Statistical Analysis of Neuronal Population Codes for Encoding Acute Pain

Zhe Chen and Jing Wang

Abstract— **To date most pain studies have focused on spinal cord or peripheral pathways. However, a complete understanding of pain mechanisms requires the study of neocortex. Using an animal model of acute pain, we investigate neural codes for pain at both single-cell and population levels. We propose a statistical framework, rooted in state space analysis, for analyzing neural ensembles recorded from the rat primary somatosensory cortex (S1) and anterior cingulate cortex (ACC) during a laser pain stimulation protocol. The state space analysis allows us to uncover a latent state process that drives the observed ensemble spike activity, and to further detect the "neuronal threshold" for pain on a single or multiple-trial basis.**

I. Introduction

Pain is the most common reason Americans access the health care system, the leading cause of disability, and a major contributor to health care costs. In basic and translational pain research, animal models are pivotal for understanding the mechanisms of pain and for the development of effective therapy for its optimal management [11]. Research over the last 50 years has resulted in a better understanding of spinal and peripheral mechanisms for pain. Nevertheless, brain circuits that regulate pain, especially the affective component of pain, remain incompletely understood. Such understanding, however, is vital to the knowledge of how the brain processes sensory and affective information, and it could lead to novel therapeutic strategies.

We use a rat model to study acute conceptive pain. Acute *nonciceptive* pain is caused by stimulation of peripheral nerve fibers that respond only to stimuli exceeding harmful intensity, e.g., thermal, mechanical, and chemical nociceptors; whereas *neuropathic* pain is caused by damage or disease affecting the somatosensory nervous system, which may be associated with abnormal sensation called dysesthesia or pain from normally non-painful stimuli (allodynia). The experience of pain in the brain derives from the combined activity of a large network of neural structures. Among many brain regions, the primary somatosensory cortex (S1) and anterior cingulate cortex (ACC) are two of the most studied areas for pain-related perception. The S1 encodes the *sensory* component of pain, whereas the ACC has been widely hypothesized to encode the *aversive* component of pain [2]. Furthermore, previous studies have established that activities in the S1 and the ACC are altered by pain stimulus [10], [19], [17], [8]. However, very few studies were devoted to population analysis in awake freely moving animals [19]. A great challenge is to identify relevant neural ensembles that drive pain behavior. The identification of such neuronal ensembles could effectively provide a "neural code" for pain. It is even more important to design algorithms that enable us to read out the neural code or pain based on the ensemble spike activity in a single trial, which will provide a basis for future brain-machine interface (BMI) applications.

Support from the US NSF-CRCNS grant IIS-130764 (Z.C.) and NIH grants K08-GM1026911 and R01-GM115384 (J.W.).

Z. Chen is with the Department of Psychiatry, and Department of Neuroscience and Physiology, New York University School of Medicine, New York, NY, USA. J. Wang is with the Department of Anesthesiology, and Neuroscience and Physiology, New York University School of Medicine, New York, NY, USA (Email: zhe.chen3@nyumc.org)

In this paper, we propose a statistical framework for detecting neuropathic pain based on ensemble spike activity recorded from rat ACC and S1 areas in response to laser noxious stimuli. Our approach employs state space analysis [3], [4] and variational inference for a latent Gaussian state space model (SSM) with Poisson observations. The state space approach allows us to perform smoothing on the latent variable that drives single-trial temporal firing rate of neuronal ensembles, and to estimate the confidence intervals. We verify our method with experimentally recorded ensemble spike data and discuss its extension for online applications.

II. Neural Codes of Acute Pain

From the perspective of neural coding of pain, we ask whether the ACC and/or S1 neuronal spike activities represent or encode the pain stimulus, at both single neuron and population levels. Typically, single neuronal spike activity is recorded from multiple trials of stimulus delivery, from which one can plot the PSTH (peri-stimulus time histogram). To identify a "pain-responsive" unit, we use the following statistical criteria: (i) Based on the spike rates confidence intervals (shaded area), if the lower bound of the post-stimulus firing rate is above the upper bound of the pre-stimulus (baseline) firing rate, we will conclude there is a significant firing rate increase and the unit is a "positive responder". (ii) Alternatively, if the upper bound of the post-stimulus firing rate is below the lower bound of the pre-stimulus firing rate, then there is a significant firing rate decrease and the unit is a "negative responder". (iii) The first time point that the spike rate meets the first or second criterion will be used to define the latency. Figure 1 presents three representative examples of both positive and negative responders of pain stimulus. In our empirical observations, we have noticed that positive, negative and non-responsive units are simultaneously present in the ACC or S1 neuronal population (which was not reported in [10], [19]). Furthermore, the latency of individual neuronal responses vary, even under the same pain stimulus intensity.

