Thalamocortical network model predicts single trial population firing rates unreliably

Eivind Hennestad

Norwegian University of Life Sciences
Department of Mathematics and Technology
Master Thesis 60 credits 2013

Preface

This thesis is the fulfillment of my Master's degree at the Norwegian University of Life Sciences. I chose to do my thesis work in computational neuroscience because I am deeply fascinated by the brain. I can sit for hours wondering about how colors arise in my mind. And why does music have a different flavor when, apparently, it is all just electrical signals going forth and back. To be able to study this using informatics and mathematics has been great.

I would like to thank my supervisor professor Gaute Einevoll for guiding me through my work. You have provided me with new perspectives and inspirations for the future. I would also like to thank my co-supervisor post doc. Birgit Kriener for constructive feedback on my writing. And thank you both for taking time to answer my questions and for bringing your good humor to our meetings.

I also want to thank my fellow students: for interesting discussions, for your lovable ways, and most importantly, for all the funny things you say and do. And I want to thank myself, for not giving up.

This last year has been a year of wonder, frustration and hopefully, new insights. Lastly, some brilliant words that sum it all up:

> *"In an honest search for knowledge, you quite often have to abide by ignorance for an indefinite period".* – Erwin Schrödinger

Ås, December 16, 2013

Eivind Hennestad

Abstract

Multielectrodes provide a powerful tool for exploring neuronal networks, but interpreting the recorded data in terms of network dynamics is difficult. Two novel modeling schemes have attempted to solve this problem, namely laminar population analysis (LPA) and the thalamocortical network model. LPA describes how to extract population firing rates from multielectrode recordings and the thalamocortical network model is estimated from these population firing rates. While these modeling methods have been successfully applied to trial-averaged data recorded from the thalamocortical loop in the rat barrel system, it is unknown whether they also work on single trial data.

This thesis aims to evaluate the thalamocortical model on single trial data. First, to learn more about the thalamocortical model, we tested it with simple input functions. Second, we looked for patterns in the single trial variability. Third, we used single trial thalamic firing rates as input in the thalamocortical model and compared the model response with cortical layer 4 population firing rates identified using results from LPA on single trial data.

We found that the thalamocortical model is unstable for large inputs. Single trial variability is large: response magnitudes typically range from two times smaller to two times larger than trial-averaged responses. Consequently, for the largest single trial inputs the model blows up. On a more general basis, the model also does not predict responses reliably due to high single trial variability in both thalamus and cortex.

In conclusion, features of the recorded activity are hidden away in trial averages, and the results demonstrate that there are many unknowns not being explained by the thalamocortical model. More effort should be put into understanding how neuronal networks handle activity in real time.

Multielektroder er et nyttig verktøy for å utforske nevrale nettverk, men det er vanskelig å tolke målte data i form av nettverksdynamikk. To nye modelleringsmetoder har forsøkt å løse dette problemet, nemlig laminær populasjonsanalyse (LPA) og en talamokortikal nettverksmodell. LPA beskriver hvordan man kan identifisere fyringsratene til nevronpopulasjoner fra multielektrodedata, og den talamokortikale nettverksmodellen er utviklet på grunnlag av disse populasjonsfyringsratene. Disse modelleringsmetodene har blitt anvendt på gjennomsnittsdata målt i hjernen til ei rotte, men det er uvisst om de også fungerer på data fra enkeltmålinger.

Denne oppgaven tar sikte på å finne ut hvorvidt den talamokortikale modellen fungerer på enkeltmålinger. For å lære mer om denne modellen har vi først testet den med enkle inputfunksjoner. Deretter har vi sett etter mønstre i variabiliteten mellom enkeltmålinger. Til slutt brukte vi talamiske fyringsrater fra enkeltmålinger som input i modellen og sammenlignet modellresponsen med de tilsvarende populasjonsfyringsratene som ble målt i lag 4 i cortex.

Vi fant ut at den talamokortikale modellen er ustabil for store input. Videre er variasjonen i enkeltmålinger stor: responsen i enkeltmålinger varierer fra to ganger mindre til to ganger større enn den gjennomsnittlige responsen. Følgelig eksploderer modellresultatet for de største inputene fra enkeltmålinger. På et mer generelt grunnlag gir modellen heller ikke pålitelige resultater på grunn av høy variasjon i enkeltmålinger, både i talamus og cortex.

