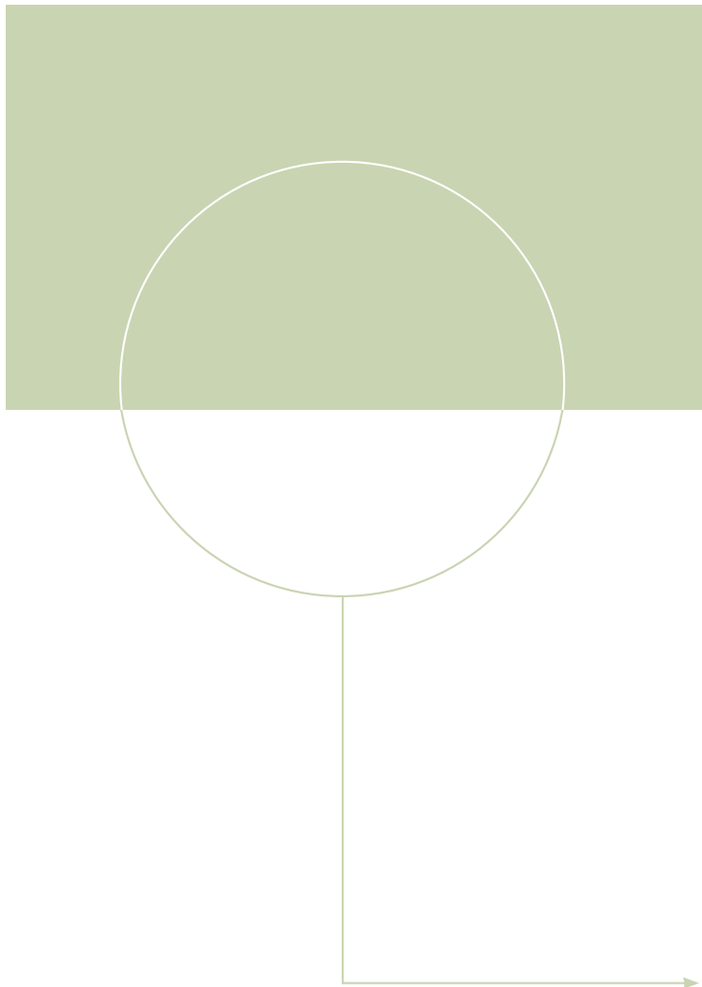

Master's thesis

2019

Atle Eskeland Rimehaug

NTNU
Norwegian University of
Science and Technology
Faculty of Natural Sciences
Department of Physics

Atle Eskeland Rimehaug

Validating the laminar population analysis for decomposing cortical multielectrode data on a model of the mouse primary visual cortex

June 2019

NTNU
Norwegian University of
Science and Technology

NTNU
Norwegian University of
Science and Technology

NTNU

Norwegian University of
Science and Technology

Validating the laminar population analysis for decomposing cortical multielectrode data on a model of the mouse primary visual cortex

Atle Eskeland Rimehaug

Applied Physics and Mathematics
Submission date: June 2019
Supervisor: Jon Andreas Støvneng
Co-supervisor: Alexander Stasik

Norwegian University of Science and Technology
Department of Physics

Abstract

Recent technological innovation has led to a vast increase in the amount neural data that can be collected at once, and spawned the need for new methods to analyze these large, multidimensional data sets. Laminar multielectrodes measures the extracellular potential (ECP) and allows recording of the activity from thousands of neurons simultaneously across the cortical layers. However, since many sources contribute to the measured ECP it can also be hard to interpret. Traditionally, the high-frequency (above 500 Hz) part of the ECP, referred to as multi-unit activity (MUA), has been used to extract spikes, while the low-frequency (below 500 Hz) part, referred to as the local field potential (LFP), has largely been ignored. Efforts to improve interpretability of the LFP by decomposing it into the different sources have often been hampered by the lack of suitable tools for performing this decomposition. One method developed to remedy this is the laminar population analysis (LPA), where the MUA and LFP are utilized jointly in the decomposition. Here, the adequacy of this method was assessed by applying it on a model of the primary visual cortex currently under development at the Allen Institute for Brain Science, and its performance was evaluated by comparing it to the performance of the standard decomposition method principal components analysis (PCA). With LPA, three out of the five laminar populations were correctly identified and separated from the other layers, while no single layer was identified and separated with PCA. Neither LPA nor PCA convincingly separated the contributions of the different laminar populations to the LFP. All the populations that LPA partitioned the LFP into were dominated by excitatory layer 5 and layer 6 cells. The only indication of a possible separation of the different contributions was that one population was more influenced by layer 5 cells than cells in other layers. With PCA, all major sources to the LFP were allocated to the first component. The inefficacy of both approaches in the decomposition of the LFP could be due to artefacts of this model, such as overly synchronous firing across layers, and the dominance of the LFP by contributions from layer 5 and layer 6 excitatory cells.

Sammendrag

Teknologisk innovasjon har ført til en stor økning i mengden nevrale data som kan innhentes på en gang, noe som har avlet et behov for nye metoder for å analysere disse store, multidimensjonale datasettene. Laminære multielektroder måler det ekstracellulære potensialet (ECP) som reflekterer aktiviteten til tusenvis av nevroner samtidig over flere kortikale lag, men de mange kildene til ECP gjør det også vanskelig å tolke. Tidligere ble den høyfrekvente (over 500 Hz) delen av ECP, kalt multi-enhet aktivitet (MEA) brukt til å ekstrahere avfyringer, mens den lavfrekvente delen (under 500 Hz) av signalet, kalt lokalt feltpotensial (LFP), stort sett ble ignorert. Forsøk på å forbedre tolkningen LFP ved å dekomponere det til de ulike kildene har ofte blitt hemmet av mangelen på passende verktøy for å gjennomføre dekomponeringen. En metode utviklet for å løse dette er laminær populasjonsanalyse (LPA), hvor MEA og LFP blir brukt sammen i dekomponeringen. I dette prosjektet ble denne metoden evaluert ved å anvende den på en modell av primær visuell korteks som er under utvikling ved Allen-instituttet for hjerneforskning. Hvor godt den fungerer ble vurdert ved å sammenligne resultatet fra dekomponeringen med resultatet fra dekomponeringen fått ved å utføre prinsipiell komponoentsanalyse (PKA) på samme datasett. Med LPA ble tre av de fem laminære populasjonene som utgjør primær visuell korteks korrekt identifisert og skilt fra de andre lagene, mens intet enkeltlag ble identifisert og skilt ut med PKA. Hverken LPA eller PKA greier å skille bidragene fra de ulike laminære populasjonene til LFP på noen overbevisende måte. Alle populasjonene som LPA delte LFP inn i var dominert av eksitatoriske lag 5- og lag 6-celler. Den eneste indikasjonen på en mulig separarering av de ulike bidragene var at én populasjon var mer påvirket av lag 5-celler enn celler i andre lag. Med PKA var alle bidragene til LFPet samlet i én komponent. Den manglende nyttigheten av begge tilnærminger i dekomponeringen av LFP kan skyldes artefakter ved denne modellen, slik som over-synkron avfyring på tvers av lagene, og dominansen av bidrag fra de eksitatoriske cellene i lag 5 og 6 til LFPet.

Contents

1 Introduction **4**

2 Theory and Methods **6**
2.1 Early visual system 6
2.2 The primary visual cortex model 7
2.3 Simulations and data processing 9
2.4 Principal components analysis 10
2.5 Laminar population analysis 11
2.5.1 Relative contribution of populations 13

3 Results **14**
3.1 Principal components analysis 14
3.2 Laminar population analysis 22
3.2.1 Synthetic data 22
3.2.2 Allen model 25
3.2.3 Number of populations 34

4 Discussion **41**
4.1 Identifying laminar populations 41
4.2 Separating the sources of contribution to the LFP 41
4.3 Determining the number of populations in LPA 43
4.4 Implications of work and future directions 43

1 Introduction

The limited number and reach of observables in neuroscience has long restricted our knowledge of the brain, but in recent decades there has been a great expansion in the technologies available to experimental neuroscientists, allowing them to both investigate familiar neural territory in greater detail, and to explore entirely new areas of the brain. One of the now well-established techniques is to measure the extracellular potential with multielectrode arrays (Einevoll et al.; 2013). Extracellular potential (ECP), as the name suggests, is the potential measured outside of the cells, and is generated by their electrochemical activity. Therefore, it can serve as a proxy for neural activity.

Multielectrode arrays pick up the potential generated from up to thousands of neurons at once, making it a suitable tool for the study of neural activity at the mesoscopic level of networks. The high-frequency part (above 500 Hz) of the ECP is referred to as the multi-unit activity (MUA), and is thought to reflect the firing of action potentials by neurons in close proximity (0.1 mm above or below (Somogyvári et al.; 2005)) to each electrode, while the low frequency part (below 500 Hz) of the ECP is referred to as the local field potential (LFP) (Einevoll et al.; 2007), and is thought to mainly reflect the transmembrane currents in synapses and dendrites in a larger area (a few hundred micrometers (Lindén et al.; 2011)) around the electrodes, though there are multiple neuronal processes contributing to the LFP (Einevoll et al.; 2007, 2013). The many sources of contribution makes the LFP particularly difficult to interpret, thus the focus of ECP analysis has traditionally often been on extracting spikes from the MUA while largely ignoring the LFP. However, the richness of information in the LFP in addition to the relative ease with which it can be measured makes it an attractive observable to use for analyzing of brain activity, if it can be reasonably decomposed into its individual sources.

Unfortunately, standard tools for data decomposition, such as principal component analysis (PCA) and independent component analysis (ICA), often do not perform well on data from multielectrode recordings. The reason is that their underlying assumptions do not necessarily align with properties of the neural populations and their contributions to the ECP that we want to uncover. In PCA the components are constructed based on covaraciance and forced to be orthogonal to each other, meaning that there cannot be any correlation across components, and they are also constructed such that the first component explains as much of the variance as possible, the second component the second most, and so on down to the last component. However, the neural activity of populations underlying the ECP is often highly correlated across all populations. ICA, on the other hand, is based on the informational independence of its components, which is enforced by assuming that the signal of the individual components is non-gaussian, and this is often not the case with the firing rates of neural populations.

The inefficacy of standard decomposition tools prompted Einevoll et al. (2007) to develop the laminar population analysis (LPA) method, which is based

on physiological rather than statistical assumptions. In the LPA framework it is assumed that the LFP arises from synaptic currents and return currents in the dendrites of postsynaptic populations after the presynaptic populations have fired action potentials, and thus makes use of the MUA and LFP jointly in the decomposition. First, the MUA is decomposed into individual, spatially non-overlapping populations of neurons and their firing rates. Then these population firing rates are used to find the temporal profile of the resulting synaptic and dendritic currents, and the spatial profile of their contributions to the LFP.

Technological development has not only facilitated advances in experimental neuroscience, it has also enabled extensive development in the theoretical research (Einevoll et al.; 2013). The increase in computational power in recent decades has made it possible to go from simulating single neurons that were biologically realistic to simulating the activity in entire networks of biologically realistic neurons. Early network models often consisted solely of spatially dimensionless point-neurons, while newer models consist of different types of neurons with their characteristic axons and dendrites incorporated into the model (Potjans and Diesmann; 2012; Hagen et al.; 2016). Point-neuron networks can reproduce many of the behaviours of real networks, but it is not possible to compute LFPs from these networks because LFPs arise from the spatial separation of current sources and sinks in an electromagnetic dipole (Einevoll et al.; 2013; Hagen et al.; 2016; Arkhipov et al.; 2018). The biorealistic models, however, which include the spatial dimensions of neurons, can produce LFPs and thus enable the use of this measure in the validation of the model.

Currently the Allen Institute for Brain Science is developing such a biophysically realistic model of the primary visual cortex in mice, an area involved in the early processing of visual input (Carandini; 2012; Kandel et al.; 2013). In this project the LPA method will be validated on data produced from simulations on this model rather than on experimental data, which was done in Einevoll et al (2007). When validating against experimental data the true distribution of populations, has to be estimated by other measures, which means that the true distribution itself is uncertain, and this can limit the degree of confidence one can put on the validation. The LPA method has previously been validated against a model consisting of a 1000 layer 5 pyramidal neurons, 40 of which were firing action potentials (Pettersen et al.; 2008), and on a thalamocortical model developed by Traub et al. (2005) (Glabska et al.; 2016), but never on a model on the scale and level of detail as the model that is being developed at the Allen institute, which consists of 45000 biologically detailed neurons of 21 different neuronal classes. Therefore it will be informative to attempt validation of the LPA method on this model, and if it proves successful the LPA method can be used in the process of validating the model against experimental data, and consequently guide future development of it.

The aim of this project is thus to validate the LPA method on data generated from simulations on this primary visual cortex model. To evaluate the performance of the LPA method, the results of its decomposition will be compared to the results of applying PCA to the same data.

2 Theory and Methods

2.1 Early visual system

Visual information processing begins at the retina, which is in the back of the eye. When light collected by the eye hits the retina, specialized photoreceptors transduce the energy of the photons into the electrical energy, and this electrical energy causes neurons called retinal ganglion cells (RGC) to fire an action potential (fig. 1) (Kandel et al.; 2013). An action potential is a sudden shift in the electrical potential across the membrane lasting for 1-3 ms, and makes out the signal that transports visual information from the eye to the rest of the brain. The signal first travels to a structure known as the lateral geniculate nucleus (LGN) in the thalamus, the brain's sensory relay center, before reaching the visual cortex (Kandel et al.; 2013). The visual cortex is divided into areas responsible for processing different aspects of the visual stimulus. The first area is referred to as the primary visual cortex (V1), and is, among other things, involved in determining the direction of movement and the orientation of bars in the visual stimulus. (Kandel et al.; 2013; Carandini; 2012) This pathway of information flow is illustrated in figure 2a.

