oscillations that are not phase locked to action or stimulus onset [7]. If this is the case in this situation, then shift and linear warping would be sufficient as the data preprocessing in this work aims to align the data in length and in respect to the stimulus of the reward zone. From these points, there is adequate justification to state that in terms of preprocessing this dataset with time warping models, linear warping should be used over shift and piecewise linear warping.

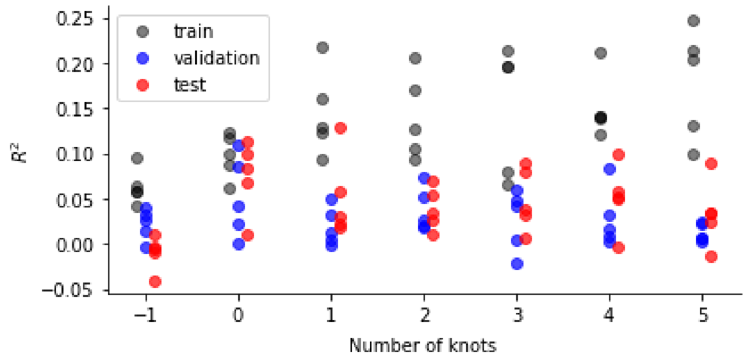

Figure 2: The cross validation performance of the shift, linear and piecewise linear warping models, measured by the coefficient of determination, R^2. The shift warping model is represented by the model with number of knots "-1" and the linear warping model is represented by the model with the number of knots "0". Significant overfitting can be seen in the piecewise linear warping models.

4 Signal and noise latents

We analyse the resulting principal components for the signal and noise subspaces, found through application of SNP-GPFA to the transformed data. Figure 3 (a) shows the resultant signal principal components when the SNP-GPFA model is passed stretched data warped by the linear model. A distinctive feature can be seen around the 23 second mark in the warped time, likely a reaction to the reward zone. Additionally, one of the PCs being completely flat is as expected due to the few neurons that have little to no activity. PC2 exhibits a sharp decline in firing rate from around the 21 to 23 second mark. This, coupled with the slight increases in firing rates of PCs 1, 3 and 4 around the same time presumably indicate the initial reactions to the differently patterned walls of the reward zone. A typical neuronal response to a paired stimulus and reward feature is for the firing rate to exhibit a sustained increase or decrease from stimulus presentation to reward, or simply a peaking at reward time [5]. PCs 1, 3 and 4 can clearly be seen to display peaking at reward time, whereas PC2 presents a sustained decrease behaviour.

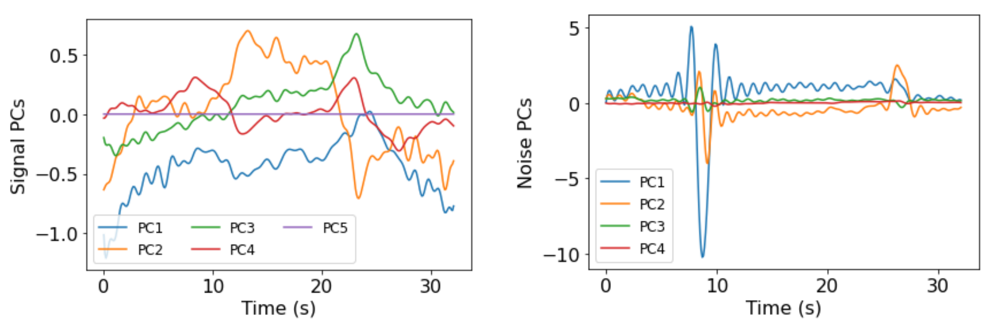

(a) (b)

Figure 3: The signal and first four noise PCs from the linearly warped stretched data. **(a)** Reactions to the visual stimulus can be seen between the 20-25 second mark. **(b)** A large feature can be seen around the 9 second mark.

The first four noise PCs can be seen in Figure 3 (b). Only the first four are visualised as PCs 5 and 6 simply lie overlapping along 0 on the x-axis. The most obvious feature of the noise PCs is the extremely large oscillation in PC1 around the 9 second mark. Apart from the smaller spike in PC2 at 26 seconds, the rest of the activity of the noise PCs is mainly consistent oscillatory action. The jump in activity of PC2 towards the end of the trials may be to do with increased neuronal activity related to the receiving of the water reward, such as noise from motor neurons acting to drink from the spout [3]. The large oscillations earlier on in the trials are harder to explain, as there is nothing different for the mouse to have a reaction to in the first half of the virtual corridor. The likely explanation is that it is a specific artifact of this particular linearly warped sub-dataset. This can be verified when comparing to the noise PCs of different linearly warped stretched sub-datasets, where this specific oscillation cannot be seen. This could possibly be due to the increased size of various other sub-datasets overwhelming any anomalous results.

5 Subspace orientations

We further examine how the noise subspace orientation changes over time. From the plots in Figure 4 (a) it is evident that for the most part, the total amount of noise is relatively evenly split between noise orthogonal and aligned to the signal subspace. There is also visibly a large amount of inconsistency in the total amount of the noise. Specifically, it can be seen that approaching the reward zone at roughly 23 seconds the level of overall noise drops significantly relative to earlier in the trials, but then proceeds to increase thereafter. Again, this latter noise may be to do with the receiving of the water reward, while the earlier noise is harder to interpret. When viewing Figure 4 (b), it can be seen that for the most part the noise is more likely to be orthogonal than aligned in general when normalized according to the total variance, contrary to what may seem the case when viewing the total L2 norm of the noise. Through the examination of these findings and those of other sub-datasets, it is found that in general the noise is more likely to be orthogonal to the signal subspace when the overall L2 norm of the noise is low.