Note that the above criteria have several limitations. First, the PSTH and confidence intervals rely on multiple trial averages; therefore, it is not applicable to single-trial analysis. Second, there is no temporal smoothing in the PSTH or Z-score. Note that although the Z-score can be computed based on time average on the single-trial firing rate, it is typically very noisy due to lack of smoothing. In addition, we would also like to obtain the uncertainty estimate of the single-trial Z-score. To address these issues, we propose a statistical analysis framework to analyze the population codes for encoding acute nonciceptive pain. Specifically, we develop a state space model (SSM) to detect the "neuronal threshold" for acute pain response on a single-trial basis.

At the single cell level, we often observe trial-to-trial variability of neuronal firing (see the raster plot). Similarly, at the population level, we also see a large variability between single trials, for either S1 or ACC neuronal ensembles. The sources of "neural" variability may be caused by (i) physical properties of the laser and skin conductance, or (ii) animal's internal brain state, or both.

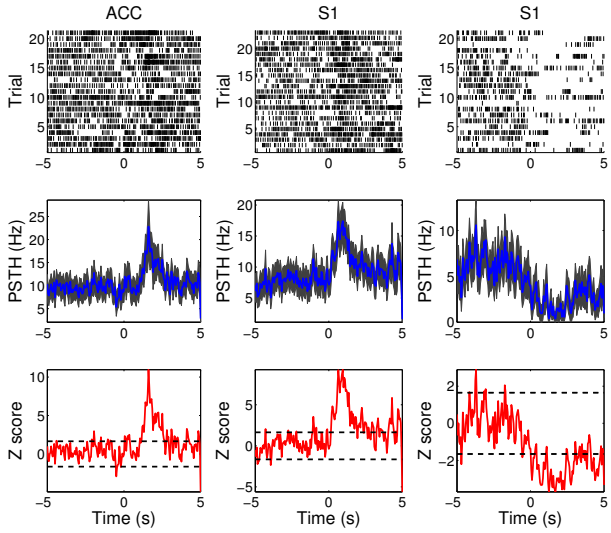

Fig. 1. (A) Spike raster plots of one ACC neuron and two S1 neurons recorded from behaving rats. Trials are aligned by the onset of paw withdrawal time. (B) PSTH (solid line) and confidence intervals (shaded area), bin size: 50 ms. (C) Z-score of the firing rate. The Z-score is defined by $z = \frac{x-\mu}{\sigma}$, where μ and σ defines the mean and standard deviation of baseline firing rate (averaged over time and trials). Horizontal dashed lines mark ± 1.65 (equivalent P-value 0.05).

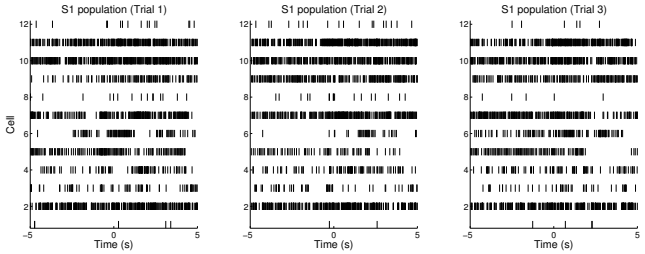

Fig. 2. Examples of spike activity from twelve S1 neuronal ensembles in three trials. Time zero marks the onset of paw withdrawal.

III. Methods

A. State Space Analysis

Let $k = 1, \ldots, T$ denote the discrete-time index of a univariate or multivariate time series, which can be continuous or discrete (e.g., spike count data) with a predefined bin size. Let $\mathbf{y}_k = [y_{1,k}, \ldots, y_{C,k}]^\top$ denote a C-dimensional neuronal population vector, with each element denoting the observed neuronal spike count.