Egenskaper i aktivitet fra enkeltmålinger gjemmes bort når man benytter seg av gjennomsnittsdata. Resultatene viser at det er mye som ikke kan forklares av den talamokortikale modellen. Det burde forskes mer på å forstå hvordan nevrale nettverk håndterer aktivitet i sanntid.

Contents

Chapter 1

Introduction

One way of trying to understand how the brain works is to present sensory stimuli (e.g. touch, light or sounds), and record neuronal responses [16]. Even though this outside-in approach has given many insights into the working of individual neurons and how they respond to sensory input, how the brain works as a whole is still a mystery. Neurons cooperate in large populations to produce even the simplest sensations, like that of the stroke of a finger on your hand. However, with 100 000 neurons per mm^3 of cortical tissue [19] it is impossible to monitor the activity of more than a handful of neurons at a time with current methods. A compromise is to record the joint activity of a large number of neurons. Such recordings are taken by inserting multielectrodes into brain tissue [5].

Multielectrodes capture the activity of thousands of neurons, but the recorded data is difficult to interpret [5]. The signal is a superposition of the activity from neurons close to the electrode and information about activity in different neuron populations, or how these populations work together, is hidden from view [27]. Using data recorded with laminar multielectrodes from the thalamocortical loop in the rat whisker system, Einevoll et al. [11] presented a technique called Laminar Population Analysis (LPA) to identify neuron populations and their firing activities.

Based on the LPA technique, Blomquist et al. [4] created mathematical models describing signal transfer between neuron populations responding to passive whisker movement in a rat. The models are fitted to experimental data and they represent a unique attempt at describing sensory processing in large neuronal networks by means of a simple framework. Due to large trial-to-trial variability in the data, these models are created from trial-averaged data.

Trial-averaging is a common approach in neuroscience because of variability in recordings. Even in controlled experiments where the system is subjected to identical stimulus from trial to trial, the recorded activity is highly variable [3, 22]. For example, spontaneous activity (firing of neurons in absence of sensory input) is often seen before, during and after stimulus.

Unlike the experimenter, the brain can not make averages over many trials and must instead make sense of single trial activity [3]. More and more research suggests that variability is an important aspect of neural computing. Spontaneous activity seems to take place in the same neural network as stimulus

evoked activity, and could represent memory replay of sensory experience. [30]. Masquelier [22] suggests that most of the observed noise is a sign of unexplained internal variables like attention, expectation or state of arousal.

In this thesis I have investigated single trial variability in laminar multi-electrode data. I used the same data that was also used in both the Laminar Population Analysis [11] and the thalamocortical model [4]. I also explored the thalamocortical model with simple input functions. Taking variability and model experience into account, the question I will try to answer is: How does the thalamocortical model work on single trials?

Chapter 2

General Background

This section is based on *Theoretical Neuroscience* by Dayan and Abbott [9] and it will present neurons and their electrical properties, describe how to record their electrical signals and how to treat these signals mathematically.

2.1 Neurons and Neuronal Activity

Neurons are discrete cells, specialized in the task of generating an electrical signal in response to (a potentially large number of) inputs, and propagating this signal rapidly and reliably over long distances. The morphology of the neuron reflects the function it has, the neuron consists of a cell body (soma) connected to widely branched structures of nerve fibers called dendrites and axons. The dendrites collect input signals from other neurons, and these inputs are propagated along the nerve fibers into the soma. If the input signal is of the right magnitude, an electrical pulse is transmitted along the neuron's axon.