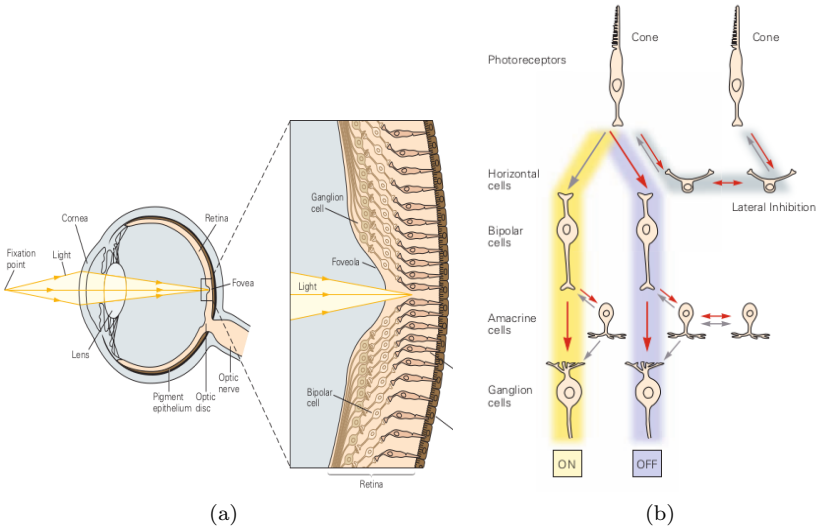

(a) (b)

Figure 1: **(a)** Illustration of eye. **(b)** Illustration of the information pathway from photoreceptors (cones) to retinal ganglion cells. (Kandel et al.; 2013)

The primary visual cortex, like the rest of the cortex, is divided into six anatomically and functionally distinct layers (illustrated in fig. 2b and 3). The layers are characterized by different numbers, size and packing density of neurons as well as idiosyncratic connectivity to other areas of the brain (Ahissar and Staiger; 2010). The first layer consists mainly of dendrites from the other

layers and contains few cells. The second layer and third layer are often pooled together into one layer, referred to as layer(s) II/III, because of the lack of a clear cytoarchitectonic border between them. This layer II/III is primarily composed of medium-sized pyramidal cells, named so because of their pyramidal-like structure, and they are connected to other cortical areas. Layer IV is composed of so-called spiny stellate and star pyramidal cells, and is the primary layer to receive projections from the LGN, though the other layers receive input from the LGN too. The fifth layer contains the largest pyramidal cells, which project to the brain stem, and the sixth layer is a more mixed population of cell types and is responsible for feedback projections to the thalamus. (Carandini; 2012; Kandel et al.; 2013)

2.2 The primary visual cortex model

The Allen Institute for Brain Science is currently developing a model of the primary visual cortex (V1) in mice based on large-scale experimental surveys performed at the institute as well as extensive literature curation (Billeh et al.; 2019). Henceforth this model will be referred to as the Allen model. The primary visual cortex is one of the most studied and best known areas of the brain, in part because it is relatively easy to access, and in part because it is easy to manipulate since it is activated by visual stimulus, and the visual stimulus can be created and controlled on a computer. The abundance of experimental data to validate against makes V1 a prime area for modelling, and a sensible starting point for large-scale, bio-realistic models.

Earlier modelling efforts were often either focused on biologically detailed models of single cells, or on simplified point-neuron networks. Efforts to combine biologically realistic neurons with large-scale networks were restricted by limitations in computational power. Developments in computer science in recent decades have given rise to super-computers that can handle the computational demands of large and detailed network models, and with these developments the models have followed (Potjans and Diesmann; 2012; Hagen et al.; 2016).

The Allen model is comprised of 230 000 neurons, whereof 52 000 are biologically detailed neurons from 21 different neuronal classes, and 178 000 are simpler, so-called leaky-integrate-and-fire (LIF) neurons (Billeh et al.; 2019). The LIF-neurons surround a core of bio-detailed neurons, and their purpose is to avoid edge effects in the model. The model is 850 µm high, corresponding to the thickness of mouse V1, and the core of detailed neurons has a diameter of 400 µm, and the total diameter of the model is 845 µm (fig. 2b and 2c). In simulations on the model the visual stimulus is conveyed through thalamcortical projections from the LGN (Billeh et al.; 2019).

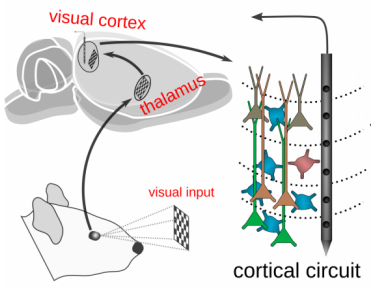

(a)

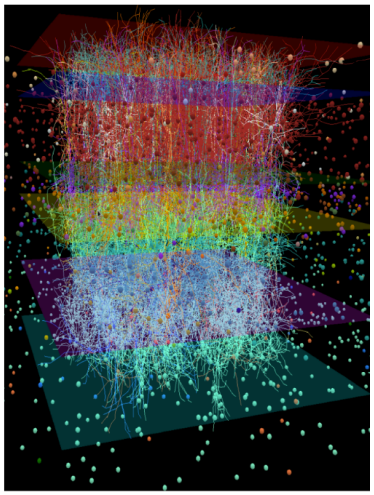

(b)

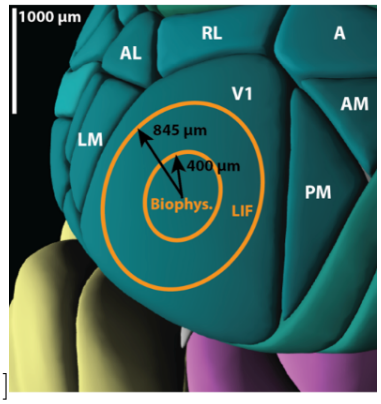

(c)

Figure 2: **(a)** Illustration of the information flow from visual stimulus to processing in visual cortex with a laminar multielectrode measuring the ECP in the cortical circuit. (). **(b)** Visualization of the biophysical network of the Allen model. **(c)** Illustration of area of mouse cortex covered by the model. (Billeh et al.; 2019)

The simpler point-neuron network models can reproduce many of the qualitative features in the functionality of biological networks, but their quantitative agreement with experiments tends to suffer (Arkhipov et al.; 2018). Thus, if the goal is a quantitative understanding of neural phenomena, more detailed models like the Allen model are probably necessary. Another property of network models with bio-realistic neurons is that one can calculate extracellular potentials (ECP) generated from their activity, which is not possible with point-neuron models. The reason for this is that to generate an electromagnatic field the current source and sink of an electromagnetic dipole has to be separated in space, and since point-neurons do not have a spatial dimensions, all currents collapse together in the single point at which the neuron resides. The biolog-

ically detailed neuron models, however, incorporate the neurons' morphologies and hence their spatial extent, so the sources and sinks of the transmembrane currents are separated and a field can be modelled.[1]

This means that bio-realistic network models can be validated against experimental data obtained by multielectrode recordings, since ECP is the observable in these measurements. ECP can be split into a high-frequency part (above 500 Hz) and a low-frequency part (below 500 Hz) (Einevoll et al.; 2007). The high-frequency part is referred to as the multi-unit activity (MUA), and is thought to reflect the firing of action potentials in neurons close to the electrode contacts (0.1 mm above or below (Somogyvári et al.; 2005)). The low-frequency part is called the local field potential (LFP), and is thought to mainly stem from the transmembrane currents in synapses and dendrites in an area spanning a few hundred micrometers around the electrode (Lindén et al.; 2011; Einevoll et al.; 2013). In addition to providing another observable for validation, the generation of ECP in models like the Allen model can aid understanding of the origins of these signals. Furthermore, they can also be used in the validation of data analysis techniques like the laminar population analysis, which was developed specifically for the analysis of laminar multielectrode recordings.

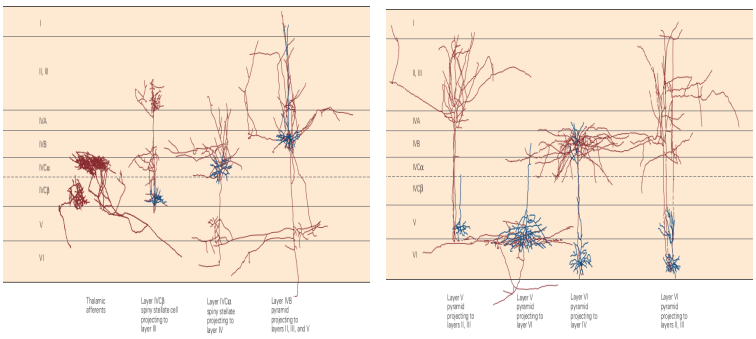

(a) (b)

Figure 3: Illustration of cell types in **(a)** layer 4 and **(b)** layer 5 and 6 of the visual cortex. IVB, IVα, IVβ refers to different subregions of layer 4. (Kandel et al.; 2013)

2.3 Simulations and data processing

The data used in this project were obtained from simulations on the Allen model performed by Espen Hagen at the Centre for Integrative Neuroplasticity (CINPLA), University of Oslo on the Jureka supercomputer at Jũelich research center in January and March 2019. The simulations were of an experiment where mice were exposed to alternating black and white flashes, each lasting for

[1] For excellent explanations of how the fields are calculated, see (Einevoll et al.; 2013) and (Holt and Koch; 1999)

250 ms, with 750 ms of gray screen between, and 500 ms of gray screen before the first flash. In total the simulations ran for 5 hours and 35 minutes for an experiment that lasted for 20.5 seconds. Two data sets from two separate simulations with the same protocol were used in the analysis. The first data set consisted of ECP (and more) from the whole model throughout the simulation, with the MUA and LFP already computed. The other data set consisted of contributions to the ECP from each of the 21 different cell classes throughout the simulation. The MUA and LFP had to be calculated from the ECP in the second data set, which was done by running it through respectively high- and low-pass filters with a cut-off frequency of 500 Hz for both. The butterworth filter function in the scipy.signal library was used to create these filters, and the "filtfilt" function in the same library was used to perform the filtering. All data processing and data analysis was done in python 3. The sampling frequency in the first data set was 1 kHz, while in the other it was 20 kHz, so after the filtering the MUA and LFP of all cell classes in the second data set were downsampled to 1 kHz to match the MUA and LFP from the whole model, which was done using the "decimate" function in scipy.signal. Lastly, since all analysis in this projects was performed at the laminar population level contributions from different cell types were summed together into common contributions from excitatory and inhibitory cells in each layer.

2.4 Principal components analysis

In the principal components analysis (PCA) framework the axes are rotated such that they point in the directions along which the the variance in the data was greatest. These new axes are the principal components, and these are ordered so that the first principal component explains as much variance as possible, the second component explains as much as possible of the remaining variance, and so on in descending order down to the Nth component in a data set with N variables. This property of maximizing the variance picked up by the first few components is what makes PCA a popular method for dimensionality reduction, since it is possible to retain most of the variance, and hence information, with fewer components. (Wolfgang and Léopold; 2015; James et al.; 2013; Peyranche et al.; 2010)

Pointing in the direction of maximum variance corresponds to pointing in the direction of maximum covariance between the variables in the original data set. This means that the components can be used to identify variables that correlate with each other, and since the underlying statistical assumption of PCA is that the components should be orthogonal, they will identify different sets of correlating variables. In the context of neural data from multielectrode recordings, this means finding co-firing neurons when applied to MUA, and finding co-occuring transmembrane currents when applied to LFP. Hence, PCA can in principle be utilized for identification of neural populations from MUA, and the distribution of synapses and dendrites from LFP (Wolfgang and Léopold; 2015; Peyranche et al.; 2010).

The output from performing PCA are a set of eigenvalues and eigenvectors.

The eigenvalues give the variance of the components, and the proportion of variance explained by a component i is given by $\Gamma_i = \frac{\lambda_i}{\sum_{i=1}^{N} \lambda_i}$, where λ_i is the eigenvalue of component i. The eigenvectors are referred to as the spatial loading scores and give the spatial profile of their respective components. The temporal profiles are obtained by taking the inner product of the original data matrix and the eigenvectors.

The resulting decomposition can be expressed as a sum of spatiotemporally separable functions:

$$\phi(z_i, t_j) = \sum_{n=1}^{N} \lambda_n f_n(z_i) g_n(t_j) \tag{1}$$

where $\phi(z_i, t_j)$ is the value of either the MUA or the LFP, depending on what the PCA is performed on, at vertical position z_i at time t_j, N is the number of components, which will be the same as the number of channels/electrodes on the multielectrode, $f_n(z_i)$ is the value of the spatial profile of component n at z_i, and $g_n(t_j)$ is the value of the temporal profile for the same component (Einevoll et al.; 2007). With normalized spatial and temporal profiles, F_n and G_n this equation can be written on the form:

$$\Phi(z_i, t_j) = \sum_{n=1}^{N} F_n(z_i) G_n(t_j) \tag{2}$$

2.5 Laminar population analysis

Laminar population analysis (LPA) is also a technique for decomposing time series data into spatiotemporally separable functions, but it is developed specifically for identifying cortical layers from laminar multielectrode recordings. Furthermore, it is based on physiological rather than statistical assumptions, which alters the form and interpretation of the spatial and temporal functions. In the LPA framework it is assumed that the recorded LFP arises primarily from synaptic currents and dendritic currents in the postsynaptic populations after action potential firing in the laminar populations, meaning that the spiking that is picked up in the MUA drives the LFP. This assumption led to the idea that the MUA and LFP could be decomposed jointly to both identify the laminar populations and the synaptic and dendritic currents caused by firing in these populations (Einevoll et al.; 2007, 2013).

First, the number of populations to decompose the data into has to be assumed a priori, and a set of parameters fixing the spatial profiles of these populations are initialized. The form of the spatial profiles is assumed to be non-overlapping trapezoids, which means that the parameters determining their distribution is their position z_i, the width of the top a_i, and the width of their slope b_i, for a population i, and the height is set to 1 for all populations. Assuming these non-overlapping spatial profiles forces the populations to be localized,

in contrast to PCA where the predicted populations can both be spatially discontinuous and overlapping. These assumptions are made because the laminar populations to be found are expected to be continuous and non-overlapping. With the parameters of the spatial profiles initialized, the firing rates are estimated by performing a pseudoinverse:

$$\mathbf{r}_{est} = (\mathbf{M}^{\dagger}\boldsymbol{\phi}_M)^{\mathrm{T}} \tag{3}$$

where $\mathbf{M}^{\dagger}$ denotes the pseudoinverse of an $N_{pop} \times N_{chan}$ matrix containing the population spatial profiles, $\boldsymbol{\phi}_M$ is the recorded MUA, and $\mathbf{r}_{est}$ is an $N_{pop} \times B$ matrix of firing rates over time for all populations, where each row corresponds to the firing rate of a single population. This firing rate estimate can then used to compute the predicted MUA given this set of spatial profile parameters:

$$\boldsymbol{\phi}_M^{est} = \mathbf{M}\mathbf{r}_{est} \tag{4}$$

The discrepancy between this estimate and the recorded MUA is evaluated by calculating the relative mean square error e_M:

$$e_M = \frac{||\boldsymbol{\phi}_M - \boldsymbol{\phi}_M^{est}||}{||\boldsymbol{\phi}_M||} \tag{5}$$

$$= \frac{\sum_{c=1}^{N_{ch}} \sum_{j=1}^{B} (\phi_M(z_c, t_j) - \phi_M^{est}(z_c, t_j))^2}{\sum_{c=1}^{N_{ch}} \sum_{i=1}^{B} (\phi_M(z_c, t_j))^2} \tag{6}$$

The parameters z_i, a_i, b_i defining the spatial profiles of the populations are then changed randomly within their constraints, a new estimate for the firing rates and the MUA is calculated with (3) and (4), and the mean square error of the new MUA estimate is computed. This algorithm is performed iteratively until the error is minimized.[2]

Since it is assumed that the presynaptic firing drives the LFP, an estimate for the temporal profile of the LFP can now be computed from the fitted firing rate. The action potential in the presynaptic populations travels along the neural fibers and across the synaptic cleft to the dendrites of the postsynaptic population, inducing synaptic and dendritic currents underlying the LFP along the way. Thus, the transfer function in the calculation of the temporal profile of the LFP has to incorporate the effects of synaptic delays, synaptic time constants, and the physiological properties of the dendrites on the action potentials as they travel up the neural fibers. This transfer function is assumed to be an exponential kernel of the form:

$$h(t_j) = \frac{1}{\tau} e^{-(t_j - \Delta)/\tau} \Theta(t_j - \Delta) \tag{7}$$

where the parameters τ and Δ are the time constants and the delays, respectively, and $\Theta(t_j - \Delta)$ is the heaviside function, imposed by causality:

[2] An optimization algorithm is used to perform this minimization procedure. The differential evolution algorithm in the OpenOpt library was chosen for the MUA.