Motivated by the work of [16], [1], we develop a latent Gaussian model with Poisson observations. We assume that the latent variable $\mathbf{z}_k \in \mathbb{R}^m$ represents the unobserved common input that drives the neuronal population firing and it follows a Markovian process

$$\mathbf{z}_k = \mathbf{A}\mathbf{z}_{k-1} + \boldsymbol{\epsilon}_k \quad (1)$$

$$\boldsymbol{\eta}_k = \mathbf{C}\mathbf{z}_k + \mathbf{d} \quad (2)$$

$$\mathbf{y}_k \sim Poisson\Big(\exp(\boldsymbol{\eta}_k)\Big) \quad (3)$$

where $\boldsymbol{\epsilon}_k \in \mathcal{N}(\mathbf{0}, \mathbf{Q})$ and $\mathbf{z}_1 \in \mathcal{N}(\mathbf{0}, \mathbf{Q}_1)$ specify temporal priors on the latent process. The state equation (1) is a first-order autoregressive (AR) model, with a state transition matrix $\mathbf{A} \in \mathbb{R}^{m \times m}$, and the observation equation is a generalized linear model (GLM) that employs the exponential link function through $\boldsymbol{\eta}_k$. The parameters $\mathbf{d}, \mathbf{C}$ or $\boldsymbol{\eta}_k$ are unconstrained, and $\mathbf{C}$ can be a full (C-by-m) matrix, which allows the possibility that the individual neuronal firing is influenced by the latent population dynamics.

In terms of model identifiability, the latent process is subject to scale and sign ambiguity. If the latent state is univariate (i.e., $m = 1$), we may compute a scaled variable from $\hat{z}_{1:T}$

$$Z = \frac{z - \text{mean of } z_{\text{baseline}}}{\text{SD of } z_{\text{baseline}}} \quad (4)$$

which is known as the Z-score. Since the Z-score is standard normally distributed, we convert it to the one-tailed P-value: [1]

$$P(Z > z) = 1 - P(Z \leq z) = 1 - \int_{-\infty}^{z} \frac{1}{\sqrt{2\pi}} e^{\frac{-u^2}{2}} du \quad (5)$$

When $Z > 1.65$ or $Z < -1.65$, it is concluded that the population significantly increases or decreases its firing (one-sided P=0.05) with respect to the pain stimulus, respectively. The normalization in (4) resolves the scale ambiguity of z_k, and the sign ambiguity can be resolved from the algebraic signs of z_k and $\mathbf{C}$.

When the latent state is multivariate (i.e., $m > 1$), we can similarly derive the P-value from the cumulative probability of the multivariate normal distribution. In MATLAB, the one-sided P-value is $1 - \texttt{mvncdf}(\mathbf{z}, \boldsymbol{\mu}, \boldsymbol{\Sigma})$, where $\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}$ denote the time-averaged mean and covariance of latent variable in baseline.

B. Inference

Let $\boldsymbol{\Theta}$ denote all unknown parameters (i.e., $\boldsymbol{\Theta} = \{\mathbf{A}, \mathbf{C}, \mathbf{d}, \mathbf{Q}, \mathbf{Q}_1\}$). Because of non-Gaussian likelihood, the E-step of traditional expectation-maximization (EM) algorithm will be intractable. Therefore, some Gaussian approximation methods can be employed [15], [16], [1], [9]. Specifically, we employ an efficient dual variational inference algorithm [9], where the unknown parameters are initialized based on a subspace method [13]. Specifically, let $\boldsymbol{z} = [\mathbf{z}_1, \ldots, \mathbf{z}_T]$ denote the joint latent state variable accumulated in time; the goal of variational inference is to optimize the variational lower bound of the marginal log-likelihood

$$\begin{aligned} \log p(\mathbf{y}_{1:T}) &= \log \int \prod_{k=1}^{T} p(\mathbf{y}_k|\boldsymbol{\eta}_k) p(\boldsymbol{z}) d\boldsymbol{z} \\ &= \log \int q(\boldsymbol{z}) \frac{\prod_{k=1}^{T} p(\mathbf{y}_k|\boldsymbol{\eta}_k) p(\boldsymbol{z})}{q(\boldsymbol{z})} d\boldsymbol{z} \\ &\geq \mathbb{E}_{q(\boldsymbol{z})}\left[\log \frac{\prod_{k=1}^{T} p(\mathbf{y}_k|\boldsymbol{\eta}_k) p(\boldsymbol{z})}{q(\boldsymbol{z})}\right] \end{aligned} \quad (6)$$

where $\boldsymbol{\eta}_k = [\eta_{1,k}, \ldots, \eta_{C,k}]$ is the natural parameter, and $p(\boldsymbol{z}) = \mathcal{N}(\mathbf{0}, \boldsymbol{\Lambda}^{-1})$ denotes the prior with a tri-diagonal covariance structure of size $(mT \times mT)$