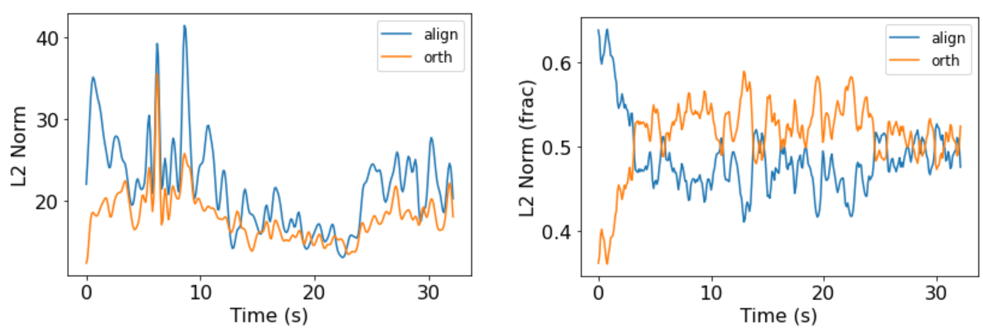

(a) (b)

Figure 4: The linearly warped stretched data noise subspace orientation averaged over all trials. **(a)** The total L2 norm of the noise that is orthogonal and aligned to the signal subspace. **(b)** Same as (a) but normalized to show the fraction of the noise activity variance.

6 Conclusion

From this work, it has been shown that neuroscience spike data with varying trial duration can be accurately transformed to be fixed in length and feature aligned, through the use of time warping transformation methods. Further, application of the SNP-GPFA model to this preprocessed data results in latents that are as expected, disentangling the true signal and noise components. This should allow for researchers in the field of neuroscience to be much more liberal with the conditions of their experiments if they wish to apply SNP-GPFA, or even standard GPFA, to their data, extending the use of Gaussian Processes in time series data modelling further. Freedom is granted for the subjects of experiments to perform in more complex scenarios, providing greater opportunity to learn more about many aspects of cognitive activity. In recommendation of the optimum preprocessing strategy, linear warping of stretched data is advised due to linear warping's better performance over shift and piecewise linear warping. However, this may not be the case for largely different datasets or when oscillations not phase locked to action or stimulus onset are important.

Future directions include application of the preprocessing methods to a dataset obtained from an experiment with different conditions and features. This could reveal weaknesses in the proposed preprocessing methods that otherwise would not become apparent. No negative societal impacts are foreseen from this work.

Acknowledgments and Disclosure of Funding

The authors declare no competing interests.

References

[1] Weights and biases. `https://wandb.ai/site`.

[2] Thomas Deneux, Attila Kaszás, Gergely Szalay, Gergely Katona, Tamás Lakner, Amiram Grinvald, Balázs József Rózsa, and Ivo Vanzetta. Accurate spike estimation from noisy calcium signals for ultrafast three-dimensional imaging of large neuronal populations in vivo. *Nature Communications*, 7, 2016.

[3] Aldo Faisal, Luc Selen, and Daniel Wolpert. Noise in the nervous system. *Nature reviews. Neuroscience*, 9:292–303, 05 2008.

[4] Stephen Keeley, Mikio Aoi, Yiyi Yu, Spencer Smith, and Jonathan W Pillow. Identifying signal and noise structure in neural population activity with gaussian process factor models. In H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 13795–13805. Curran Associates, Inc., 2020.

[5] Janelle MP Pakan, Valerio Francioni, and Nathalie L Rochefort. Action and learning shape the activity of neuronal circuits in the visual cortex. *Current Opinion in Neurobiology*, 52:88–97, 2018. Systems Neuroscience.

[6] Carsen Stringer, Marius Pachitariu, Nicholas Steinmetz, Charu Bai Reddy, Matteo Carandini, and Kenneth D. Harris. Spontaneous behaviors drive multidimensional, brainwide activity. *Science*, 364(6437):eaav7893, 2019.

[7] Alex H. Williams, Ben Poole, Niru Maheswaranathan, Ashesh K. Dhawale, Tucker Fisher, Christopher D. Wilson, David H. Brann, Eric M. Trautmann, Stephen Ryu, Roman Shusterman, Dmitry Rinberg, Bence P. Ölveczky, Krishna V. Shenoy, and Surya Ganguli. Discovering precise temporal patterns in large-scale neural recordings through robust and interpretable time warping. *Neuron*, 105(2):246–259.e8, 2020.

[8] Byron M Yu, John P Cunningham, Gopal Santhanam, Stephen Ryu, Krishna V Shenoy, and Maneesh Sahani. Gaussian-process factor analysis for low-dimensional single-trial analysis of neural population activity. In D. Koller, D. Schuurmans, Y. Bengio, and L. Bottou, editors, *Advances in Neural Information Processing Systems*, volume 21. Curran Associates, Inc., 2008.

A Appendix

Code available at `https://github.com/ArjanDhesi/SNP_GPFA_with_preprocessing`