In neurons, electrical currents consist of flowing ions, not electrons, and these flowing ions are regulated by the neuronal cell membrane. The cell membrane is almost non-permeable, but it is equipped with ion pumps and gated ion channels that selectively let ions pass back and forth from the extracellular medium. The extracellular medium is a conductive solution containing, among other things, ions and charged macromolecules. The ion pumps constantly pump ions in and out of the cell to create a concentration gradient and a potential difference across the cell membrane. When a neuron is inactive there is a higher concentration of negative charge inside the cell than outside, resulting in a potential difference across the cell membrane, *a membrane potential*, of around -70 mV. When neurons receive input through the dendrites this membrane potential will either *depolarize* (decrease) or *hyperpolarize* (increase). This change in membrane potential is the key to understanding how a neuron processes input because the gated ion channels mentioned earlier are very sensitive to changes in potential.

If the membrane potential is depolarized below a certain threshold level, a chain reaction is initiated which results in a propagating spike in the membrane potential. The first step in this chain reaction is an opening of the gates in sodium channels. Due to the higher concentration of sodium ions outside the

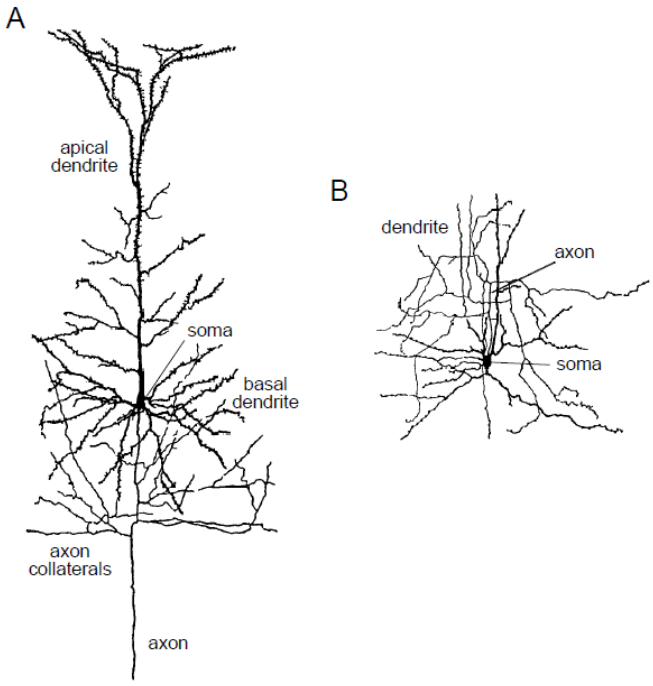

Figure 2.1: Sketches of Two Cortical Neurons. A) A pyramidal cell, an excitatory cell of the cerebral cortex. The pyramidal cell has a pyramidal shaped cell body, apical dendrites reaching upwards in the cortex, basal dendrites reaching out from the soma and both locally and distantly projecting axons. B) A stellate cell. Stellate cells are local interneurons recognizable by their star shaped appearance. They mainly branch out locally. Figure is adapted from [9].

cell, sodium ions flow into the cell, raising the membrane potential up to +40 mV. At this level of depolarization, the sodium gates close, but potassium gates open and potassium ions rush out of the cell (due to the concentration difference of potassium) bringing the membrane potential back to the resting potential. This sudden spike in the membrane potential is called an action potential. The action potential is the output signal of any neuron and it is propagated along the axon to other neurons. Typically, a neuron generates/fires several spikes every second, and a sequence of spikes is called a spike train.

The action potential spreads through the axon and is transmitted to dendrites of other neurons through a *synapse*. A synapse is a connection between an axon terminal and a protrusion on the dendrite called a spine. The connetion is not by direct contact, the axon and the dendrite are separated by a small space called the synaptic cleft (see Figure 2.2). When an action potential reaches the axon terminal, chemical messengers called neurotransmitters are released into the synaptic cleft where they diffuse towards receptor sites on the dendritic spine. There are many different neurotransmitters and all have specialized functions, but especially relevant for neuronal signaling are those neurotransmitters that