$$\boldsymbol{\Lambda} = \begin{pmatrix} \mathbf{Q}_1^{-1} + \mathbf{A}\mathbf{Q}^{-1}\mathbf{A}^\top & \mathbf{A}^\top\mathbf{Q}^{-1} & & \\ \mathbf{Q}^{-1}\mathbf{A} & \mathbf{Q}^{-1} + \mathbf{A}\mathbf{Q}^{-1}\mathbf{A}^\top & \mathbf{A}^\top\mathbf{Q}^{-1} & \\ & \ddots & \ddots & \ddots \end{pmatrix}$$

and $q(\boldsymbol{z}) = \mathcal{N}(\boldsymbol{z}|\boldsymbol{m}, \boldsymbol{V})$ denotes a variational Gaussian posterior with mean $\boldsymbol{m}$ and covariance $\boldsymbol{V}$.

[1] In MATLAB (The MathWorks Inc., Natick, MA) implementation, we have the one-sided P-value as $1 - \texttt{normcdf}(Z, 0, 1)$, where `normcdf` denotes a normal cumulative distribution function.

In variational inference, we optimize the lower bound of $\log p(\mathbf{y}_{1:T})$, which has the following form

$$\mathcal{L}(\boldsymbol{m},\boldsymbol{V}|\boldsymbol{\Theta}) = \frac{1}{2}\Big(\log|\boldsymbol{V}|) - \mathtt{tr}[\boldsymbol{\Lambda}\boldsymbol{V}] - \boldsymbol{m}^\top\boldsymbol{\Lambda}\boldsymbol{m}\Big) + \sum_{k=1}^{T}\mathbb{E}_{q(\boldsymbol{z})}[\log p(\mathbf{y}_k|\boldsymbol{\eta}_k)] \quad (7)$$

In the E-step of the EM algorithm, the dual variational inference of $\{\boldsymbol{m},\boldsymbol{V}\}$ uses a convex dual optimization method [9] (with T variational parameters) embedded with forward-backward Kalman smoothing (which computes the sufficient statistics $\hat{\mathbf{z}}_{k|T} = \mathbb{E}\big[\mathbf{z}_k|\mathbf{y}_{1:T}\big]$, $\mathbf{P}_{k|T} = \mathbb{E}\big[(\mathbf{z}_k - \hat{\mathbf{z}}_{k|T})(\mathbf{z}_k - \hat{\mathbf{z}}_{k|T})^\top\big]$ and $\mathbf{P}_{k,k-1|T} = \mathbb{E}\big[(\mathbf{z}_k - \hat{\mathbf{z}}_{k|T})(\mathbf{z}_{k-1} - \hat{\mathbf{z}}_{k-1|T})^\top\big]$). At the backward step, it computes

$$\mathbf{W}_k = \mathbf{P}_{k|k}\mathbf{A}^\top\mathbf{P}_{k+1|k}^{-1} \quad (8)$$

$$\hat{\mathbf{z}}_{k|T} = \hat{\mathbf{z}}_{k|k} + \mathbf{W}_k\Big(\hat{\mathbf{z}}_{k+1|T} - \hat{\mathbf{z}}_{k+1|k}\Big) \quad (9)$$

$$\mathbf{P}_{k|T} = \mathbf{P}_{k|k} + \mathbf{W}_k\Big(\mathbf{P}_{k+1|T} - \mathbf{P}_{k+1|k}\Big)\mathbf{W}_k^\top \quad (10)$$

$$\mathbf{P}_{k,k-1|T} = \mathbf{W}_{k-1}\mathbf{P}_{k|T} \quad (11)$$

In the M-step, we optimize $\{\mathbf{C},\mathbf{d}\}$ using Newton's method to solve $\frac{\partial\mathcal{L}}{\partial\mathbf{C}} = 0$ and $\frac{\partial\mathcal{L}}{\partial\mathbf{d}} = 0$, and estimate $\{\mathbf{A},\mathbf{Q},\mathbf{Q}_1\}$ with a closed-form solution as follows

$$\hat{\mathbf{A}} = \Big(\sum_{k=2}^{T}\mathbb{E}\big[\mathbf{z}_k\mathbf{z}_{k-1}^\top|\mathbf{y}_{1:T}\big]\Big)\Big(\sum_{k=2}^{T}\mathbb{E}\big[\mathbf{z}_{k-1}\mathbf{z}_{k-1}^\top|\mathbf{y}_{1:T}\big]\Big)^{-1} \quad (12)$$

