Journal of Neuroscience Methods 258 (2016) 1–15

Contents lists available at ScienceDirect

Journal of Neuroscience Methods

journal homepage: www.elsevier.com/locate/jneumeth

Computational Neuroscience

Extracting spatial–temporal coherent patterns in large-scale neural recordings using dynamic mode decomposition

Bingni W. Brunton[a,b,c,*], Lise A. Johnson[d,e], Jeffrey G. Ojemann[e], J. Nathan Kutz[b]

[a] Department of Biology, University of Washington, Seattle, WA 98195, USA
[b] Department of Applied Mathematics, University of Washington, Seattle, WA 98195, USA
[c] Institute for Neuroengineering, University of Washington, Seattle, WA 98195, USA
[d] Center for Sensorimotor Neural Engineering, University of Washington, Seattle, WA 98195, USA
[e] Department of Neurological Surgery, University of Washington, Seattle, WA 98195, USA

HIGHLIGHTS

- Dynamic mode decomposition (DMD) extracts dynamically coherent patterns from large-scale neuronal recordings.
- Multiple, distinct sleep spindle networks are identified by DMD as measured in subdural array recordings.
- Sleep spindle networks are characterized by different cortical distribution patterns, carrying frequencies and durations.

ARTICLE INFO

Article history:
Received 7 April 2015
Received in revised form
17 September 2015
Accepted 20 October 2015
Available online 31 October 2015

Keywords:
Dynamic mode decomposition
Spatiotemporal modes
Electrocorticography
Feature extraction
Sleep spindles

ABSTRACT

Background: There is a broad need in neuroscience to understand and visualize large-scale recordings of neural activity, big data acquired by tens or hundreds of electrodes recording dynamic brain activity over minutes to hours. Such datasets are characterized by coherent patterns across both space and time, yet existing computational methods are typically restricted to analysis either in space or in time separately.
New method: Here we report the adaptation of dynamic mode decomposition (DMD), an algorithm originally developed for studying fluid physics, to large-scale neural recordings. DMD is a modal decomposition algorithm that describes high-dimensional dynamic data using coupled spatial–temporal modes. The algorithm is robust to variations in noise and subsampling rate; it scales easily to very large numbers of simultaneously acquired measurements.
Results: We first validate the DMD approach on sub-dural electrode array recordings from human subjects performing a known motor task. Next, we combine DMD with unsupervised clustering, developing a novel method to extract spindle networks during sleep. We uncovered several distinct sleep spindle networks identifiable by their stereotypical cortical distribution patterns, frequency, and duration.
Comparison with existing methods: DMD is closely related to principal components analysis (PCA) and discrete Fourier transform (DFT). We may think of DMD as a rotation of the low-dimensional PCA space such that each basis vector has coherent dynamics.
Conclusions: The resulting analysis combines key features of performing PCA in space and power spectral analysis in time, making it particularly suitable for analyzing large-scale neural recordings.

1. Introduction

Advances in technology and infrastructure are delivering the capacity to record signals from brain cells in much greater numbers and at even faster speeds. This deluge of data is central to answering many critical open questions in neuroscience and motivates the continued development of computational approaches to analyze, visualize, and understand large-scale recordings of neural activity. Fortunately, the activity of complex networks of neurons can often be described by relatively few distinct patterns (for instance, Broome et al. (2006), Byron et al. (2009), Churchland et al. (2012) and Machens et al. (2010)). Identifying these spatial–temporal patterns enables the reduction of complex measurements through projection onto coherent structures, where it is tractable to build dynamical models and apply machine learning tools for pattern

* Corresponding author at: Department of Biology, University of Washington, Box 351800, Seattle, WA 98195, USA. Tel.: +1 2062219330.
E-mail address: bbrunton@uw.edu (B.W. Brunton).

http://dx.doi.org/10.1016/j.jneumeth.2015.10.010
0165-0270/

analysis. Here we introduce dynamic mode decomposition (DMD) as a novel approach to explore spatial–temporal patterns in large-scale neural recordings. The method combines well-characterized advantages from two of the most powerful data analytic tools in use today: power spectral analysis in time and principal components analysis (PCA) in space.

Measurements of neural activity from tens to hundreds of simultaneously recorded channels are traces in time that probe a network with complex dynamics; one principled way to make sense of such dynamic networks is with modal decomposition (Holmes et al., 1998). Modal decomposition has been successfully applied in almost every discipline of science and engineering because it makes tractable the analysis of very high-dimensional data, reducing them to combinations of relatively few distinct patterns, or *modes*. A particularly popular modal decomposition tool is PCA, which derives modes ordered by their ability to account for energy or variance in the data (Jolliffe, 2005). PCA has already been widely applied in the study of high-dimensional biological systems; however, it suffers from a few well known drawbacks. In particular, PCA is a static technique and does not model temporal dynamics of time-series data explicitly, so it often performs poorly in reproducing dynamic data, such as recordings of neural activity.