$h(t_j - \Delta < 0) = 0$. The temporal profile of all populations can then be computed with the following convolution:

$$(\mathbf{h} \circledast \mathbf{r}_{est})(t_j) = \sum_{k=-\infty}^{\infty} \mathbf{h}(t_k)\mathbf{r}_{est}(t_j - t_k) \tag{8}$$

where the kernel h is assumed to be the same for all populations. This can now be used to compute an estimate for the LFP $\boldsymbol{\phi}_L^{est}$:

$$\boldsymbol{\phi}_L^{est} = \mathbf{L}(\mathbf{h} \circledast \mathbf{r}_{est}) \tag{9}$$

where $\mathbf{L}$ is an $N_{pop} \times N_{chan}$ matrix of the LFP spatial profiles for each population. These spatial profiles represent the contributions of synaptic and dendritic currents in each population to the total LFP, and, in contrast to the MUA spatial profiles, they are unconstrained and fitted non-parametrically. The parameters that are fitted here are the kernel parameters τ and Δ. The LFP estimate $\boldsymbol{\phi}_L^{est}$ is optimized by changing the kernel parameters randomly within their constraints, and calculating a new estimate for the temporal profile and the LFP according to (8) and (9), until the relative mean square error between the estimated LFP and the recorded LFP is minimized. This relative mean square error is calculated analagous to (6).[3]

The results of the laminar population analysis are spatial distributions of the cortical layers, estimates for their firing rates, spatial distribution of the contributions of the synaptic and dendritic currents produced by each laminar population to the LFP, and a time constant and delay for the transport of population firing and the LFP.

2.5.1 Relative contribution of populations

To determine how many populations to assume when performing the LPA it can be useful to calculate the relative contribution of the populations to the MUA and LFP estimates. If one (or more) populations contribute substantially less than the others, this population can potentially be disregarded, and thus the assumed number of populations should be reduced. The relative contributions are calculated as follows:

$$W_k = \frac{(\sum_{i=1}^{N_{chan}} S_{i,k}) \times (\sum_{j=1}^{N_{time}} T_{j,k})}{\sum_{k=1}^{N_{pop}} ((\sum_{i=1}^{N_{chan}} S_{i,k}) \times (\sum_{j=1}^{N_{time}} T_{j,k}))} \tag{10}$$

where W_k is the relative weight of population k, $S_{i,k}$ is the value of the spatial profile of population k at channel number i, and $T_{j,k}$ is the value of the temporal profile of population k at time point j.

[3]The optimization algorithm used to perform this minimization here was the Broyden–Fletcher–Goldfarb–Shanno algorithm (BFGS), which is the default, in the minimize function of the scipy.optimize library.

3 Results

In the following the results from applying PCA and LPA to both the MUA and the LFP from the Allen model simulations will be presented. The results from applying the PCA will serve as a benchmark for evaluating the performance of LPA.

3.1 Principal components analysis

Multi-Unit Activity. The total variance explained by the first three principal components was 94.2%, and the fourth principal component explained less than 1% of the variance (tab. 1). Since the contribution of components beyond the first three was so small, it was assumed that they were unlikely to correspond to populations, thus only the first three components were kept. The approximation of the MUA by these three components is shown in figure 4b, and a qualitative comparison with the original MUA data generated with the Allen model for both white and black stimulus (fig. 4a) suggests that the MUA is well rendered by these three components. Henceforth, results are only shown for the white flash for simplicity.

Table 1: The proportion of the variance in MUA explained by the first four principal components.

PC	1	2	3	4
Proportion of variance	89.9 %	2.5 %	1.7 %	0.9%

The spatial profiles of the components give the localization of the populations in V1 predicted by these components (figure 5b). These profiles can be compared to the layer borders in the Allen model, which are based on the average depths of each layer in mouse V1 reported in the literature and represents the ground truth for the distribution of laminar populations. Performing this comparison qualitatively, we see that first component is mainly situated in layer 5 and 6, with its peak value residing in layer 5, and that the second and third component are more evenly and similiarly distributed across all layers, though they too peak in layer 5. This suggests that at most two populations are distinguished and separated by performing PCA on MUA: A deep population represented by component 1 which merges layer 5 and layer 6, and a more shallow population represented by component 2 or 3 which is comprised of layer 2/3 and layer 4. However, the substantial overlap in the spatial profiles of the components makes it hard to confidently assert a certain separation of the populations from this qualitative analysis.

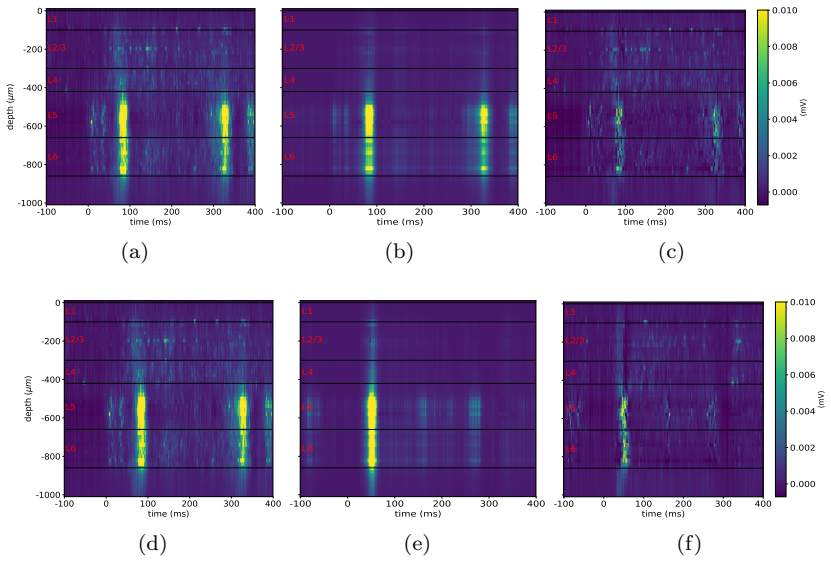

Figure 4: MUA before and throughout the white (top panel) and black (bottom panel) flash. Stimulus onset at 0 ms and stimulus offset at 250 ms. **(a)** and **(d)**: MUA of Allen model, **(b)** and **(e)**: MUA estimated by first three PCs. **(c)** and **(f)**: The absolute difference between the MUA of the Allen model and this estimate.

To explore the identification of populations further and more quantitatively, the correlation between the mean MUA of each component and the mean MUA of cells in each layer was calculated (fig. 5c and table in fig. 5d). The mean was taken from onset of the white flash to 250 ms after the offset. The first component correlates most strongly and positively with layers 5 and 6, and correlates negatively with layers 2/3 and 4. Both the second and third component correlates positively with layers 2/3 and 4, but the correlations are only significant for the third component. The second component correlates negatively and significantly with layer 5, while the third component correlates negatively and significantly with layer 6. This supports the interpretation that the first component mainly captures the deeper layers, and the two other components capture the upper layers. The lack of significant correlations with the upper layers for the second component, however, means that the third component is the most viable candidate to represent the populations in the upper layers. None of the components correlate significantly with layer 1.

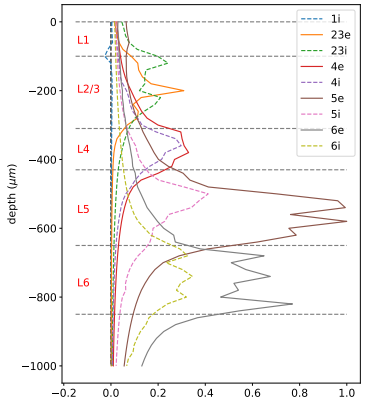

(a)

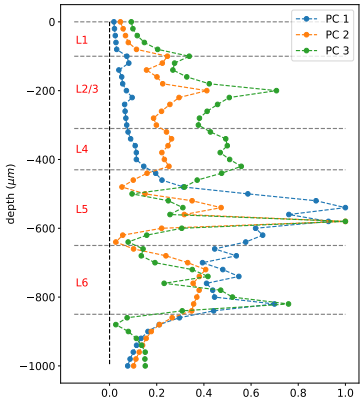

(b)

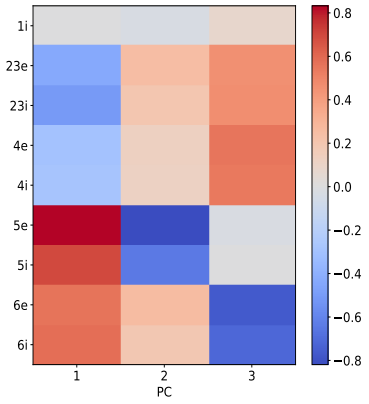

(c)

Spatial Correlation

Layer	PC 1	PC 2	PC 3
1i	0.00	-0.03	0.07
2/3e	-0.43**	0.25	0.45***
2/3i	-0.51***	0.20	0.46***
4e	-0.29*	0.13	0.55***
4i	-0.28*	0.12	0.54***
5e	0.83***	-0.81***	-0.02
5i	0.69***	-0.63***	0.01
6e	0.55***	0.25	-0.77***
6i	0.57***	0.18	-0.72***

(d)

Figure 5: (**a**) Mean MUA for excitatory (e) and inhibitory (i) cells of each layer. (**b**) Spatial profiles of the first three principal components, normalized to their individual peak values for visibility. (**c**) Correlation matrix of mean MUA of each layer with mean MUA of each component. Mean taken from onset of white flash to 250 ms after flash offset. (**d**) Correlation coefficients for the matrix, $^*p < 0.05$, $^{**}p < 0.01$, $^{***}p < 0.001$.

The same analysis was done for the temporal profiles of the components, which would correspond to the population firing rates, to investigate whether any timing differences in the spiking activity of the layers could be discerned with PCA. Studying of the average MUA of cells in each layer (fig. 6a) reveals that firing in the upper layers (2/3 and 4) precedes firing in the lower layers (5 and 6) by about 20 to 30 ms. Furthermore, the timing in the firing rates of the first component matches the timing in the firing rates of layers 5 and 6, while the firing rates of the second and third component matches the firing rates in

layers 2/3 and 4. This is in line with the classification of components based on localization in space. The correlations between the temporal profiles of the components and the firing rates of the cells in different layers also supports this analysis (fig. 6b and table in fig. 6c).

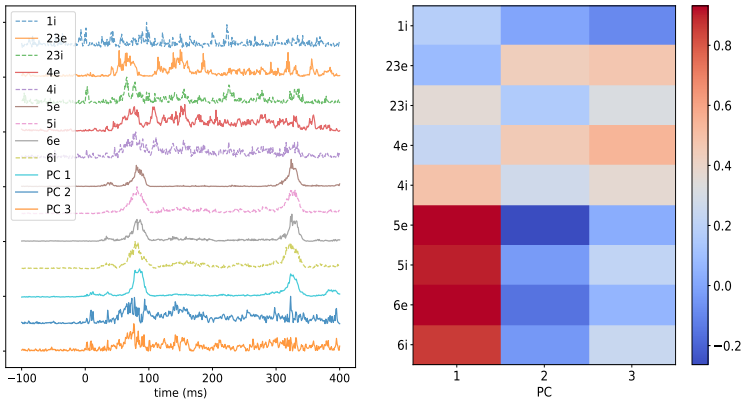

(a) (b)

Temporal Correlation

Layer	PC 1	PC 2	PC 3
1i	0.17***	-0.00	0.09
2/3e	0.09	0.44***	0.47***
2/3i	0.36***	0.16***	0.32***
4e	0.23***	0.45***	0.54***
4i	0.49***	0.27***	0.37***
5e	0.93***	-0.26***	-0.03
5i	0.90***	-0.03	0.22***
6e	0.93***	0.15***	-0.06
6i	0.86***	0.04	-0.25***

(c)

Figure 6: **(a)** Temporal profiles of excitatory (e) and inhibitory (i) cells from each layer with temporal profiles of components. All normalized to their individual peak values. **(b)** Correlation matrix for temporal profiles of cells from layers with temporal profile of components. **(c)** Correlation coefficients for the matrix in **(b)**, $^{*}p < 0.05$, $^{**}p < 0.01$, and $^{***}p < 0.001$.

Local Field Potential The first three components obtained from performing PCA on the LFP make out 99.8% of the total variance (tab. 2), and the approximation of the LFP by these components fits well with the original LFP from the Allen model (fig. 7). Excitatory cells from layers 5 and 6 make out substantially more of the contribution to the LFP than cells in the other layers (fig. 9 and 10), which suggests that using three components may just be over-

fitting. The third component was still retained, however, in case it captured the some of the contributions to the LFP better than the other components, as it did for the MUA.

Table 2: The proportion of variance in the LFP explained by the first three principal components.

PC	1	2	3
Proportion of variance	97.0%	1.9 %	0.86 %

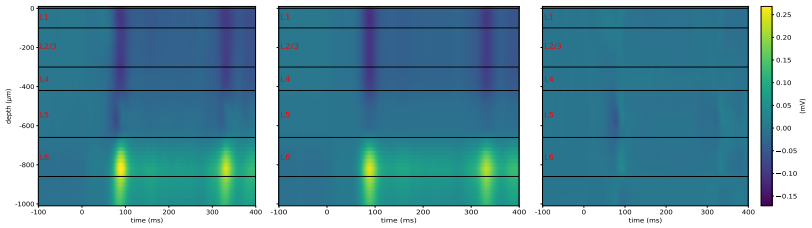

(a) (b) (c)

Figure 7: LFP before and throughout the white flash. Stimulus onset at 0 ms and stimulus offset at 250 ms. **(a)** LFP of Allen model, **(b)** LFP estimated by first three PCs. **(c)** The absolute difference between the LFP of the Allen and the estimate from the first three PCs.

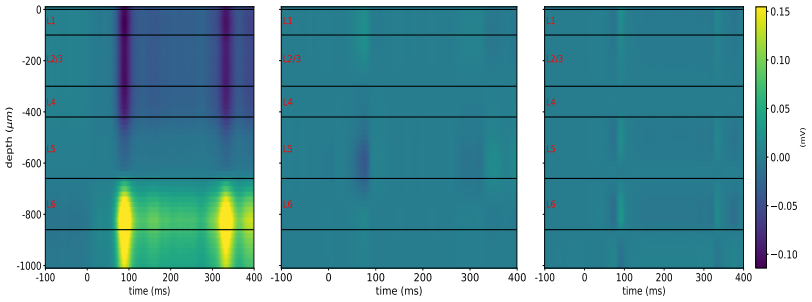

(a) PC 1 (b) PC 2 (c) PC 3

Figure 8: LFP of first three principal components

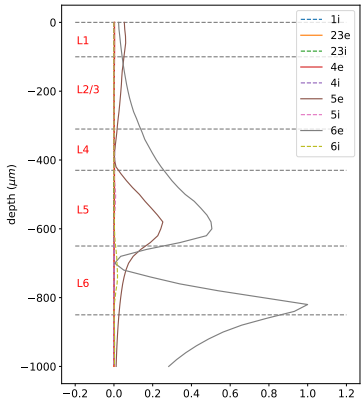

(a)

Proportion of variance

Layer	[%]
1i	0.0
2/3e	0.5
2/3i	0.0
4e	1.1
4i	0.6
5e	25.7
5i	0.2
6e	70.5
6i	1.3

(b)

Figure 9: Variance in LFP of cells from each layer. **(a)** Variance across depth, normalized to peak value for all cells. **(b)** Proportion of the sum of variances of all cells across all channels.