$$\hat{\mathbf{Q}} = \frac{1}{T-1}\sum_{k=2}^{T}\Big(\mathbb{E}\big[\mathbf{z}_k\mathbf{z}_k^\top|\mathbf{y}_{1:T}\big] - \hat{\mathbf{A}}\mathbb{E}\big[\mathbf{z}_k\mathbf{z}_{k-1}^\top|\mathbf{y}_{1:T}\big]\Big) \quad (13)$$

$$\hat{\mathbf{Q}}_1 = \mathbb{E}\big[\mathbf{z}_1\mathbf{z}_1^\top|\mathbf{y}_{1:T}\big] - \hat{\mathbf{z}}_{1|T}\hat{\mathbf{z}}_{1|T}^\top = \mathbf{P}_{1|T} + \hat{\mathbf{z}}_{1|T}\hat{\mathbf{z}}_{1|T}^\top \quad (14)$$

We run the EM algorithm for a fixed number of iterations or the incremental increase of objective function is less than 10^{-4}. The algorithm often converges within a few hundreds of iterations.

Notes: (i) The above inference procedure is equally applicable to single-trial or multiple-trial data, in the latter stage the log likelihood will be summed over many trials. (ii) To select the dimensionality of $\mathbf{z}_k$, in the case of multiple trials we use cross-validation; whereas in the case of single trials, we can use the Bayesian information criterion (BIC). In most single trials, it is found that $m = 1$ yields the best result. (iii) For $m > 1$, there are sign, scale, and permutation ambiguity for displaying $\hat{\mathbf{z}}_k$.

C. Simplified Models

Modeling multi-unit activity: We can sum up all spike activity of sorted single units and obtain the multi-unit activity (MUA). Since there are potentially positive and negative responders in population, this strategy will only work if the majority of units show significant increase or decrease in firing rate and share a similar latency statistic. Since the sum of Poisson variables is still Poisson, we can then use a simplified univariate SSM

$$z_k = az_{k-1} + \epsilon_k \quad (15)$$

$$\sum_{c=1}^{C} y_{c,k} \sim Poisson\Big(\exp(cz_k + d)\Big) \quad (16)$$

When C is large, the Gaussian approximation will become more accurate due to the law of large numbers.

Linear Gaussian SSM: We can approximate Poisson observations with half-rectified Gaussian observations, which yields the linear dynamical systems (LDS). For inference, we use the standard Kalman smoothing and EM algorithm [6]. This holds for both univariate/multivariate state and observation equations.

D. Quantitative Assessment

In batch data analysis, assuming that the parameters in $\boldsymbol{\Theta}$ are stationary across trials, and the trial variability is induced by latent process $\mathbf{z}_{1:T}$; we can quantify the between-trials variability by computing the averaged Kullback-Leibler (KL) divergence

$$\begin{aligned} D_{\mathrm{KL}} &= \frac{1}{2}\mathrm{KL}\big(\mathbf{z}_{1:T,i}\|\mathbf{z}_{1:T,j}\big) + \frac{1}{2}\mathrm{KL}\big(\mathbf{z}_{1:T,j}\|\mathbf{z}_{1:T,i}\big) \\ &= \frac{1}{2T}\sum_{k=1}^{T}\mathrm{KL}\big(\mathbf{z}_{k,i}\|\mathbf{z}_{k,j}\big) + \mathrm{KL}\big(\mathbf{z}_{k,j}\|\mathbf{z}_{k,i}\big) \\ &= \frac{1}{4T}\sum_{k=1}^{T}\Big(\mathrm{tr}\Big(\mathbf{P}_{k|T,j}^{-1}\mathbf{P}_{k|T,i}\Big) + \mathrm{tr}\Big(\mathbf{P}_{k|T,i}^{-1}\mathbf{P}_{k|T,j}\Big) - 2m \\ &\quad + (\mathbf{z}_{k|T,j} - \mathbf{z}_{k|T,i})^\top\mathbf{P}_{k|T,j}^{-1}(\mathbf{z}_{k|T,j} - \mathbf{z}_{k|T,i}) \\ &\quad + (\mathbf{z}_{k|T,i} - \mathbf{z}_{k|T,j})^\top\mathbf{P}_{k|T,i}^{-1}(\mathbf{z}_{k|T,i} - \mathbf{z}_{k|T,j}) \\ &\quad + \log\Big(\frac{\det|\mathbf{P}_{k|T,j}|}{\det|\mathbf{P}_{k|T,i}|}\Big) + \log\Big(\frac{\det|\mathbf{P}_{k|T,i}|}{\det|\mathbf{P}_{k|T,j}|}\Big)\Big) \end{aligned} \quad (17)$$

where the subscripts i and j denote the trial indices.