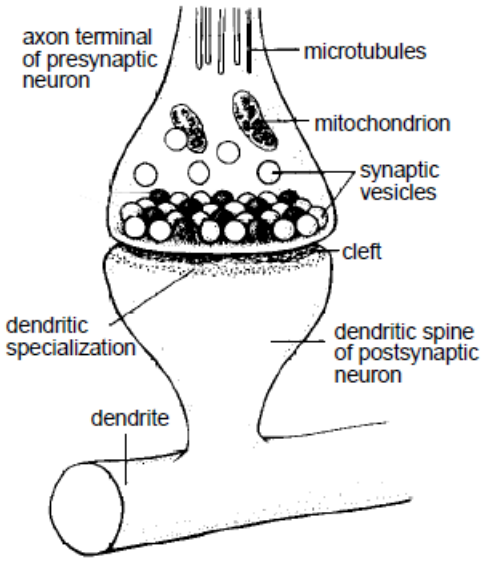

Figure 2.2: Illustration of a Synapse. A synapse is a specialized connection between axons and dendrites. Synaptic vesicles in the axon terminal contain neurotransmitters that will be released into the synaptic cleft in response to an action potential. On the dendritic side, specialized receptors respond to neurotransmitters, e.g. by opening ion channels. Figure is copied from [9].

cause opening of ion channels. This in turn leads to ions flowing through the membrane, and depending on the direction of this current flow, the membrane potential will either depolarize or hyperpolarize.

If the ions flowing through the dendritic membrane depolarize the membrane potential, the neuron comes closer to firing an action potential and is said to be *excited*. If the membrane potential is instead hyperpolarized, the neuron goes further away from firing threshold and is said to be *inhibited*. In general a neuron must be excited by many other neurons to be depolarized enough to be able to fire an action potential [21]. Neurons are highly interconnected and form thousands of synaptic connections with other neurons and only a small fraction of the synapses are activated at the same time.

2.2 Neuron Populations and Sensory Processing

The human brain is made up of about 100 billion neurons. These take part in different brain structures – of which the cerebral cortex, or *cortex*, and thalamus will be considered here. Thalamus is a walnut-sized structure in the middle of the brain. All sensory input (except for olfactory) goes through thalamus before ending up in the cortex. The cortex is the outer surface of the brain and it consists of a 2–4 mm thick sheath of neurons. In humans it is highly folded to maximize its area within in a limited volume. Cortex is responsible for sensory processing and motor control and is also related to higher cognitive thinking, and these different functions are located in different regions. Zooming in on one region will reveal subregions with clusters of highly interconnected neurons specialized at very specific processing. For example, in the visual cortex, many such clusters are connected to tiny parts of the visual field, similar to the organization of pixels in a digital image. Since these clusters are seen in

many different parts of cortex, a common hypothesis is that they represent basic functional circuits. These circuits are referred to by many as cortical columns because they often have a cylindrical shape. It is also common to divide cortex into six horizontal layers (I-VI) (Figure 2.3), with neurons having different sizes and shapes in different layers.

There is a great variety of cortical neurons, but they can be divided into two main classes, namely excitatory and inhibitory neurons. Dale's law states that a neuron either has an excitatory or an inhibitory effect on all its postsynaptic neurons. Another common way to classify neurons is based on their morphology. Again there are two main cortical types, pyramidal and stellate neurons, as shown in Figure 2.1. Their function is suggested from their appearance, pyramidal neurons collect input both from above and from the vicinity of the soma and typically project their axon downwards, while stellate neurons interact more locally.

The coding of information in the cortex is usually not performed by individual neurons. One neuron typically codes for certain features of a sensory signal, but many other neurons are needed to provide a context. For example, in the cricket, several neurons are needed to give a representation of the direction of incoming air currents. Crickets have hair sensors connected to neurons that fire whenever the hair is deflected. There are thousands of these hairs, and all of the corresponding sensory neurons send their action potentials to a set of interneurons. Some of these interneurons only respond to certain wind directions while others respond to all wind directions. However, with a set of only four neurons responding to wind directions in four different directions (e.g. left, right, forwards, backwards), accurate information about wind direction can be inferred. In the mammalian cortex, thousands of neurons are thought to participate in even simple sensory processing.