By comparing the LFP of each component (fig. 8) to the LFP contributions from cells in each layer (fig. 10) we see that the first component captures the LFP generated by excitatory cells in layer 6 (fig. 10h), while the second component might capture LFP generated by excitatory layer 5 cells (fig. 10f). The interpretation of the third component is less clear as its profile does not fit contributions from some layers markedly better than others. The average correlations over both space and time between the components and the LFP contributions of each layer supports the view that the first component captures the LFP contributions of layer 6, as it correlates positively with layer 6 cells, and negatively with cells in layers 1 to 4 (fig. 11). But it also seems to capture LFP contributions from excitatory layer 5 cells, as both the temporal and spatial correlation with these cells are positive and stand out from the other correlations. The second and third component, however, do not correlate notably more with some cells than others, with the possible exception of a positive temporal correlation between excitatory layer 4 cells and the second component. Together, these results indicate that only the first component contains notable contributions to the LFP, and the dominant contributions from excitatory layer 5 and layer 6 cells seem to be merged into this component.

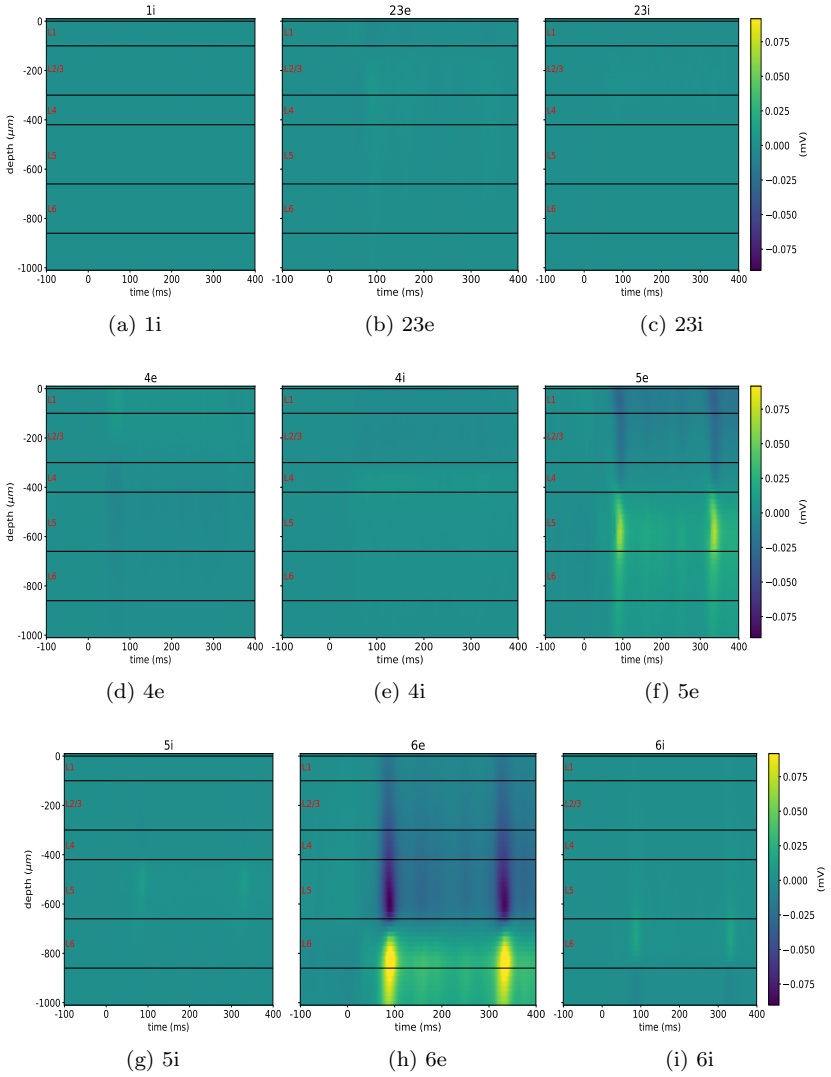

(a) 1i (b) 23e (c) 23i

(d) 4e (e) 4i (f) 5e

(g) 5i (h) 6e (i) 6i

Figure 10: LFP of excitatory (e) and inhibitory (i) cells in each layer.

The spatial correlation was found by calculating the correlation between LFP of each component and LFP generated by the different cells across all channels at each time step, and then averaging over time in the end. The temporal correlation was found by calculating the correlation between the time vectors of the LFP of each component and of LFP generated by the cells at each channel on the multielectrode, and then averaging over the number of channels. This was done because the LFP can take on both positive and negative values, so averaging over time or space *before* calculating the correlation can lead to mis-

leading results as the positive and negative values can cancel each other out. Hence, the reported correlations are average correlations, and it is not meaningful to calculate the statistical significance for an average correlation, so it is not possible to assess whether the correlations are significantly different from zero. Therefore these correlations should be interpreted with caution and the correlation between a the LFP of a component and the LFP of the cell types should first and foremost be considered relative to other correlations.

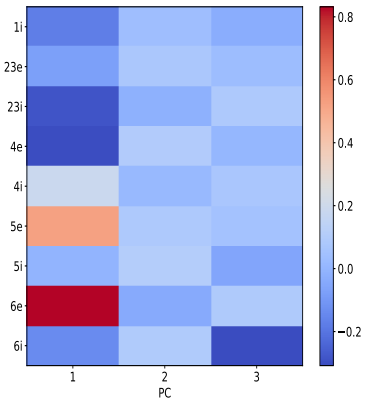

(a)

Spatial Correlation

Layer	PC 1	PC 2	PC 3
1i	-0.17	0.04	-0.03
2/3e	-0.07	0.08	0.04
2/3i	-0.29	-0.01	0.09
4e	-0.30	0.10	0.00
4i	0.19	0.15	0.07
5e	0.52	0.86	0.06
5i	-0.00	0.11	-0.06
6e	0.83	-0.04	0.09
6i	-0.14	0.10	-0.31

(b)

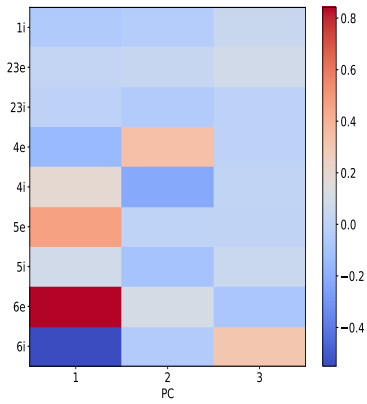

(c)

Temporal Correlation

Layer	PC 1	PC 2	PC 3
1i	-0.06	-0.04	0.05
2/3e	0.02	0.03	0.09
2/3i	-0.00	-0.05	0.00
4e	-0.15	0.34	0.00
4i	0.19	-0.22	0.01
5e	0.46	0.00	0.00
5i	0.09	-0.10	0.05
6e	0.84	0.11	-0.08
6i	-0.55	-0.05	0.30

(d)

Figure 11: Correlation in spatial distribution **(a, b)** and temporal profile **(c, d)** of LFP generated by cells in each layer and LFP attributed to each component in PCA.

3.2 Laminar population analysis

As a first test of the LPA approach it was performed on synthetic data to see if it could recover the constructed spatial profiles and firing rates of two populations. Once the adequacy of the LPA method had been established on this data, it was used on ECP data generated with the Allen model, as was done with PCA (see section 4.1). In the following the results of performing LPA on synthetic data will be presented first before presenting the results from doing it on the Allen model.

3.2.1 Synthetic data

Multi-Unit Activity The synthetic MUA was generated by first constructing two spatial profiles with trapezoid forms at distinct depths in a hypothetical cortex (fig. 14a) and firing rates of an exponential form with random time constants between 1 ms and 10 ms and random delays between 0 ms and 300 ms (fig. 13a). Then the outer product of these spatial profiles and firing rates was computed to obtain the MUA. The resulting MUA is shown in 12a. The MUA estimated by performing LPA on this data recovers its qualitative features (fig. 12b, 12c), and the relative mean square error e_M of this estimate was smaller than 0.001, demonstrating a good fit.

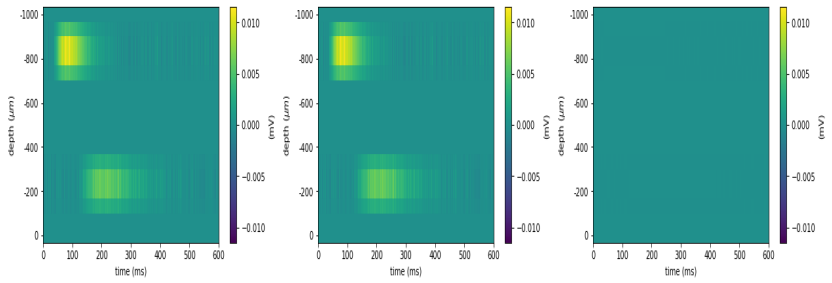

(a) (b) (c)

Figure 12: **(a)** Synthetic MUA. **(b)** MUA estimated by LPA. **(c)** Difference between the synthetic MUA and MUA estimated by LPA.

The original firing rates and the firing rates estimated by LPA are plotted together in figures 13b and 13c, and the constructed spatial distribution and LPA-estimated spatial distribution of the populations are plotted together in figures 14b and 14c. In all cases the overlap between the LPA-estimate and the original profiles is complete. Hence, LPA is able to recover the spatial distribution and the firing rates of populations for this synthetic MUA data.

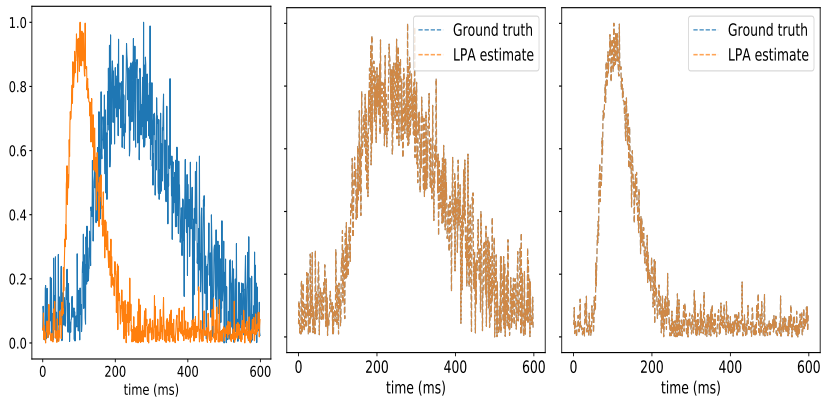

Figure 13: **(a)** Synthetic firing rates for the two populations. **(b)** and **(c)** Firing rates of the first and second population, respectively, together with the estimates from LPA.

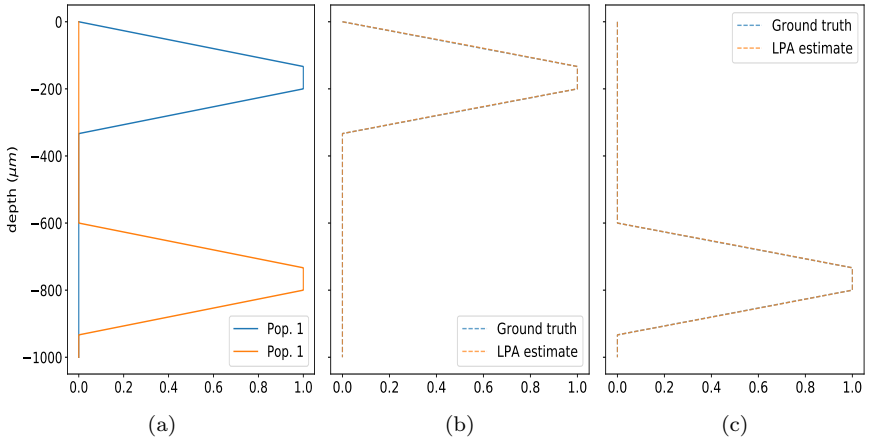

Figure 14: **(a)** Spatial MUA profile of the two populations. **(b)** and **(c)** Spatial profile of first and second population, respectively, together with the estimates from LPA.

Local Field Potential The population firing rates constructed for the MUA were convolved with an exponential kernel with a time constant $\tau = 50\,\text{ms}$ and delay $\Delta = 20\,\text{ms}$, values thought to be within a realistic range for the travelling of pre-synaptic action potentials to the postsynaptic dendrites. Under the assumptions of LPA, this convolution should produce the temporal profile of

the LFP (fig. 16a). The spatial profiles of the LFP contributions from the two synthetic populations were constructed as sinusoids with different initial phases (fig. 17a). To generate the synthetic LFP the outer product of the spatial profiles and the postynaptic potentials was computed (fig. 15a).

The LFP estimated from performing LPA on the synthetic LFP data reproduces the qualitative features of the synthetic data (fig. 15b and 15c), and the relative mean square error e_L of this estimate was 0.058, indicating a good fit. The temporal and spatial profiles of the populations estimated in LPA reproduced the synthetic temporal and spatial profiles used to construct the LFP (fig. 16b, 16c, 17b, and 17c). It also reproduced the kernel parameters $\tau = 50\,\text{ms}$, and $\Delta = 20\,\text{ms}$.

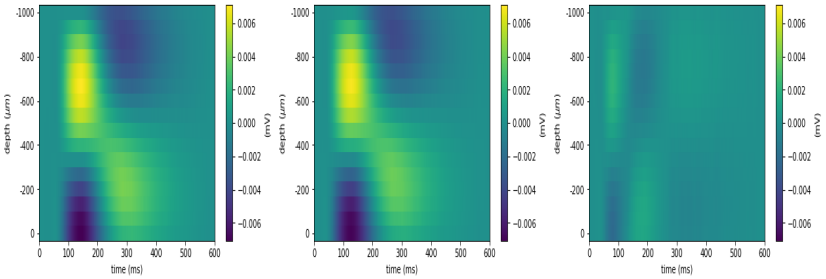

(a) (b) (c)

Figure 15: **(a)** Synthetic LFP. **(b)** LFP estimated by LPA. **(c)** Difference between the synthetic LFP and LFP estimated by LPA.

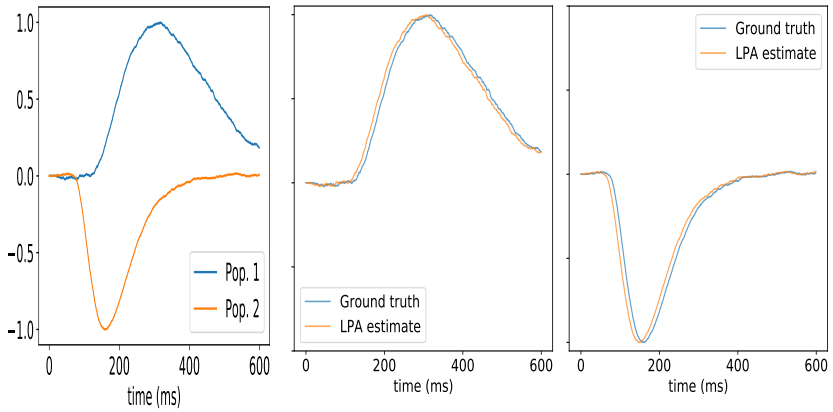

(a) (b) (c)

Figure 16: **(a)** Temporal profiles of the LFP for the two populations. **(b)** and **(c)**: Temporal profiles of the LFP for the first and second population, respectively, together with the estimates from LPA.