For multiple trials, we compute the averaged "leave-one-out" log marginal likelihood. Given unseen ensemble spike data of a new trial, the log marginal likelihood of new data $\tilde{\mathbf{y}}$ is given by

$$\log p(\tilde{\mathbf{y}}_{1:T}|\boldsymbol{\Theta}) = \log\int p(\tilde{\mathbf{y}}_{1:T},\mathbf{z}_{1:T}|\boldsymbol{\Theta})d\mathbf{z}_{1:T} \quad (18)$$

$$= \sum_{k=1}^{T}\log p(\tilde{\mathbf{y}}_k|\tilde{\mathbf{y}}_{1:k-1}) = \sum_{k=1}^{T}\log\int p(\tilde{\mathbf{y}}_t|\mathbf{z}_k)p(\mathbf{z}_k|\tilde{\mathbf{y}}_{1:k-1})d\mathbf{z}_k$$

In the case of linear Gaussian SSM, equation (18) is analytically integrable. For the Poisson observation model, we can use Laplace approximation (e.g., [14]) or Monte Carlo approximation for (18)

$$\log p(\tilde{\mathbf{y}}_{1:T}|\boldsymbol{\Theta}) \approx \frac{1}{N_p}\sum_{i=1}^{N_p}\sum_{k=1}^{T}\log p(\tilde{\mathbf{y}}_k|\mathbf{z}_k^{(i)}) \quad (19)$$

where $\mathbf{z}_k^{(i)} \sim \mathcal{N}(\hat{\mathbf{z}}_{k|k-1},\mathbf{P}_{k|k-1})$ are i.i.d. samples drawn from $p(\mathbf{z}_k|\tilde{\mathbf{y}}_{1:k-1})$ estimated from an online Kalman filter.

IV. Experimental Data and Results

A. Neurophysiology

Tetrodes were constructed from four twisted 12.7 μm polyimide coated microwires (Sandvik) and mounted in an 8 tetrode VersaDrive (Neuralynx). A spacing distance of 0.508 mm separated tetrodes aligned in a 1.35 mm diameter circle. Electrode tips were plated with gold to reduce electrode impedances to 100-500 kΩ at 1 kHz. As described previously, rats were anesthetized with isoflurane (1.5-2%). The skull was exposed and a 2.5-mm diameter hole was drilled above the target region. A durotomy was performed before tetrodes were slowly lowered unilaterally into the ACC or S1 with the stereotaxic apparatus. All procedures in this study were performed in accordance with the NYU School of Medicine IACUC and the NIH guideline.

Noxious stimulation via a 473 nm blue diode-pumped solid-state laser (150-250 mW intensity) was applied to plantar surface of the hind paw contralateral to the brain recording site in freely moving rats. All recording sessions consisted on average 30 trials with variable inter-trial intervals (45$\pm$15 s). Using video (30 frame/s) tracking, the onset of noxious pain was identified, indicated by the paw withdrawal. Trials were aligned to the initiation of paw

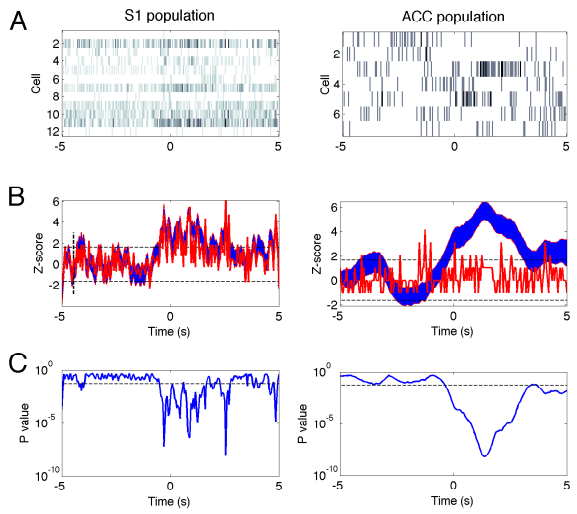

Fig. 3. Single-trial analysis based on S1 population (*left*, 12 units) and ACC population (*right*, 7 units). (A) Population spike count between $[-5, 5]$ s, with 0 indicating paw withdrawal. Bin size 50 ms. (B) Estimated Z-score from the latent state variable ($m = 1$, equation 4). Horizontal dashed lines mark the thresholds of significant zone. Vertical dashed line marks the laser onset. Blue shaded area marks the confidence intervals, and red curve marks the Z-score computed from multi-unit spike count directly. (C) Equivalent P-value derived from the Z-score (equation 5).

withdrawal. Behavioral response latency was determined from the time of laser onset to withdrawal.