2.3 Recording of Activity

A very common tool for measuring electrical activity in the brain is an electrode. The electrode is inserted into brain tissue to record electrical potential changes in single neurons or in groups of neurons. A sharp electrode (<1 μm) can penetrate the membrane of a neuron (electrode placed inside cell) and directly measure the membrane potential, giving a precise description of the activity in that neuron [16]. Extracellular electrodes (placed near cells) on the other hand, measure changes in the potential of the extracellular medium due to currents going into, or out of neurons. This also gives information about the activity of a neuron, but in a more indirect way. Since neurons are densely packed in the extracellular medium – unless an extracellular electrode is placed very close to a neuron, it will pick up the extracellular potential due to transmembrane currents from many different neurons. Also, when the contact surface of the electrode is large, it will sample a larger area in the extracellular medium and pick up contributions from more neurons.

The potential recorded by an electrode is usually divided into a low-frequency part and a high-frequency part, the local field potential (LFP) and the multi-unit

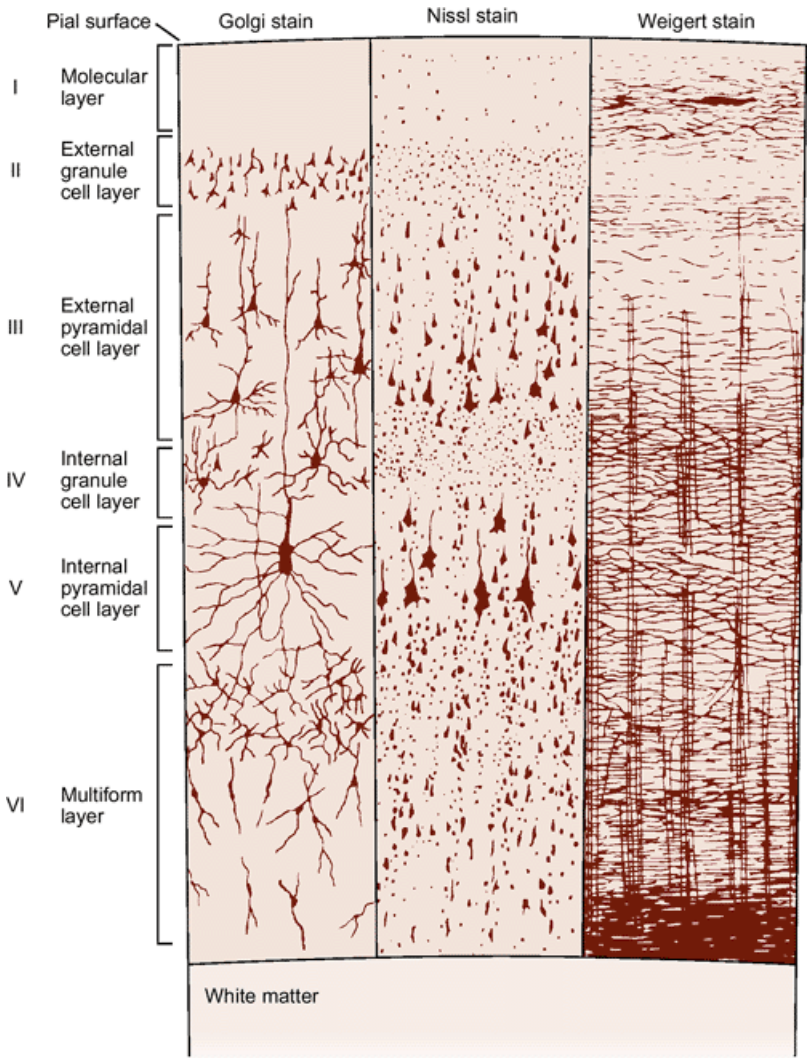

Figure 2.3: Neurons in Cortical Layers. This figure shows neurons of the cortex under different staining methods. When stained, a chemical solution is impregnated in a neuron, making it visible under light microscopy. Golgi staining reveals neuronal somas and dendrites (left), Nissl staining reveals cell bodies (center) and Weigert staining reveals myelinated axons (right). Different neuron types are visible in the different layers. Figure is copied from [17]

activity (MUA) respectively. The LFP is thought to arise from transmembrane currents in dendrites due to synaptic processing. These processes are relatively slow (low frequency), much slower than the generation of action potentials which are thought to give rise to the MUA [27].