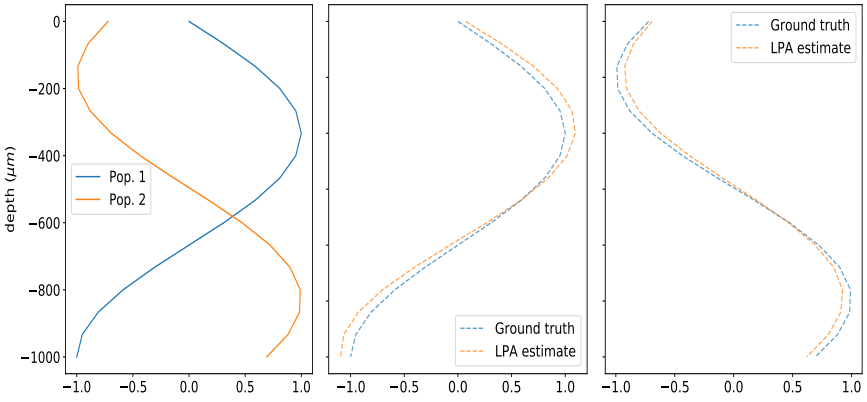

(a) (b) (c)

Figure 17: **(a)** Spatial LFP profile of the two populations. **(b)** and **(c)**: Spatial profile of first and second population, respectively, together with the estimates from LPA.

3.2.2 Allen model

The number of populations to be found has to be assumed in advance in the LPA algorithm. To get a more direct comparison with the performance of PCA it was first assumed that the number of populations to be found was three.

Multi-Unit Activity The LPA estimate of the MUA from Allen model reproduces the qualitative features of the original data (fig. 18), and the relative mean square error of this estimate was 0.083. The spatial profiles reveal that: layers 2/3 and 4 are essentially merged into one population, henceforth referred to as population 1; a region of layer 5 makes out a population on its own, henceforth referred to as population 2, and the whole of layer 6 also makes out a population on its own, henceforth referred to as population 3 (fig. 19a). Computing the correlation between the mean MUA of each population and the mean MUA of cells in each layer supports these conclusions. The first population correlates most strongly and positively with both excitatory and inhibitory cells in layers 2/3 and 4, the second population correlates most strongly and positively with both cell types in layer 5, and the third population correlates most strongly and positively with both cell types in layer 6 (fig. 19b and table in fig. 19c). The mean is taken from the onset of the white flash to 250 ms after the offset of the flash.

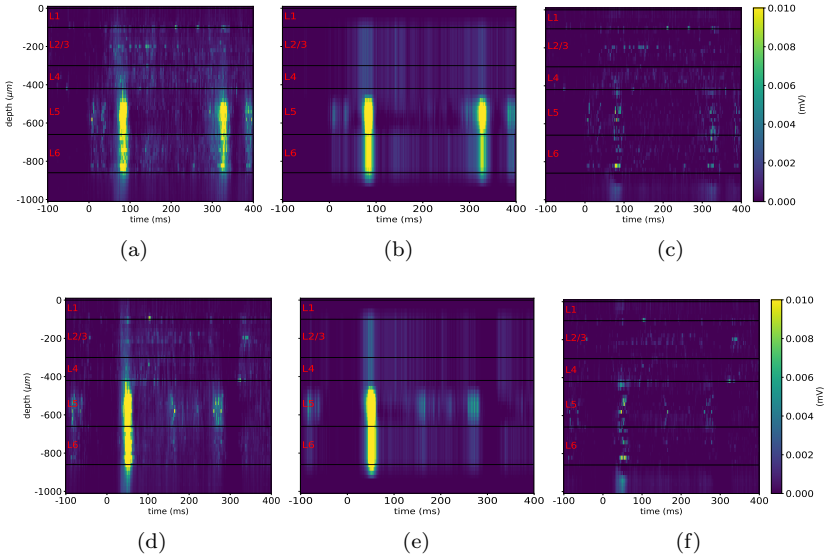

Figure 18: MUA before and throughout the white (top panel) and black (bottom panel) flash. Stimulus onset at 0 ms and stimulus offset at 250 ms. **(a)** and **(d)**: MUA of Allen model. **(b)** and **(e)**: MUA estimated three populations in LPA. **(c)** and **(f)**: The absolute difference between the MUA of the Allen and the estimate from LPA.

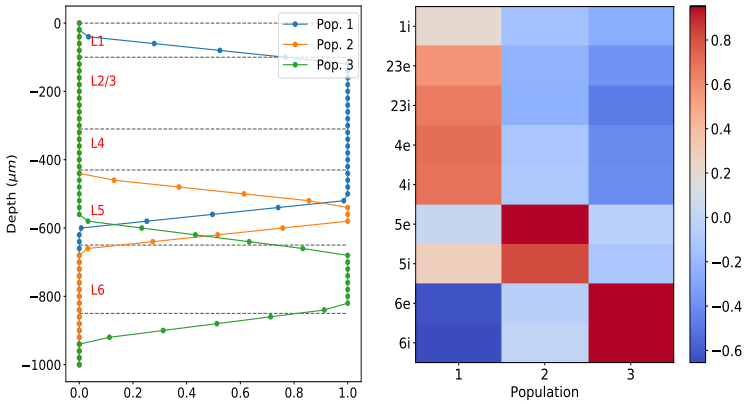

(a) (b)

Spatial Correlation

Layer	Pop 1	Pop 2	Pop 3
1i	0.21	-0.15	-0.26
2/3e	0.56***	-0.23	-0.38**
2/3i	0.66***	-0.24	-0.48***
4e	0.70***	-0.12	-0.41**
4i	0.69***	-0.11	-0.41**
5e	0.03	0.94***	-0.04
5i	0.28*	0.81***	-0.12
6e	-0.63***	-0.05	0.95***
6i	-0.65***	-0.00	0.95***

(c)

Figure 19: **(a)** Spatial profile of the the three populations estimated by LPA. **(b)** Correlation matrix of mean MUA of each layer with mean MUA of each LPA population. Mean taken from onset of white flash to 250 ms after flash offset. **(c)** Correlation coefficients for the matrix, $^*p < 0.05$, $^{**}p < 0.01$, $^{***}p < 0.001$.

To investigate whether any differences in the timing of spiking activity across the layers could be discerned with LPA, the temporal profiles of the populations were compared to the temporal profiles of the layers. The firing rates of population 1 matches the firing rates of layers 2/3 and 4, while the firing rates of populations 2 and 3 match the firing rates of layers 5 and 6, indicating that the populations can uncover the timing difference in firing rates between the upper (2/3 and 4) and the deeper (5 and 6) layers. The correlations supports this qualitative assessment, as the first population correlates more strongly and positively with layers 2/3 and 4, while populations 2 and 3 correlate more strongly and positively with layers 5 and 6 (fig. 20c and table in fig. 20d). The firing activity in layers 5 and 6 is highly synchronous in this model, thus the firing rates of both population 2 and 3 correlate very strongly with the firing rates of both layers, even though population 2 should only correspond to layer 5 cells

and population 3 should only correspond to layer 6 cells. The same is also to some extent true for population 1, though this population correlates slightly less with the layers 5 and 6 than population 2 and 3 do.

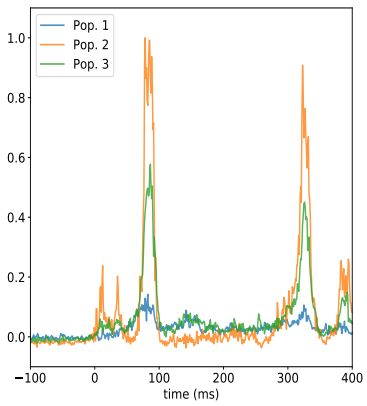

(a)

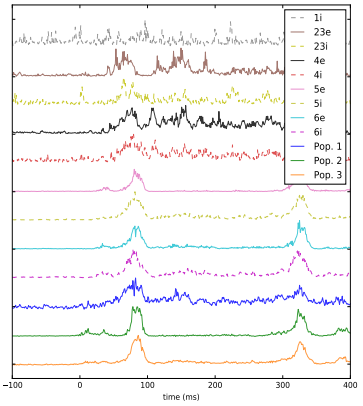

(b)

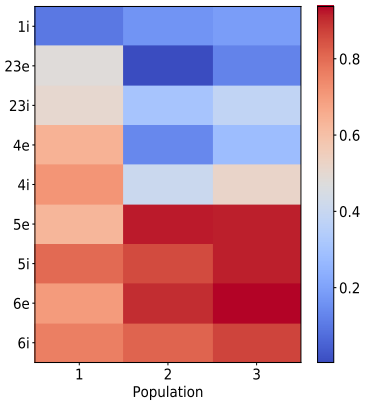

(c)

Temporal Correlation

Layer	Pop 1	Pop 2	Pop 3
1i	0.10*	0.16***	0.19***
2/3e	0.48***	0.00	0.13**
2/3i	0.50***	0.30***	0.38***
4e	0.64***	0.14**	0.28***
4i	0.71***	0.41***	0.52***
5e	0.63***	0.92***	0.91***
5i	0.80***	0.85***	0.91***
6e	0.70***	0.90***	0.93***
6i	0.76***	0.81***	0.86***

(d)

Figure 20: **(a)** Firing rates of cells in each layer and three populations estimated from LPA. **(b)** Correlation matrix between the firing rates of layers and firing rates of populations. **(c)** Correlation coefficients for this matrix, $^*p < 0.05$, $^{**}p < 0.01$, $^{***}p < 0.001$.

Local field potential The LFP estimate by LPA reproduces the qualitative features of the LFP generated with the Allen model (fig. 21), and the relative mean square error of this estimate was 0.083. Comparing the LFP attributed to each population (fig. 22) to the LFP generated from different layers in the Allen model (fig. 10), it appears that the first population (fig. 22a) captures the

LFP generated from excitatory layer 6 cells, while the second population (fig. 22b) captures LFP from both layer 5 and layer 6 excitatory cells (fig. 10f and 10h). The profiles of these two LPA-populations are quite similar (fig. 23b), but the LFP of the first population is negative in the layer 5 region, while the LFP of the second population is close to zero in this area. Since the LFP of excitatory layer 6 cells is negative in layer 5 and the LFP of excitatory layer 5 cells is positive, and the strength of their response is comparable, it may be that the positive contribution from layer 5 cells cancels out the negative contribution from layer 6 cells in the layer 5 region for the second population.

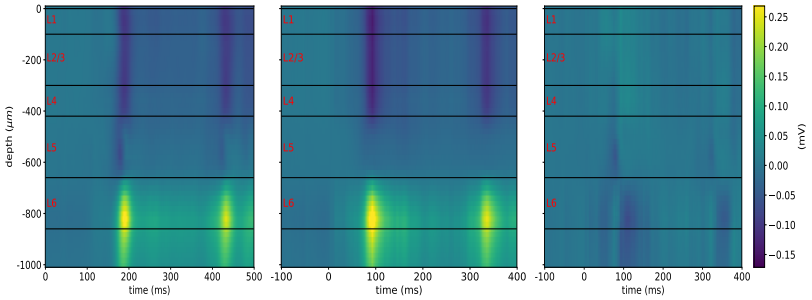

(a) (b) (c)

Figure 21: LFP before and throughout the white flash. Stimulus onset at 0 ms and stimulus offset at 250 ms. **(a)** LFP of Allen model, **(b)** LFP estimated three populations in LPA. **(c)** The absolute difference between the LFP of the Allen and the estimate from LPA.

The average spatial and temporal correlations between the population LFPs and the LFP generated by each layer supports the view that the first population is most influenced by excitatory layer 6 cells (fig. 24), as it correlates most strongly and positively with these cells. However, the average correlations also suggest that the first population is influenced by excitatory layer 5 cells and possibly inhibitory layer 4 cells too, since it correlates strongly and positively with them as well. However, the LFP generated by layer 4 inhibitory cells is probably too weak be a significant influence on this population (fig. 9, 10e, and 25e). The average correlations supports the interpretation that the second population is influenced by both layer 5 and layer 6 excitatory cells, as it correlates most strongly with these. Furthermore, the temporal correlation (fig. 20c and table in fig. 20d) between excitatory layer 5 cells and the second population is substantially stronger than it is between these cells and the first population, suggesting that the second population is indeed more influenced by layer 5 cells than the first population.

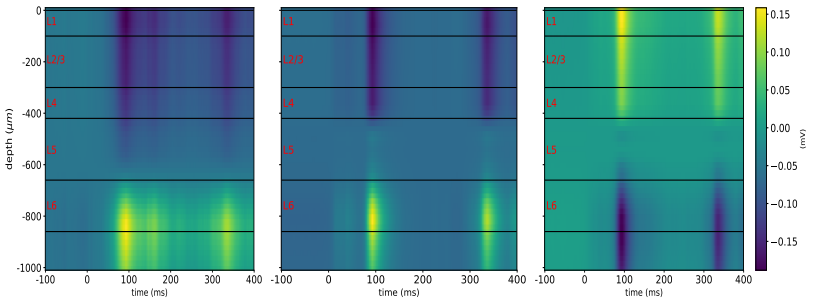

(a) (b) (c)

Figure 22: LFP of each population. **(a)** Population 1. **(b)** Population 2. **(c)** Population 3.

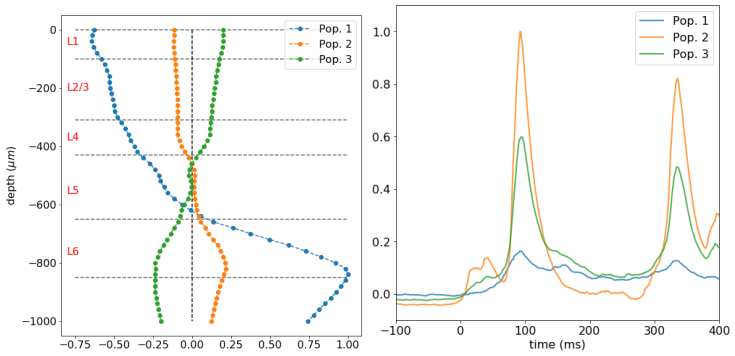

(a) (b)

Figure 23: **(a)** Spatial profile of the contributions of each population to the LFP. **(b)** Temporal profile of the LFP of each population computed by convolving firing rates of laminar populations with an optimized temporal kernel.

