Learning dynamic functional connectivity networks from infant magnetoencephalography data

Inferring brain interactions is an important problem in neuroscience with far-reaching applications, including understanding neurological disorders such as autism. An open area of research is studying how brain connectivity evolves in the early stages of learning and development. We approach this problem through the lens of *functional connectivity* which describes the temporal correlation between patterns of activity in two brain regions, and is an important measure for studying how the brain processes information [1]. Specifically, we are interested in learning functional connections from neuroimaging recordings collected via magnetoencephalography (MEG), a popular technique that boasts high temporal resolution with fairly good spatial resolution [2].

Our particular case involves two key modifications to the original setup of learning a functional connectivity network from MEG data. The first is that we are interested in dynamic connectivity, as studying how these connections change over time could yield insight that wouldn't be possible from studying static networks. The high temporal resolution of MEG recordings allows us to study the dynamics of the data even over short time intervals. Additionally, we are interested in learning these connections from infant recordings to improve our understanding of how the brain changes during early learning and development. Analyzing infant data poses a number of additional problems, including (i) high intra-subject variability due to noise from head size and head movements during recording sessions as well as (ii) high inter-subject variability as a result of large variation in anatomy and early brain development across subjects.

A promising technique for learning dynamic functional connections from MEG data comes from a recently proposed state-space modeling approach, which learns a linear dynamical system over estimates of cortical signals and uses the learned dynamics as a measure of directed connectivity between particular regions of interest, or ROIs [3]. Specifically, the ROI signals follow a time varying vector autoregressive model

$$u_t = A_t u_{t-1} + \epsilon_t \qquad \epsilon_t \sim N(0, Q), \tag{1}$$

where $u_t \in \mathbb{R}^p$ describes the activity across p ROIs at timepoint t, and $A_t \in \mathbb{R}^{p\times p}$ describes the dynamics within and between ROIs at time t. Given the ROI signals, the sensor recordings are obtained according to

$$y_t = C u_t + \eta_t \qquad \eta_t \sim N(0, R) \tag{2}$$

where $y_t \in \mathbb{R}^n$ is a vector of recorded values across n MEG sensors, $C \in \mathbb{R}^{n\times p}$ is a known matrix that describes the transformation from p ROI locations to n sensors, and $R \in \mathbb{R}^{n\times n}$ is a noise covariance matrix that is a combination of variation in ROI activity and noise from the sensors themselves. Previous work developed this state-space modeling approach to simultaneously learn connectivity while performing source localization of signals from MEG sensor recordings [3]. Two significant advantages of this method over traditional approaches are: (1) avoiding the two-step procedure that first involves an ill-posed mapping from sensor space to source space, introducing numerous artifacts in the source-localized signals used to infer connectivity in the second stage; and (2) capturing directed connectivity that changes over short time scales.

Although a promising first step, Yang et al. [3] only apply the state-space model to data from a single subject performing a visual task. Such an approach is ill-suited to studies of auditory behavior where the interactions are much weaker, necessitating integrating data from multiple subjects to boost the signal-to-noise ratio. This poses a non-trivial extension since different subjects can have different channels corrupted (which are typically discarded prior to analysis). To handle this situation, we extend the model in Eqs. (1) and (2) to learn a set of shared dynamics while making use of subject-specific structural and sensor noise information. Preliminary experiments on simulated data indicate that this state-space modeling approach works well when incorporating structural information from adult subjects. Performing the same analysis with *identical simulated dynamics*—but now modeled using *infant structural information*—demonstrates that integrating across subjects is even more critical in this case; see Fig. 1. We see that when analyzing a single infant subject, the inferred directed connections are much further from the true connections than in the case of a single adult subject. Considering multiple infant subjects mitigates this issue. The reason for these differences relative to the adult analysis is that infant structural information is inherently noisier and thus more difficult to learn known connectivity from, necessitating sharing of information between subjects.

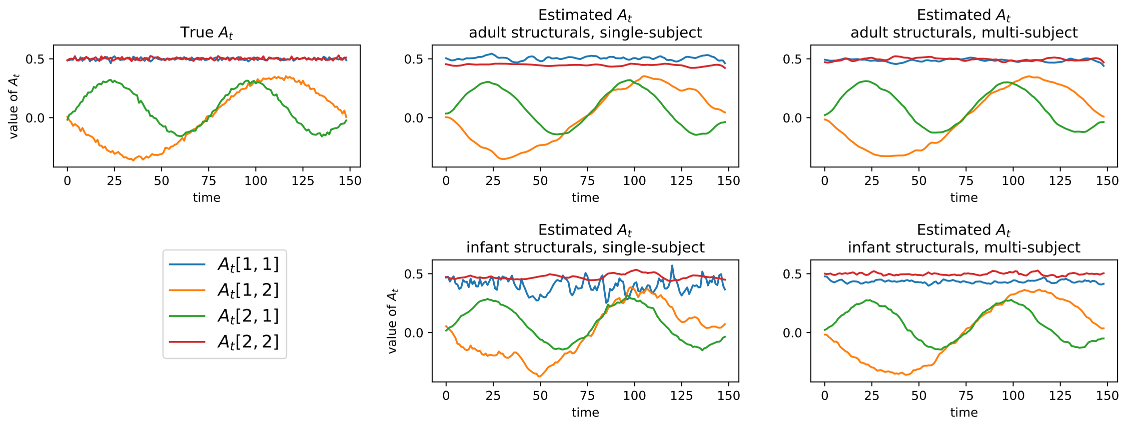

Figure 1: Results for simulated dynamics, where the plots show (from left to right, top to bottom): the true dynamics; learned dynamics using adult structural information, one subject; learned dynamics using adult structural information, two subjects; learned dynamics using infant structural information, one subject; and learned dynamics using infant structural information, two subjects. Despite using the same data and region labels corresponding to the infant brain, the learned dynamics corresponding to the infant structurals are noisier than for the adult structurals. By sharing dynamics across two infant subjects while maintaining independent structural and sensor noise information, we can better recover the true connectivity as compared to the single-subject infant case.

Our synthetic experiments demonstrate that infant structural information is inherently noisier than that of adults, resulting in the need to share information across subjects to better separate the meaningful signal from the noise. However, sharing information without modeling inter-subject variability can fail to uncover meaningful patterns that are specific to certain subjects, but not globally expressed. This motivates the development of a hierarchical model which can learn some global set of dynamics that includes shared features across all subjects, while also learning subject-specific dynamics that are instances of the global trend but also contain useful intra-subject variation. We hope to pursue the development of such a hierarchical model in our future work, and apply it to newly collected infant MEG data to attempt to learn these networks while better modeling issues specific to the infant brain. The state-space model explored herein provides a natural framework upon which to build such extensions.

References

[1] B. Horwitz. The elusive concept of brain connectivity. *NeuroImage*, 19(2):466–470, 2003.

[2] P. C. Hansen, M. L. Kringelbach, and R. Salmelin. *MEG: An Introduction to Methods*. Oxford University Press, 2010.

[3] Y. Yang, E. M. Aminoff, M. J. Tarr, and R. E. Kass. A state-space model of cross-region dynamic connectivity in MEG/EEG. *30th Conference on Neural Information Processing Systems*, 2016.