2.4 Firing Rates

Spike trains can be described mathematically as a series of idealized pulses. The action potential is similar from neuron to neuron; it is a sharp pulse of depolarization in the membrane potential (∼100 mV), it lasts around 1 ms, it is reliably propagated along axons, and it has the same effect on axon terminals. And most importantly, it is an all or nothing process. Therefore, it can be regarded as an idealized event and the essential thing to include when describing this event is the time t at which it occurs. This is accomplished with the *neural response function*:

$$\rho(t) = \sum_{i=1}^{n} \delta(t - t_i) \tag{2.1}$$

The sum is over n spikes and t_i denotes the time of the ith spike. Here δ denotes the Dirac delta function. The Dirac delta function approaches zero everywhere except where its argument is zero; there it goes to infinity. It is defined so that its integral is one whenever $t = 0$ is within the limits of integration (otherwise it is zero):

$$\int \delta(t)dt = 1 \tag{2.2}$$

In real neurons, spike trains vary from trial to trial in a seemingly stochastic way, even when the same stimulus is presented. Therefore it is often more convenient to treat spiking activity in a probabilistic way. This is achieved by the *time-dependent firing rate*, [9], a measure of how frequent a neuron fires action potentials. It is defined by counting spikes over a time interval and dividing by the duration of this time interval:

$$r(t) = \frac{1}{\Delta t} \int_{t}^{t+\Delta t} \rho(t)dt \tag{2.3}$$

This definition uses the fact that integrating the neural response function ($\rho(t)$) over a time interval is equivalent to counting spikes. The time interval should ideally go towards zero ($\Delta t \to 0$), but in practice the value of Δt should be large enough to ensure that there are a sufficient number of spikes within the time interval. If the time interval is too short the values of $r(t)$ would always be either zero or one.

The definition of the firing rate doesn't provide a unique way to approximate $r(t)$ since the time interval in eq. (2.3) is arbitrary. Another approach to approximate the firing rate is to use a *linear filter* and *kernel*. In this method a "window" function is slid along a spike train as demonstrated in Figure 2.4 C) and D). The value of $r(t)$ at a time t then depends on the number of spikes seen through a window located at time t, and the shape of the window. Sliding

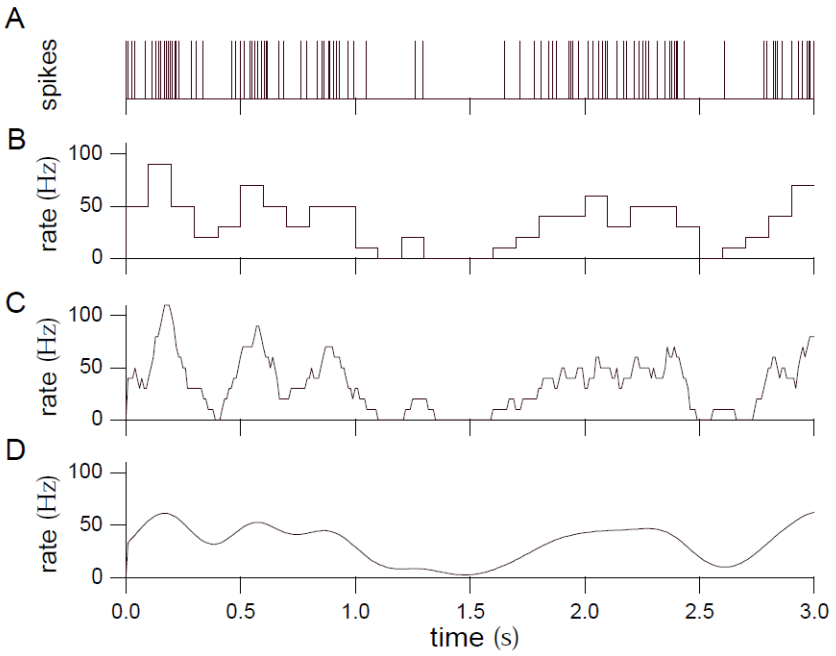

Figure 2.4: Illustration of firing rate approximation. A) An example of a spike train. B) Discrete firing rate obtained by counting spikes within a time interval and dividing by the time interval ($\Delta t = 100$ ms). C) Firing rate obtained by sliding a rectangular window function along the spike train with $\Delta t = 100$ ms. D) Same as in C), but the window function is a gaussian function with $\sigma_t = 100$ ms. Figure is adapted from [9].

a window function over the neural response function gives an approximation for $r(t)$:

$$r_{approx}(t) = \int_{-\infty}^{\infty} w(\tau)\rho(t-\tau) \tag{2.4}$$

$w(t)$ denotes the "window" function and is usually called a filter kernel. The method of integrating the product of a kernel and a function along the real line is called linear filtering.