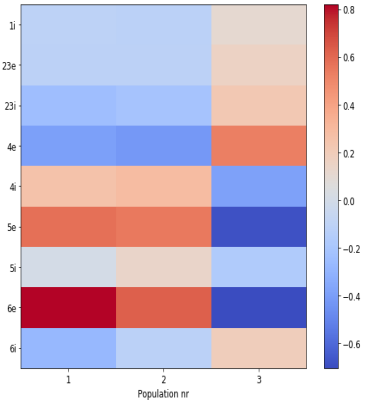

(a)

Spatial Correlation

Layer	Pop 1	Pop 2	Pop 3
1i	-0.16	-0.14	0.17
2/3e	-0.06	-0.05	0.05
2/3i	-0.27	-0.24	0.26
4e	-0.34	-0.36	0.47
4i	0.24	0.26	-0.35
5e	0.53	0.50	-0.65
5i	0.02	0.15	-0.19
6e	0.78	0.56	-0.65
6i	- 0.14	0.00	0.05

(b)

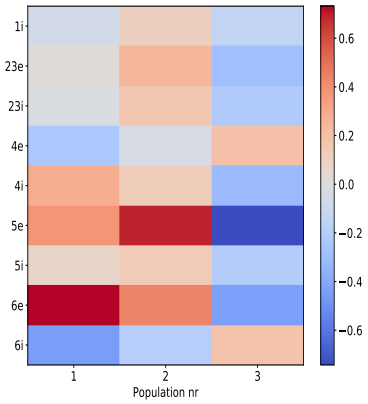

(c)

Temporal Correlation

Layer	Pop 1	Pop 2	Pop 3
1i	-0.06	0.11	-0.14
2/3e	0.00	0.24	-0.29
2/3i	-0.03	0.16	-0.22
4e	-0.24	-0.05	0.20
4i	0.29	0.12	-0.31
5e	0.37	0.69	-0.74
5i	0.07	0.14	- 0.21
6e	0.73	0.43	- 0.44
6i	-0.45	-0.20	0.18

(d)

Figure 24: Correlation between the spatial distribution **(a** and **b)** and temporal profiles **(c** and **d)** of LFP from each layer and LFP of each population.

That the first population is seemingly only influenced by layer 5 and 6 cells is not quite in line with what one would expect from the classification of populations based on the MUA, since the first population should contain LFP caused by firing of action potentials in layer 2/3 and 4 too (see section on multi-unit activity above). Although it is possible that the LFP generated by cells in layers 5 and 6 is caused by input from layers 2/3 and 4, it seems more plausible that it is caused by the firing in layers 5 and 6 itself, since the MUA was also substantially stronger in these layers (fig. 18. See also appendix section B.). Furthermore, one would still expect at least some positive correlation between the first population and the LFP generated by layers 2/3 and 4 if it is the firing of action potentials in these layers that causes the LFP attributed to population 1, but both the spatial and temporal correlations are largely negative. To

assess this qualitatively as well, the LFP generated from each layer is plotted on individual scales for visibility in figure 25, and the LFP generated by layers 2/3 and 4 are summed in figure 26. The LFP in response to the stimulus 80-90 ms generated by layers 2/3 and 4 is indeed largely localized in the upper layers, and additionally, it is too weak to be source of the LFP attributed to the first population (fig. 26b).

That the second population is influenced by layer 5 cells is in line with expectations, since this should contain some LFP caused by firing of action potentials in this layer. The overlap between the LFP generated by layer 5 and layer 6 excitatory cells means that it is going to appear to be influenced by layer 6 cells when it picks up the LFP generated by layer 5 cells, even though in reality it might not be notably influenced by the layer 6 cells. However, it seems that the influence from layer 6 cells on this population is greater than what one would expect solely from this overlap, since the LFP in response to the stimulus 80-90 ms is close to zero in the layer 5 region, where the LFP from layer 5 cells is strongly positive, and there is a positive contribution to the LFP of this population further down in layer 6 than the positive contribution from layer 5 cells stretches (fig. 25f). The third population (fig. 22c) seems to just be a mirror image of the second population, both qualitatively and based on the correlation patterns. Its positive LFP in response to the stimulus in the upper layers could indicate that it picks up the LFP generated in the upper layers after receiving input from the firing in layer 6, but the amplitude of the LFP in this population is too large relative to the amplitude of the LFP actually generated by the upper layers (fig. 26a). The fact that the third population appears to be a mirror image of the second population could also be an indication that assuming three populations may amount to overfitting, since it seems that it does not carry any information that is not already represented by the first two populations.

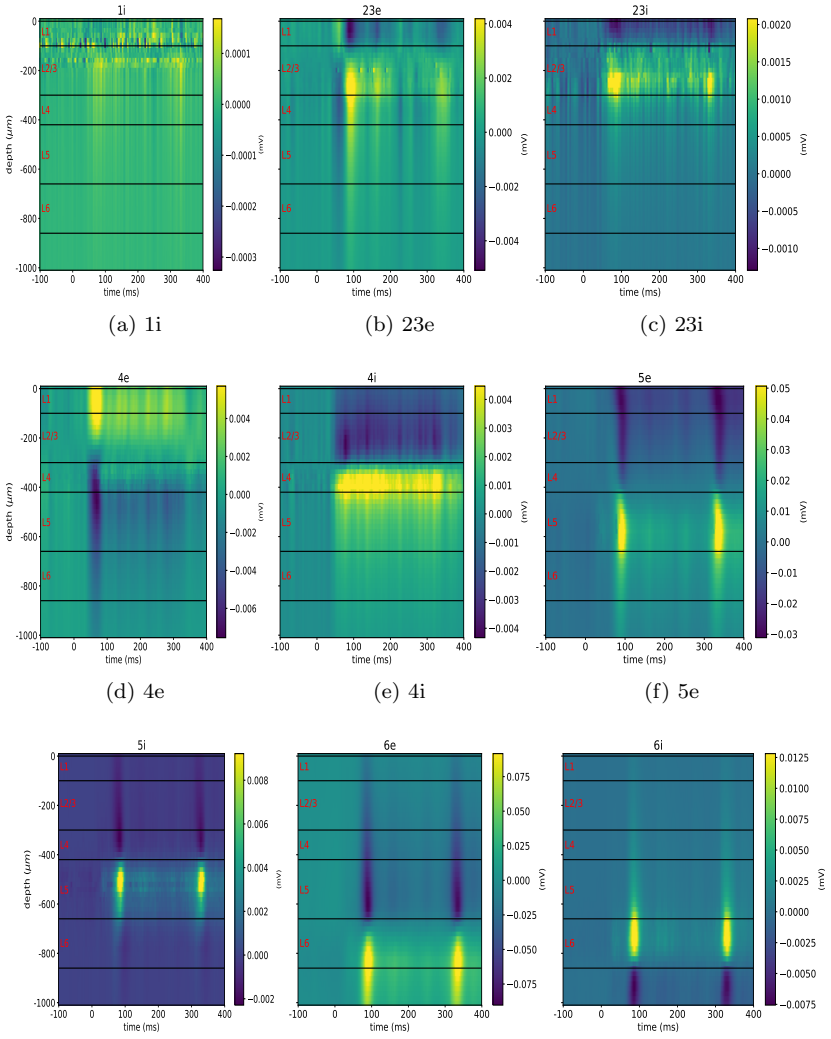

(a) 1i (b) 23e (c) 23i

(d) 4e (e) 4i (f) 5e

(g) 5i (h) 6e (i) 6i

Figure 25: LFP of excitatory (e) and inhibitory (i) cells in each layer. Scales are individual for visibility.

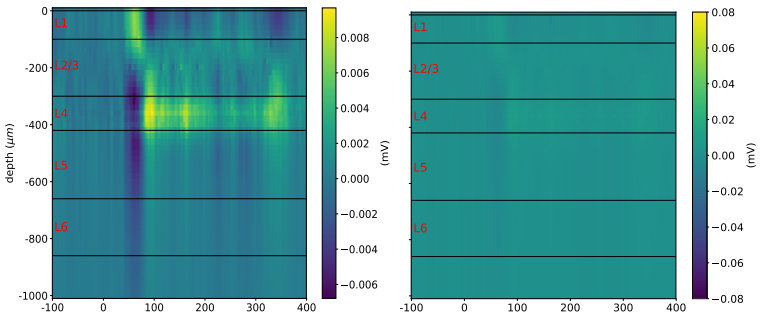

Figure 26: Summed LFP of layers 2/3 and 4 **(a)** on a small scale for visibility, and **(b)** on the scale that was used in figure 10

3.2.3 Number of populations

In the analysis above, three populations was chosen in order to have a more 1-to-1 comparison with the results from PCA. However, this need not be the most reasonable division of populations. There are five layers in the primary visual cortex (six if you layers 2 and 3 are not grouped together), so in principle there are five potential populations to be found, but some layers may contribute so little or have dynamics that are so similar to the neighbouring layers that they cannot be distinguished from each other, and can potentially be viewed as one population. One way to determine how many populations to assume is to compute the relative mean square error of the LPA estimates as a function of the number of populations, and set the population number to be the number at which the error begins to plateau. The reason for using the population number where the error begins to plateau instead of the population number that gives the lowest error is that the error will always decrease with more populations, since there are more parameters to fit in the optimization procedure, but at some point the algorithm is just fitting noise and more populations does provide valuable contributions. The error in the MUA estimate begins to decrease more slowly after three populations assumed, and hardly decreases at all from five populations to six (fig. 27a). This suggests that three, four, or five populations are the best candidates for the number of populations to assume. On the other hand, the error in the LFP estimate also begins to decrease less after three populations (fig. 27b), but as we saw above, when three populations was assumed the third population did not appear to contain any information that was not already represented by the first two populations, so even though the fit is better, the extra population need not add anything to the analysis. In the following the results from doing LPA with different numbers of populations assumed will be presented to study the effect of this assumption, and to evaluate what would be the best assumption for the Allen model.

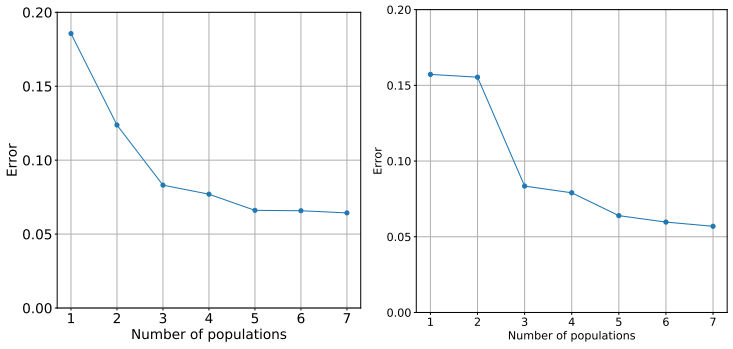

(a) (b)

Figure 27: Relative mean square error of the LPA estimate as a function of the number of populations for **a** MUA and **b** LFP

MUA

4 populations The two populations that covered parts of layer 5 and the whole of layer 6 when three populations was assumed are also present when four populations is assumed (fig. 19a and 28a), but the population that covered layer 2/3 and 4 is now split into two populations - one covering the whole of layer 2/3, and one covering parts of layer 4 and parts of layer 5. However, the correlations with the mean MUA of the layers shows that the populations covering layer 2/3 still correlates positively and significantly with layer 4, and the second population correlates more strongly with layer 5 cells than layer 4 cells (fig. 28b and table in fig. 28c). Hence, this division into four populations does not follow the layer borders of the Allen model as faithfully as the division into three populations.

As for the temporal dynamics, the qualitative form and timing of stimulus response in the firing rates of the first population resembles the firing rates of both layer 2/3 and layer 4 cells, while the timing of the stimulus response of the second population seems to be more in line with layers 5 and 6 (fig. 29a). This is also demonstrated by temporal correlations, with the first population correlating most strongly with the upper layers (though the correlations are actually equally strong with inhibitory cells in the deeper layers) (fig. 29b and 29c). The second population does correlate more strongly with layer 4 cells than populations 3 and 4 do, but it still correlates most strongly with cells in layers 5 and 6.

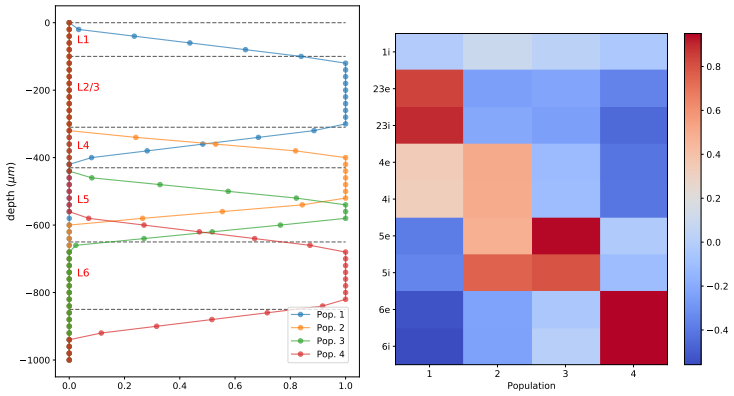

(a) (b)

Spatial Correlation

Layer	Pop 1	Pop 2	Pop 3	Pop 4
1i	0.41**	-0.15	-0.15	-0.27
2/3e	0.85***	-0.23	-0.23	-0.38**
2/3i	0.91***	-0.17	-0.24	-0.48***
4e	0.34*	0.50***	-0.12	-0.41**
4i	0.31*	0.52***	-0.11	-0.41**
5e	-0.37**	0.46***	0.94***	-0.02
5i	-0.36**	0.75***	0.79***	-0.11
6e	-0.53***	-0.24	-0.05	0.95***
6i	-0.56***	-0.24	0.00	0.95***

(c)

Figure 28: **(a)** Spatial MUA profile of the four populations estimated by LPA. **(b)** Correlation matrix of mean MUA of each layer with mean MUA of each LPA population. Mean taken from onset of white flash to 250 ms after flash offset. **(c)** Correlation coefficients for the matrix, $^*p < 0.05$, $^{**}p < 0.01$, $^{***}p < 0.001$.

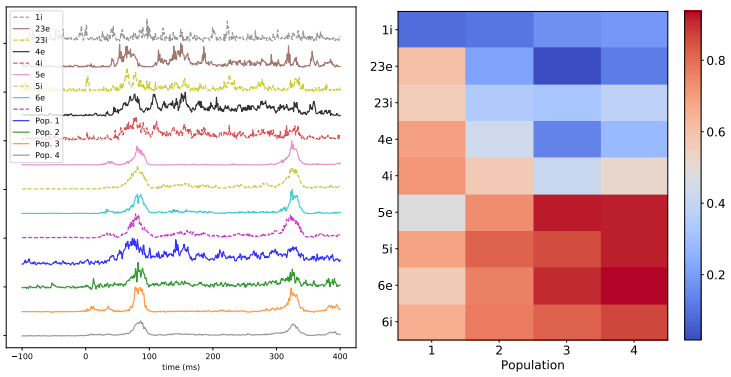

(a) (b)

Temporal Correlation

Layer	Pop 1	Pop 2	Pop 3	Pop 4
1i	0.09	0.11*	0.17***	-0.18***
2/3e	0.60***	0.21***	0.02	-0.12**
2/3i	0.56***	0.34***	0.32***	-0.37***
4e	0.69***	0.43***	0.14**	-0.28***
4i	0.71***	0.57***	0.42***	-0.52***
5e	0.48***	0.73***	0.92***	-0.91***
5i	0.68***	0.82***	0.85***	-0.91***
6e	0.57***	0.76***	0.90***	0.93***
6i	0.65***	0.77***	0.81***	0.86***

(c)

Figure 29: **(a)** Firing rates of cells in each layer and three populations estimated from LPA. **(b)** Correlation matrix between the firing rates of layers and firing rates of populations. **(c)** Correlation coefficients for this matrix, $*p < 0.05$, $**p < 0.01$, $***p < 0.001$.