Neural dynamics are well known to be characterized by dynamic oscillations at many frequency bands, which are implicated in a variety of neural functions (Buzsáki and Draguhn, 2004; Fries, 2005; Raghavachari et al., 2001; Uhlhaas and Singer, 2010). Most tools analyzing the frequency content of a signal are related to the Fourier transform, which transforms time-varying signals into a spectrum in the frequency domain. Importantly, the power spectrum can be computed efficiently using the fast Fourier transform (FFT) algorithm (Welch, 1967), whose efficient implementation has contributed to its ubiquitous use. One example of a modal decomposition in time that goes beyond the Fourier transform is empirical mode decomposition (EMD), which computes intrinsic oscillatory modes from time-varying data (Huang et al., 1998). EMD has been used to analyze neural data, including cortical local field potential (Liang et al., 2005) and EEG (Sweeney-Reed and Nasuto, 2007). There are several extension of frequency-domain analyses that also support spatial structures (Rehman and Mandic, 2009; Rudrauf et al., 2006)

A relatively new modal decomposition method is DMD (Rowley et al., 2009; Schmid and Sesterhenn, 2008; Schmid, 2010). DMD was developed initially to study experiments and simulations in the fluid mechanics community, where it was introduced to reduce very high-dimensional dynamic data into relatively few coupled spatial–temporal modes. Importantly, it has been shown that DMD is related to Koopman spectral analysis, motivating its usefulness in characterizing dynamics of nonlinear systems (Budišić et al., 2012; Rowley et al., 2009). Beyond fluid mechanics, DMD has recently been applied to the fields of robotics (Berger et al., 2015) and disease modeling (Proctor and Eckhoff, 2015). In the context of analyzing neural recordings, DMD modes can be thought of as coherent structures in the neural activity.

1.1. Summary of computational developments

In Section 2, we describe a set of adaptations of the DMD that make it useful in the extraction of spatial–temporal patterns from neural recordings. The base DMD algorithm is given in Section 2.1, where we also note its relationship to more established methods and compare DMD modes to PCA modes for an illustrative synthetic dataset. Notably, in contrast to experiments and simulations in fluid mechanics, neural recordings often have fewer measurements (channels of electrodes) than time snapshots, so in Section 2.2 we describe the construction of an augmented data matrix. We give intuition and recommendations for how to choose a set of parameters such that the extracted DMD modes are interpretable. For instance, DMD modes are useful as features in machine learning algorithms that undercover stereotyped patterns in the data, an attribute we leverage for a specific example described in Section 3.3. Examples of ECoG data decomposed by DMD are given in Section 2.3. Section 2.4 describes the DMD spectrum, which has units easily interpretable in comparison with traditional power spectral analyses. Next in Section 2.5, we characterize to what extent spatial–temporal modes extracted by DMD from human subdural recordings are robust to noise and subsampling.

For very large datasets whose dimensionality strains typical computing resources, DMD may be readily implemented using standard linear algebra routines to take advantage of cluster computing (for example, see Freeman et al. (2014)). We suggest that DMD may be useful in understanding spatial–temporal coherent patterns in data of escalating scale in neuroscience, including non-invasive and invasive measurements such as functional MRI, MEG, neurophysiological recordings with electrode arrays, and optical imaging of neural activity.

1.2. Summary of experimental demonstrations

To demonstrate DMD's applicability to large-scale neural recordings, we analyzed sub-dural electrode array recordings from human subjects in two different contexts.

1.2.1. Sensorimotor maps

First, we validated the DMD approach to derive sensorimotor maps based on a simple movement task. Our sensorimotor maps show statistically significant changes in activation over the sensorimotor cortex in two frequency ranges. These changes are distinct for movements of the hand and tongue, and they are consistent with results previously described by Miller et al. (2007).

1.2.2. Sleep spindle networks

Next, we leveraged DMD in combination with unsupervised clustering techniques to detect and characterize spindle networks present during sleep; a method to automatically extract these networks had not been described previously in the literature. Sleep spindles are distinctive, transient oscillations around 14 Hz that are characteristic of non-rapid eye movement (NREM) sleep, and their presence is commonly used to classify sleep stages (De Gennaro and Ferrara, 2003). Spindles have been the subject of scientific investigation since the early 1930s and their mechanisms of generation are now quite well understood (Steriade et al., 1993). In brief, sleep spindles oscillations are generated in the thalamus and their electrographic signature arises from thalamacortical connections. Even so, the role these transient oscillatory events play in brain function remains unclear. A line of evidence suggests that sleep spindles facilitate the consolidation of recently acquired memories (Clemens et al., 2005; Eschenko et al., 2006; Gais et al., 2002; Johnson et al., 2012). This hypothesis is supported by recent work demonstrating that sleep spindles can be locally, rather than globally, synchronous events (Johnson et al., 2012; Nir et al., 2011).

Historically, sleep spindles have been scored by experts on scalp EEG data. Spindles vary in amplitude, duration, central frequency, and often concur with other regularly observed sleep features. Automated detection algorithms typically rely on band-pass filtering the signal followed by an amplitude threshold on some moving average window (for instance, Ray et al. (2010) and Schimicek et al. (1994)). Recently, a number of these algorithms were evaluated against experts and crowd-sourced spindle detectors (Warby et al., 2014). It is important to point out that all of these approaches only address spindle detection. The reliable identification and characterization of spatial networks of electrodes showing synchronous spindle activity has remained a challenge. The structure of sleep