A Probabilistic Modeling Approach for Uncovering Neural Population Rotational Dynamics

Shamim Nemati[1], Scott W. Linderman[1] and Zhe Chen[2]
[1] Harvard University; [2] Massachusetts Institute of Technology

1 Summary

An important task in computational neuroscience is to understand neural representations of population codes, using either Bayesian approach or dynamical systems approach. We provide a probabilistic interpretation of the jPCA algorithm, recently proposed by Churchland et al. [1] for revealing the rotational patterns and dynamics in neuronal population activity. Using the jPCA technique, the authors showed that neural firing rates during reaching tasks (a non-rhythmic activity) exhibit a quasi-periodic pattern. Here we show that the original jPCA algorithm can be formulated as a linear algebra problem, and therefore be solved using the well-known technique of polar decomposition. Moreover, we show that the technique can be cast into a probabilistic framework which allows for application of Bayesian methods for model identification and parameter fitting. Moreover, missing data also can be handled within this framework using an expectation-maximization approach. Our approach provides an alternative yet more robust approach for dimensionality reduction and visualization in analyzing noisy neural population codes.

2 The jPCA Tranform

Consider the a multivariate time series of (smoothed) neuronal firing rates $\mathbf{r}(t) = [r_1(t), \cdots, r_N(t)]$, where N denotes the total number of recorded neurons. Let $R \in \mathbb{R}^{N\times CT}$ be a matrix of firing rates for T time-steps and over C reaching trials (possibly including multiple conditions). The jPCA algorithm first projects the columns of R onto a lower dimensional space using PCA, resulting in a reduced data matrix $X \in \mathbb{R}^{K\times CT}$ ($K \ll N$). Next, the data is modeled as a linear time-invariant (continuous) dynamical system of the form $\dot{\mathbf{x}}(t) = M\mathbf{x}(t)$, where the linear transformation matrix M is constrained to be skew-symmetric (i.e., $M^\intercal = -M$). The jPCA algorithm proceeds by projecting each column of the data matrix X onto the eigenvectors of the M matrix. These eigenvectors come in complex conjugate pairs (corresponding to a complex conjugate pair of eigenvalues). Thus, given a pair of eigenvectors $\{v_k, \bar{v}_k\}$, the k-th jPCA projection plane axes are defined as $u_{k,1} = v_k + \bar{v}_k$ and $u_{k,2} = j(v_k - \bar{v}_k)$, which are then suitably normalized [1]. To see why jPCA reveal rotational dynamics, recall the solution to the jPCA continuous differential equation is given by $\mathbf{x}(t) = e^{Mt}\mathbf{x}(0)$. If M is skew-symmetric then e^M is orthogonal; thus, describing rotation of the initial condition $\mathbf{x}(0)$ over time.

2.1 Equivalence of jPCA and Polar Decomposition in Discrete Time

In practice, neuronal firing rates are binned and the resulting time-series are in discrete time. For instance, Churchland et al. [1] used a 10-ms bin, and approximate the continuous time derivatives using a first difference method: $\frac{\mathbf{x}[t+\Delta t]-\mathbf{x}[t]}{\Delta t} = M\mathbf{x}[t]$, with the discrete time index $t = k\Delta t$, $\Delta t = 10$ ms. For simplicity let us assume that $\Delta t = 1$. Then a possible discrete analog of jPCA dynamical system is $\mathbf{x}[t+1] = (I + M)\mathbf{x}[t]$. Note that $I + M$ is a first order (Taylor's series) approximation to the orthogonal transformation $Q = e^M$. However, one may directly solve a discrete dynamical system of the form $\mathbf{x}[t+1] = Q\mathbf{x}[t]$ over the space of orthogonal Q matrices. This is equivalent to solving the following constrained linear problem:

$$Q^* = \underset{Q}{\operatorname{argmin}} \; ||A - Q||_F, \; \textit{subject to } QQ^\intercal = I \;, \tag{1}$$

where $A = (X_{t+1}X_t^\intercal)(X_tX_t^\intercal)^{-1}$ is the least square solution to the unconstrained problem $\mathbf{x}[t+1] = A\mathbf{x}[t]$. The solution to the constrained optimization in Eq. (1) is given by the orthogonal matrix factor of celebrated *polar decomposition* of matrix A, namely $A = QP$. In summary, $I + M$ is an approximation to Q, which can be effectively computed using polar decomposition. Notably, if λ_k and v_k are an eigenvalue and eigenvector of matrix M, then $1 + \lambda_k$ and v_k are the approximate eigenvalue and eigenvector of matrix Q. That is, the jPCA eigenvectors span the same space as the eigenvectors of the polar decomposition factor Q.