5 populations With five populations assumed the spatial profiles of the first four populations are identical to the spatial profiles when four populations was assumed (fig. 28a and 30a). The fifth population is localized outside the layers of the model, in the area that is only populated by LIF-neurons that cannot produce extracelullar potentials, and hence should not be identified by an ECP based method like LPA, so this alone is a strong indication that assuming five populations is not reasonable. Additionally, the correlation patterns of the fifth population are just the same as the the correlation patterns of the fourth population, only weaker (fig. 30b). To check whether this fifth population could be predicted by LPA even though four populations was the true number of populations, synthetic data with four populations as the ground truth was constructed by copying the LPA estimate on the Allen model with four populations assumed. That way the data would largely be the same (with the exception of

the discrepancy between the LPA estimate of the MUA and the MUA from the Allen model), but the correct number of populations to assume is now known to be four populations. Then, when the LPA was performed on this synthetic data with five populations assumed, the fifth population was again localized below layer 6 (fig. 31). Together, this strongly suggests that five populations is not a reasonable assumption for this data.

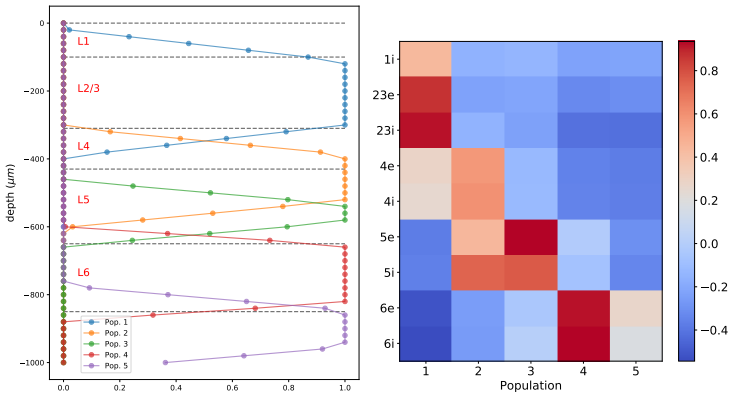

(a) (b)

Figure 30: **(a)** Spatial MUA profiles for 5 populations. **(b)** Correlation matrix of mean MUA of each layer with mean MUA of each LPA population. Mean taken from onset of white flash to 250 ms after flash offset.

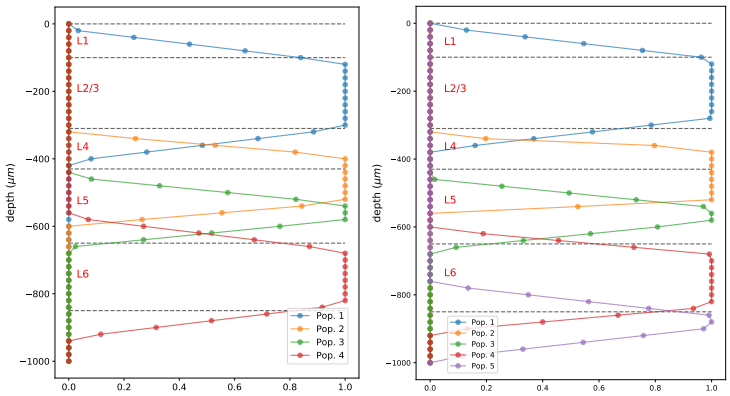

Figure 31: **a** Spatial MUA profile of synthetic data with four populations. **b** Spatial MUA profiles of populations from LPA performed on synthetic data with four populations as the ground truth and five populations assumed.

Determining how many populations to assume when the ground truth is not known. The analysis above of how many populations to assume was in large part based on comparing the predictions of the LPA to the ground truth, but LPA is intended for use on experimental data too, where the ground truth is not known. In this case, the most reasonable number of populations can potentially be assessed by computing the relative contribution of populations to the MUA and LFP estimates using equation (10).

For the LPA estimate of the MUA with three or four populations assumed, none of the populations contribute significantly less than the others (fig. 32a and 32b). When five populations are assumed, however, the fifth population contributes almost half as much as the population that contributes second least (fig. 32c). This aligns with the conclusion above that the fifth population can be discarded, and is a first indication that this approach can in fact be used to determine how many populations it is most reasonable to assume.

As for the LFP, none of the populations contribute substantially less under the assumption of three populations, but under the assumption of four populations the second population contributes between one quarter and one eighth as much as the other populations (fig. 33). Above we saw that only two populations - excitatory layer 5 and layer 6 cells - stand for almost all the variance in LFP in the Allen model, which means that we would expect that only two populations made substantial contributions to the LFP in the LPA estimate as well. Hence, this suggests that only looking at the relative contributions is not enough to determine the number of populations. However, since the sources to the LFP appeared to not be well separated into the different populations, it is difficult to assess what is a sensible number of populations to assume for the LFP in this analysis.

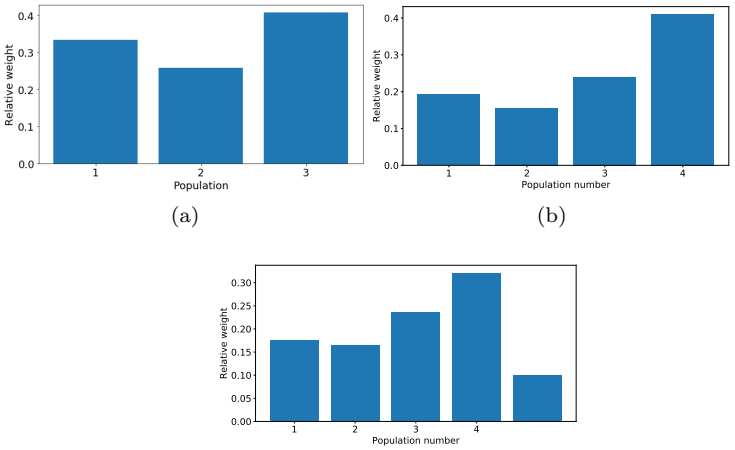

(a) (b)

(c)

Figure 32: Relative contribution of each population to the MUA estimate from LPA performed on MUA data from the Allen model with three **(a)**, four **(b)**, or five **(c)** populations assumed.

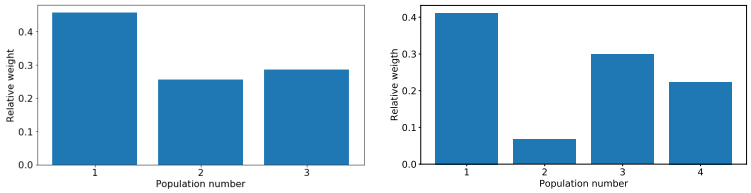

(a) (b)

Figure 33: Relative contribution of each population to the LFP estimate from LPA performed on LFP data from the Allen model with three **(a)** or four **(b)** populations assumed.

4 Discussion

4.1 Identifying laminar populations

The selected components in PCA and the populations in LPA both reproduced the original MUA well, with the first and the third principal component explaining 91.6% of the variance, and a relative mean square error of the LPA estimate equal to 0.083. Nevertheless, the laminar populations of the Allen model were not equally well identified by the the two approaches. The principal components only distinguished between upper and lower layers. Layers 2/3 and 4 were merged in the third principal component, while layers 5 and 6 were merged in the first principal component (fig. 5 and 6). Hence, it did not discern any single layers. In contrast, three of the laminar populations were correctly identified and distinguished from the other layers with LPA. When four populations was assumed, layers 2/3, 5, and 6 were identified by the first, third and fourth population, respectively. Layer 4 was not accurately identified because it was only partly covered by the second population, and this population also partly covered layer 5 (fig. 28 and 29). When three populations was assumed, the identification of layer 5 and layer 6 was the same, but layer 2/3 and 4 were now merged into one population (fig. 19 and 20).

The reason the second population covers parts of layer 4 and parts of layer 5 could be that there is substantial cross-layer correlation in the MUA, which makes them hard to distinguish. However, this possibility was investigated, and though MUA in neighbouring layers such as layer 4 and 5 do correlate more than non-neighbouring layers, the spatial correlations are not notably stronger between layers 4 and 5 than they are between other neighbouring layers (see appendix section A). A second possibility could be that the amplitudes in the MUA generated by layers 4 and 5 are the same in the area around the boundary (see appendix section B). This explanation has yet to be investigated quantitatively, so at this point in the analysis it is not clear why the second population is not confined within the boundaries of a layer, like the others are.

4.2 Separating the sources of contribution to the LFP

The excitatory cells in layers 5 and 6 stood for most of the variance in the LFP (fig. 9). This caused the decomposition by both PCA and LPA to be dominated by these two sources. The first principal component correlated positively and strongly with both layers, and explained 97% of the variance in the LFP. Neither of the other components correlated notably more with some of the layers than others. Hence, the sources to the LFP were not separated with PCA. In LPA the LFP of the first population mainly reflected the LFP generated by layer 5 and layer 6 excitatory cells, even though it was expected to also contain some of the LFP generated by layers 2/3 and 4, since under the assumption in LPA it is the firing in these layers that should cause the LFP in this population. That the first populations was almost only influenced by layers 5 and 6 could be due to synaptic input to these layers from layers 2/3 and 4, but it seems more

likely that the LFP from layers 5 and 6 originated mainly from transmembrane currents caused by their own firing, since the MUA was also much stronger in these layers (fig. 18). However, at this point in the analysis this cannot be determined, and it needs to be explored in more detail before any conclusions can be drawn.

The second population also contained LFP from both excitatory layer 5 and layer 6 cells, but was more influenced by the layer 5 cells than the first population. For this population it was in line with the classification based on MUA that it contained the LFP generated by layer 5, but the influence by LFP from layer 6 cells was greater than what could be plausibly be explained by the overlap of the LFP from layer 5 and layer 6 cells. Thus, the LFP generated by layer 5 was not fully separated from the LFP generated by layer 6 in this population. This could also be because the layer 6 cells receive synaptic input from layer 5 cells, but in any event, since the same populations were the main sources in population 1, this does not indicate that the LPA separates the different sources to the LFP well for this model. The third population was just a mirror image of the second population. It did not contain any LFP generated from layer 6 cells, even though based on the classification of laminar populations the firing in this layer should be the driver of the LFP in this population. Furthermore, the amplitude of its LFP was too large for it to represent LFP generated in upper layers after receiving synaptic input from layer 6. Thus, apart from the slightly greater influence of layer 5 cells on the LFP of the second population, the different sources to the LFP were not well separated with LPA either.

One reason the sources might be hard to separate is that both the firing of action potentials and the generation of LFP in response to the stimulus are highly synchronous between layers 5 and 6 (fig. 20b and 10f and 10h). In PCA the components are constructed based on covariance, and if the LFP generated by layers 5 and 6 is very synchronous, they are also going to covary (see appendix section A), and consequently be placed in the same component. In LPA, the population firing rates are convolved with the same temporal kernel to find the temporal profile of the LFP, so if the firing rates of the populations representing layers 5 and 6 are highly synchronous, the temporal profile of the population LFPs are going to be highly synchronous too (fig. 23a). Since the spatial profiles of the populations are fitted without constraints from the recorded LFP and the calculated temporal profiles, different populations can get similar, or even the same, spatial profiles when the temporal profiles are very similar. It appears that this occurred here. The second and third population, which were supposed to correspond to LFP caused from firing in layers 5 and 6, respectively, had almost identical spatial profiles, only with the signs reversed (fig. 23b). The first population had a slightly different spatial profile, with more negative scores in the layer 5 region, and its temporal profile was also slightly less synchronous with the other two populations (fig. 23a). However, it was still quite synchronous, and since the LFP generated by layers 5 and 6 was so dominant, this synchrony was possibly enough for it to get a spatial profile that reflected only the LFP from these two populations and none of the LFP generated by layers 2/3 and 4.

Glabska et al. (2016) developed a generalized LPA method, where the different populations could have individual and unique temporal kernels. Applying this method could reduce the synchrony in the temporal profiles, and possibly lead to a better separation of the sources to the LFP.

4.3 Determining the number of populations in LPA

Determining the most reasonable number of populations to assume based on the relative contributions of populations largely corresponded well the ground truth. When three or four populations was assumed for the MUA, none of the populations contributed substantially less than the others. When five populations was assumed, the fifth population contributed between half and one fourth as much as the others. There are five layers in the Allen model, but layer 1 contributes little to the ECP, and it was not expected to be identified as a separate population, thus a priori one might expect four populations to be the most reasonable assumption. Layer 2/3 and 4 were so similar, however, both in amplitude and timing of their activity, that they should possibly be considered as one population in this configuration of the model.

For the LFP it was more difficult to determine what would be a reasonable number of populations since the LFP from different sources were not well separated in LPA. It could be that the separation of sources would be better with two populations assumed, but since the MUA and LFP is decomposed jointly, two populations would have to be assumed for the MUA as well. Considering that the error of the estimate had not yet begun to plateau with two populations assumed, and that the classification of populations was sensible with three populations, it seems that it is not reasonable assume two populations instead of three. This is something that could be explored further in an extension of this project. Furthermore, in this project, the assessment whether a population contributed substantially less than others was rather qualitative. If it is going to be used in the future, some criterion on how little and how much less than the others a population has to contribute for the assumed number of populations to be reduced.

4.4 Implications of work and future directions

With the LPA method validated on a large-scale, biolocically detailed model, it could be used to identify laminar populations and sources to the LFP in future experimental studies. This would build on previous efforts to validate the LPA on either experimental data (Einevoll et al.; 2007) or smaller-scale, simpler models (Pettersen et al.; 2008; Glabska et al.; 2016) and potentially strengthen the case for using LPA in the decomposition of laminar multielectrode data. LPA did not fully recover the spatial distribution of the laminar populations, and it did not separate the different sources to the LFP, but it identified the laminar populations more precisely than PCA. The performance of LPA was probably affected by artefacts of the Allen model, such as the overly synchronous firing of action potentials across layers, and the dominance of the LFP generated

by layer 5 and 6 excitatory cells. This model is currently under development, so in a later configuration of the model the performance of LPA may be enhanced. Thus, one extension of this work could be to perform LPA on experimental data and compare the results, which could then guide future development of the Allen model. Furthermore, in (Einevoll et al.; 2007) the population LFPs estimated by LPA was used as a stepping stone to infer connectivity between the cortical layers, but the validity of these inferences were difficult to assess on experimental data where the ground truth is unknown. Hence, this is also something that can be attempted with this large-scale model with known connectivity.

Appendix

A: Cross-layer correlations

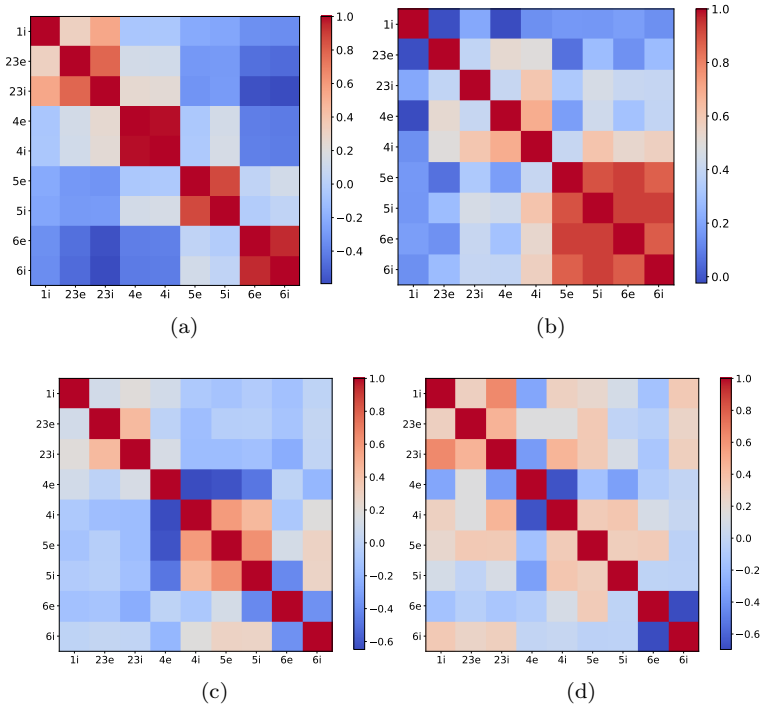

Figure 34: Correlations between layers. **(a)** Spatial correlations between mean MUA. **(b)** Temporal correlations between average MUA. Average taken over all channels. **(c)** Spatial correlations between mean LFP of layers. **(d)** Temporal correlations between average LFP.