Before recording, animals were given a 30-min period to habituate to a plastic recording chamber ($38 \times 20 \times 25$ cm^3). Tetrodes were lowered in steps of 60 μm before each session recording. Rats were connected to recording equipment (Open Ephys) via an RHD2132 amplifier board and Serial Peripheral Interface cable (Intan Technologies). Signals were monitored and recorded from 32 low-noise amplifier channels at 30 kHz, band-passed filtered (0.3 to 7.5 kHz). Spike sorting and data analysis were performed using Offline Sorter (Plexon), NeuroExplorer (Nex Technologies), and MATLAB (MathWorks). Clusters were identified as spikes using a number of different features, including timestamp, nonlinear energy, principal (1st to 3rd) components, and negative peak amplitude of each channel. All sorted single units (putative pyramidal neurons and interneurons) are included in population analysis.

B. Single Unit Analysis

In single cell analysis (for a multiple-trial setting), under 200 mW laser intensity we have found that the average percentages of positive and negative responsive units are respectively 47.5% and 2.5% among ACC populations ($n = 39$ cells), as well as respectively 54.3% and 22.9% among S1 populations ($n = 49$ cells). These ratios also change with varying laser intensities. Since S1 population encodes sensory component of pain, and ACC population encodes aversive or affective component of pain, some of responsive S1 neurons may encode pain non-specific responses (e.g., thermoneutral vs. thermonoxious components).

C. Population Analysis

In population analysis, we apply our proposed methods to analyze ACC and S1 ensemble spike data in both single and multiple-trial settings. The detection of significant change in P-value is robust for $m = 1$ or $m = 2$ (thus we use $m = 1$ for remaining analyses). In single-trial analysis, the algorithm converges very fast in most of cases. For instance, in our non-optimized MATLAB implementation at an iMac computer (3.5 GHz Intel Core i7, 16 GB RAM), the CPU time is <4 seconds for 12 units with 200 temporal bins. We define the latency as the first time point that the Z-score reached the statistical significance level relative to the onset of paw withdrawal in response to pain. If the time occurs prior to onset of withdrawal, the latency value is allowed to be negative. In comparison, we also compute the Z-score obtained from the multi-unit spike count activity (red curve, Fig. 3B), which is termed as the "empirical method". It is found that the Z-score computed from the latent variable is much smoother than the one computed directly from the empirical method. When all units show the same trend of change (e.g., increase) in firing rate, both methods can detect pain responses (Fig. 3, left panels); nevertheless, when only some but not all units exhibit a significant firing rate change, the empirical method fails to identify the pain episode, whereas our method can detect pain successfully (Fig. 3, right panels).

In our preliminary investigation, we note that the results derived from the linear Gaussian SSM are qualitatively similar and has a comparable latency statistic (results not shown due to space limitation). However, the computational speed is 4-5 fold faster than the Poisson SSM. Therefore, from a practical point of view, we may favor the speed with a little compromise of accuracy.

In multiple-trial analysis, we compute the between-trials variability for both ACC and S1 populations. In the Fig. 2 example, we obtained the pairwise D_{KL} statistics is 0.22 ± 0.19. We also compare the latency statistics computed from ACC and S1 populations under different laser intensity. Detailed experimental results will be reported elsewhere.

V. Discussion

A naive implementation of the full or simplified SSMs has a linear computational complexity of $\mathcal{O}(\ell m^3 T)$, where ℓ denotes the number of EM iterations until convergence. Nevertheless, in online detection applications, it is preferred to analyze ensemble spike data on the fly. One possible solution is to use some initial recordings as training set for model identification, and in the testing phase to run a recursive Kalman filter (with $\mathcal{O}(m^3 + m^2 C + m C^2 + C^3)$ complexity) or a steady-state Kalman filter (with $\mathcal{O}(m^2 + mC)$ complexity). The computational reduction is significant when either m or C is large. In a non-stationary scenario, the parameters can be updated in an online or semi-batch manner. A sequential change-point (pain) detection algorithm is under investigation.

In our statistical analysis, the latent variable $\mathbf{z}_k$ is modeled as a linear Gaussian Markovian process, the smoothing is determined by the transition matrix $\mathbf{A}$ and noise covariance matrix $\mathbf{Q}$. To increase the flexibility of smoothing, we can introduce a Gaussian process (with zero mean and a covariance function $K(\mathbf{z}_k, \mathbf{z}_l)$, parameterized by some hyperparameters) to the state equation [5], [18]. In addition, the state space analysis is useful for dimensionality reduction and visualization for neuronal ensembles.