2.2 A Probabilistic Model for jPCA

Here we show that an appropriately constrained complex linear dynamical system (CLDS) [2] is equivalent to the jPCA. The jPCA projection axes emerge as maximum-likelihood parameter estimations of the observation matrix in the CLDS, where the state captures the rotational dynamics within the latent space. Moreover, it provides a probabilistic interpretation of the

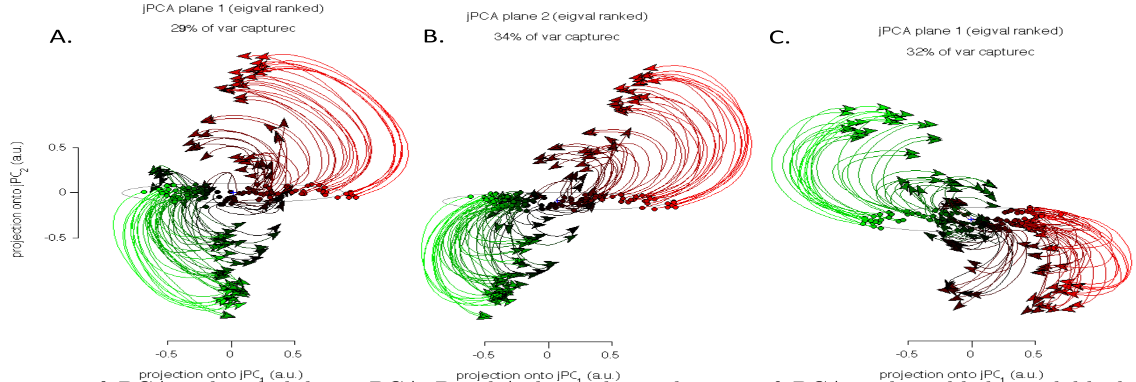

Figure 1: Comparison of jPCA and probabilistic jPCA. Panel A shows the application of jPCA to the publicly available data of Churchland et al. [1]. Panel B is the result of applying the probabilistic jPCA using the expectation-maximization algorithm. Panel C shows the application of the probabilistic jPCA to the original firing rate data – thus skipping the PCA step and instead setting the latent state dimension to the PCA dimension of six.

jPCA, with the extra flexibility to incorporate the state and measurement noise, priors on the involving matrices, and automatic determination of the state dimension (e.g., using ARD), among others. Let $x[t] \in \mathbb{R}^{K\times 1}$ be a column vector of (PCA transformed) firing rates, and define the complex linear dynamical system with rotational dynamics of the form

$$\begin{aligned} \mathbf{z}[t+1] &= D\mathbf{z}[t] + \mathbf{w}[t] \\ \mathbf{x}[t] &= V\mathbf{z}[t] + \mathbf{v}[t] \;, \end{aligned} \tag{2}$$

where the latent state $z[t] \in \mathbb{R}^{2L\times 1}$ is a vector of L pairs of complex conjugate latent factors, D a diagonal matrix with complex conjugate pair entries of the form $e^{j2\pi f}$ (with f being the frequency of oscillations), $\mathbf{w} \sim \mathcal{CN}(0, W)$ and $\mathbf{v} \sim \mathcal{CN}(0, V)$ ($\mathcal{CN}(.)$ denotes the multivariate complex normal distribution). It is straightforward to show that Eq. (2) can emerge from diagonalization of a dynamical system with rotational dynamics (i.e., an orthogonal transition matrix), and therefore is the generative model of the jPCA. Learning the parameters of this model, for instance via the expectation-maximization (EM) algorithm, and with the constraint that D is diagonal with complex conjugate pair elements of absolute value one (a linear least square problem), yields the jPCA bases as the columns of the observation matrix V.

3 Experimental Results

We applied the proposed probabilistic jPCA algorithm to the publicly available data of Churchland et al. [1] recorded from the M1 region of monkey motor cortex. We used the complex EM algorithm [2] for learning the parameters of the probabilistic model of Eq. (2), starting from a polar decomposition-based initialization. Figure 1 panels A and B show that the jPCA algorithm and the proposed probabilistic model yield slightly better results (in terms of explained variance) when the latent dimension is set to be equal to the observed data dimension (here the PCA-transformed firing rate data, $K = 6$). Figure 1 panel C shows the application of probabilistic jPCA to the raw multivariate firing rate data (218 dimensions), but using a latent dimension of order six ($L = 3$). Since the the probabilistic jPCA explicitly models the noise, it still succeeds in recovering the rotational dynamics within the population firing rate data. In addition, we have applied the proposed approach to synthetic motor population neuronal data (simulated from a balanced network with excitatory and inhibitory neurons [3]) and found similar observations (results not shown).

4 Conclusion

Probabilistic Bayesian modeling framework has provided us with a wealth of techniques for inference, model selection, learning, and handling of missing data. We provide a probabilistic interpretation of the jPCA technique, which may inspire developments of future techniques for analysis of neural population dynamics. An intriguing extension of this work is to combine the probabilistic state-space model of the jPCA with non-gaussian observation models of spiking activity (using the GLM-framework). This may in turn allow us to uncover the functional network connectivity responsible for emergence of rotational dynamics in neuronal population activity, and to provide insight into the underlying neural mechanism.

References

[1] M. M. Churchland, J. P. Cunningham, M. T. Kaufman, J. D. Foster, P. Nuyujukian, S. I. Ryu, and K. V. Shenoy, "Neural population dynamics during reaching," *Nature*, vol. 487(7405):51-6, 2012.

[2] L. Li and B. A. Prakash, "Time series clustering: Complex is simpler!" in *Proceedings of the 28th International Conference on Machine Learning (ICML-11)*, 2011, pp. 185–192.

[3] G. Hennequin, "Stability and amplification in plastic cortical circuits," Ph.D. dissertation, ÉCOLE POLYTECHNIQUE FÉDÉRALE DE LAUSANNE, 2013.