B: Mean MUA amplitude

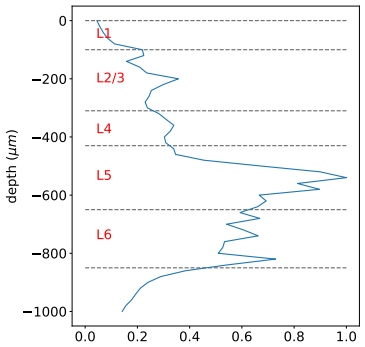

Figure 35: Mean MUA across layers. Mean taken from stimulus onset to 250 ms after stimulus offset.

C: Robustness of solution

To investigate the robustness of the solution the LPA-algorithm was run with 20 to 50 different sets of initial values for the parameters to be optimized, where the sets were given random values within the constraints. The solutions do vary with the initialization (fig. 36), but how much they vary, and whether the intitialization can change the identification of laminar populations and the sources contributing LFP depends on the assumed number of populations. The spatial profiles of the laminar populations vary enough with the initialization to affect the identification of laminar populations when the assumed number of laminar populations is four. (fig. 37c and 37d). In three out of the 20 different solutions shown in figure 37c a population is localized below layer 6, outside the model area, and one population is localized in this area even when only the 10 solutions with the lowest error were selected. When three populations is assumed, on the other hand, the solution is more stable, which is as expected since there are fewer parameters to optimize and thus fewer degrees of freedom. The predicted spatial profiles do not follow the layer boundaries as faithfully in all cases, but the classification of the populations remains the same. When only the 10 solutions with the lowest error are selected, the spatial profiles are the same as the best solution, with the exception of some fluctuations at the lower end of the populations covering layer 2/3 and 4 (fig. 37b). The picture is the same for the temporal profiles, with robust solutions when three populations are assumed, and more variability when four populations is assumed (fig. 38).

As for the LFP, the solutions are robust both when three and four populations are assumed (fig. 39 and 40), with the partial exception of the spatial profile of one population when four populations are assumed (fig. 39b). However, qualitatively the profile of this population is still the same, and furthermore, the reason this population varies is likely that four populations assumed is beyond

the number of populations that make significant contributions to the LFP.

Since the LPA does not find the same solutions regardless of the initial values of the parameters, it should be done several times with different initializations to avoid ending up in a local minimum. Optimization algorithms that might explore a greater area of the solution space should also be tried.

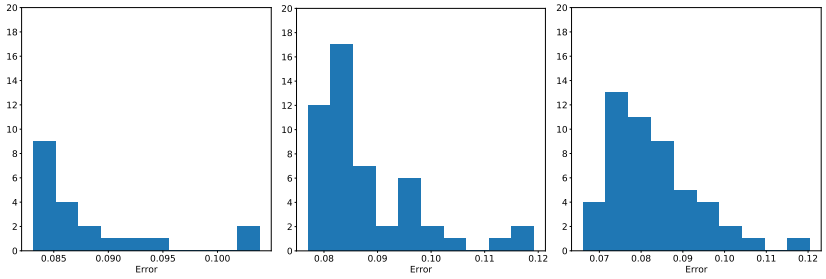

(a) (b) (c)

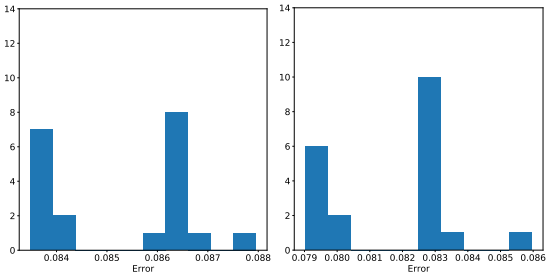

(d) (e)

Figure 36: Histogram for number of initializations that results in a given error for **(a)**, **(b)** and **(c)** MUA with three, four, and five populations, respectively, and **(d)** and **(e)** with three and four populations, respectively.

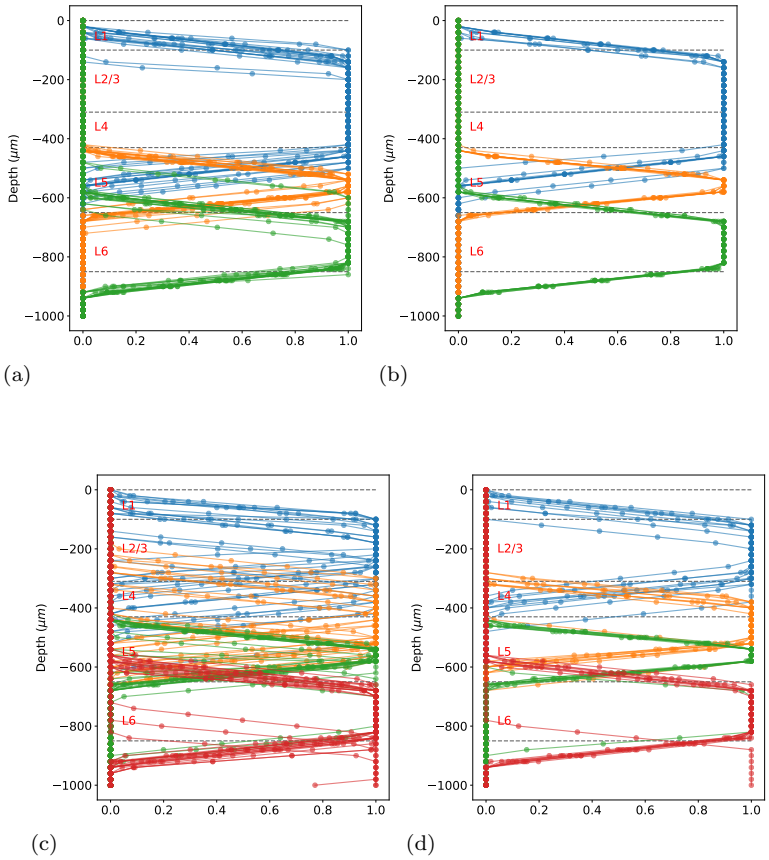

Figure 37: Spatial profiles predicted from **(a)** 20 randomly selected initializations for with three populations, **(b)** the 10 predictions that resulted in the lowest error with three populations, **(c)** 20 randomly selected initializations with four populations, **(d)** the 10 predictions that resulted in the lowest error with four populations.

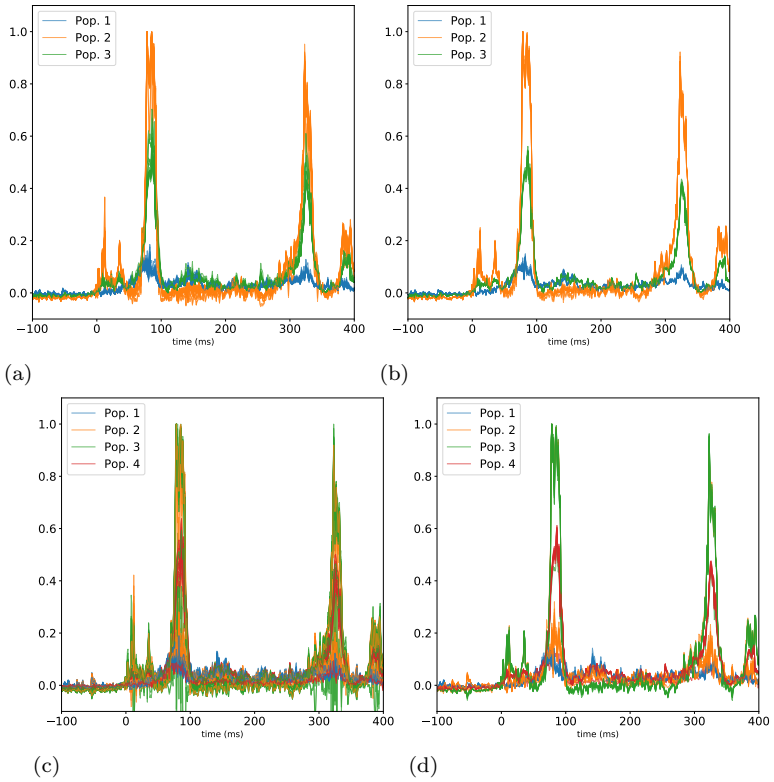

Figure 38: Firing rates predicted from **(a)** 20 randomly selected initializations for with three populations, **(b)** the 10 predictions that resulted in the lowest error with three populations, **(c)** 20 randomly selected initializations with four populations, **(d)** the 10 predictions that resulted in the lowest error with four populations.

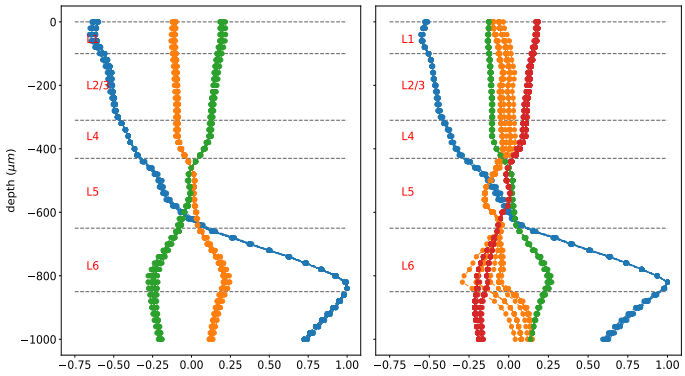

(a) (b)

Figure 39: Spatial profiles of the contributions to the LFP for 20 different initializations with **(a)** three populations and **(b)** four populations.

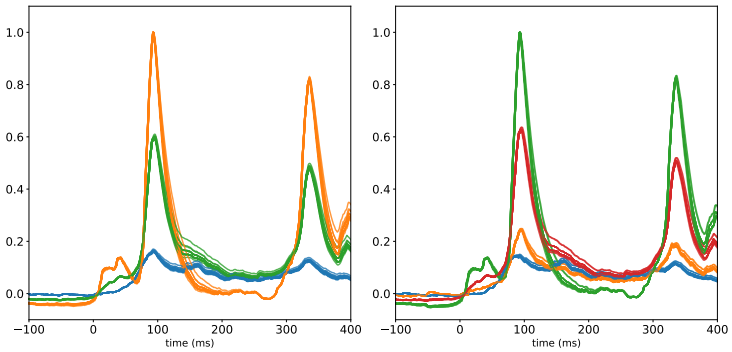

(a) (b)

Figure 40: Temporal profiles of the LFP for 20 different initializations with **(a)** three populations and **(b)** four populations.

References

Ahissar, E. and Staiger, J. (2010). S1 laminar specialization, *Scholarpedia* **5**(8): 7457. revision #150515.

Arkhipov, A., Gouwens, N. W., Billeh, Y. N., Gratiy, S., Iyer, R., Wei, Z., Xu, Z., Abbasi-Asl, R., Berg, J., Buice, M. et al. (2018). Visual physiology of the layer 4 cortical circuit in silico, *PLoS computational biology* **14**(11): e1006535.

Billeh, Y. N., Cai, B., Gratiy, S. L., Dai, K., Iyer, R., Gouwens, N. W., Abbasi-Asl, R., Jia, X., Siegle, J. H., Olsen, S. R. et al. (2019). Systematic integration of structural and functional data into multi-scale models of mouse primary visual cortex, *bioRxiv* p. 662189.

Carandini, M. (2012). Area V1, *Scholarpedia* **7**(7): 12105. revision #137292.

Einevoll, G., Pettersen, K., Devor, A., Ulbert, I., Halgren, E. and Dale, A. (2007). Laminar population analysis: Estimating firing rates and evoked synaptic activity from multielectrode recordings in rat barrel cortex, *Journal of neurophysiology* **97**: 2174–90.

Einevoll, G. T., Kayser, C., Logothetis, N. K. and Panzeri, S. (2013). Modelling and analysis of local field potentials for studying the function of cortical circuits, *Nature Reviews Neuroscience* **14**(11): 770.

Glabska, H. T., Norheim, E., Devor, A., Dale, A. M., Einevoll, G. T. and Wójcik, D. K. (2016). Generalized laminar population analysis (glpa) for interpretation of multielectrode data from cortex, *Frontiers in neuroinformatics* **10**: 1.

Hagen, E., Dahmen, D., Stavrinou, M. L., Lindén, H., Tetzlaff, T., van Albada, S. J., Grün, S., Diesmann, M. and Einevoll, G. T. (2016). Hybrid scheme for modeling local field potentials from point-neuron networks, *Cerebral Cortex* pp. 1–36.

Holt, G. R. and Koch, C. (1999). Electrical interactions via the extracellular potential near cell bodies, *Journal of computational neuroscience* **6**(2): 169–184.

James, G., Witten, D., Hastie, T. and Tibshirani, R. (2013). *An Introduction to Statistical Learning*, 4 edn, Springer, pp. 231–233 and 374–380.

Kandel, E. R., Schwartz, J. H., Jessell, T. M., of Biochemistry, D., Jessell, M. B. T., Siegelbaum, S. and Hudspeth, A. (2013). *Principles of Neural Science*, 5 edn, McGraw-hill New York, pp. 556–576.

Lindén, H., Tetzlaff, T., Potjans, T. C., Pettersen, K. H., Grün, S., Diesmann, M. and Einevoll, G. T. (2011). Modeling the spatial reach of the lfp, *Neuron* **72**(5): 859–872.

Pettersen, K. H., Hagen, E. and Einevoll, G. T. (2008). Estimation of population firing rates and current source densities from laminar electrode recordings, *Journal of computational neuroscience* **24**(3): 291–313.

Peyranche, A., Benchenane, K., Khamassi, M., Wiener, S. I. and Battaglia, F. P. (2010). Principal component analysis of ensemble recordings reveals cell assemblies at high temporal resolution, *Journal of Computational Neuroscience* **29**: 309–325.

Potjans, T. C. and Diesmann, M. (2012). The cell-type specific cortical microcircuit: relating structure and activity in a full-scale spiking network model, *Cerebral cortex* **24**(3): 785–806.

Somogyvári, Z., Zalányi, L., Ulbert, I. and Érdi, P. (2005). Model-based source localization of extracellular action potentials, *Journal of neuroscience methods* **147**(2): 126–137.

Wolfgang, K. H. and Léopold, S. (2015). *Applied Multivariate Statistical Analysis*, 4 edn, Springer, pp. 319–327.