Acknowledgments

We thank T. Manders, A.P.S. Tong and Q. Zhang for collecting experimental data used here.

REFERENCES

[1] L. Buesing, J. H .Macke, and M. Sahani, "Spectral learning of linear dynamics from generalized-linear observations with application to neural population data," *NIPS*, 2012.

[2] M. C. Bushnell, M. Ceko and L. A. Low, "Cognitive and emotional control of pain and its disruption in chronic pain," *Nat. Rev. Neurosci.* 14: 502-511, 2013.

[3] Z. Chen, "An overview of Bayesian methods for neural spike train analysis," *Comput. Intell. Neurosci.*, Article 251905, 2013.

[4] Z. Chen and E. N. Brown, "State space models," *Scholarpedia*, 8(3):30868, 2013.

[5] J. P. Cunningham, B. M. Yu, K. V. Shenoy and M. Sahani, "Inferring firing rates from spike trains using Gaussian processes," *NIPS*, 2007.

[6] Z. Ghahramani and G. E. Hinton, "Parameter estimation for linear dynamical systems," Tech. Rep. CRG-TR-96-2, Univ. Toronto, 1996.

[7] M. Lee, T. R. Manders, S. E. Eberle, C. Su, J. D'amour, R. Yang, H. Y. Lin, K. Deisseroth, R. C. Froemke and J. Wang, "Activation of corticostriatal circuitry relieves chronic neuropathic pain," *J. Neurosci.*, 35(13):5247-5259, 2015.

[8] X-Y. Li, N. Wang, Y-J. Wang, Z-X. Zuo, K. Koga, F. Luo and M. Zhuo, "Long-term temporal imprecision of information coding in the anterior cingulate cortex of mice with peripheral inflammation or nerve injury," *J. Neurosci.* 34(32): 10675-10687, 2014.

[9] M. E. Khan, A. Y. Aravkin, M. P. Friedlander, and M. Seeger, "Fast dual variational inference for non-conjugate latent Gaussian models," *Proc. 30th ICML*, 2013.

[10] C. C. Kuo and C. T. Yen, "Comparison of anterior cingulate and primary somatosensory neuronal responses to noxious laser-heat stimuli in conscious, behaving rats," *J. Neurophysiol.* 94: 1825-1836, 2005.

[11] J. S. Mogil, "Animal models of pain: progress and challenges," *Nat. Rev. Neurosci.* 10: 283-294, 2009.

[12] L. Paninski, Y. Ahmadian, D. G. Ferreira, S. Koyama, K. R. Rad, M. Vidne, J. T. Vogelstein and W. Wu, "A new look at state-space models for neural data," *J. Comput. Neurosci.*, 29(1-2): 107-126, 2009.

[13] D. Pfau, E. A. Pnevmatikakis, and L. Paninski, "Robust learning of low-dimensional dynamics from large neural ensembles," *NIPS*, 2014.

[14] H. Shimazaki, S. Amari, E. N. Brown and S. Gruen, "State-space analysis of time-varying higher-order spike correlation for multiple neural spike train data," *PLoS Comp. Biol.* 8(3): e1002385, 2012.

[15] A. Smith and E. N. Brown, "Estimating a state-space model from point process observations," *Neural Comput.*, 15: 965–991, 2003.

[16] A. Smith and E. N. Brown, "State-space algorithms for estimating spike rate functions," *Comput. Intell. Neurosci.*, Article 426539, 2010.

[17] C. J. Vierck, B. L. Whitsel, O. V. Favorov, A. W. Brown and M. Tommerdahl, "Role of primary somatosensory cortex in the coding of pain," *Pain*, 154: 334-344, 2013.

[18] B. M. Yu, J. P. Cunningham, G. Santhanam, S. I. Ryu, K. V. Shenoy and M. Sahani, "Gaussian-process factor analysis for low-dimensional single-trial analysis of neural population activity," *J. Neurophysiol.* 102: 614-635, 2009.

[19] Y. Zhang, N. Wang, J-Y. Wang, J-Y. Chang, D. J. Woodward and F. Luo, "Ensemble encoding of nociceptive stimulus intensity in the rat medial and lateral pain systems," *Molecular Pain*, 7: 64, 2011.

[20] M. Zhuo, "Neuronal mechanism for neuropathic pain," *Molecular Pain*, 3: 14, 2007.