Firing rates can be approximated from experimental data, but they also represent a quantity that can be predicted using firing rate models. In the next chapter, firing rate models based on experimentally extracted firing rates from neuron populations are presented.

Chapter 3

Specific Background

This chapter describes extraction of population firing rates from laminar electrode data and the application of a firing-rate model to these population rates.

3.1 Example System: Rat Barrel Cortex

The whisker system in rats provides a good framework for studying structure and function of cortical columns. When a whisker on the rat's snout is deflected, a signal is sent along nerves through the brain stem and thalamus into a structure called barrel cortex (illustrated in Figure 3.3). Barrel cortex is located in the somatosensory cortex in rats, and it got its name because neurons in barrel cortex are arranged in columns that look like barrels [36]. These "barrels" are ordered in a grid-like pattern and this pattern corresponds to the way whiskers are ordered on the rat's nose. Unsurprisingly, flicking a whisker at one position in this grid will lead to firing of neurons in a corresponding barrel. In fact, all major whiskers are connected to their "own" barrel, and Figure 3.1 illustrates this mapping from whiskers to barrels [12]. The neurons that make up a barrel form connections upwards and downwards in a columnar fashion, making barrels a prime example of the concept of a cortical column. The relationship between stimulus and neuronal responses in a cortical column can therefore easily be studied by stimulating whiskers and recording neuronal responses.

Equivalents of the barrels are also found in earlier steps of the neural pathway [12]. In the brain stem, whisker signals are distributed to four clusters of neurons called the trigeminal nuclei. In three of them, the whisker map is visible as darker spots surrounded by lighter areas when stained, and the darker spots are named *barrelettes*. From the trigeminal nuclei, whisker signals are distributed to many different parts of the brain, of which the ventral posteromedial nucleus (VPm) in thalamus is relevant with regards to sensory processing. In the VPm, the barrel/barrelette equivalents are also present, and they are called barreloids. From these barreloids, whisker signals are transferred mainly to layer IV in barrel cortex [12].

The rat uses its whiskers for orienting and exploring its immediate surroundings and diverse forms of information are therefore represented in barrels. Barrel

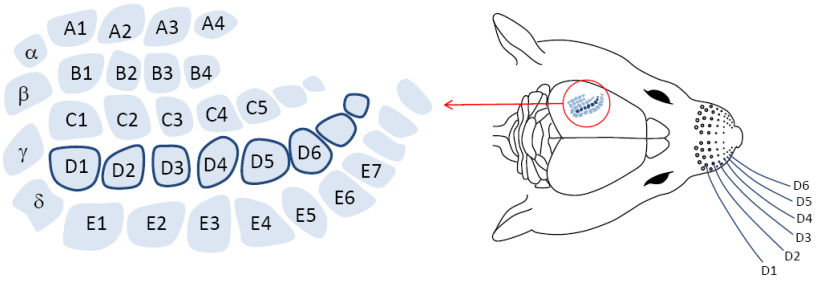

Figure 3.1: Whisker mapping to barrel cortex. The pattern of whisker placement on a rat's snout is mapped onto the rat barrel cortex. Whiskers D1, D2, ..., D6 project signals to the corresponding barrels D1, D2, ..., D6. Image is copied from Neural Lab (web site) [1]

neurons are for example sensitive to amplitude, velocity and direction of whisker deflection [35, 28, 34]. The sensitivity of neurons to these stimulus features are preserved through the neural pathway from whisker to cortex [31, 32]. This preservation of information through the neural pathway makes it easy to study the signal at different relays and compare the transformation from one to the next.

3.2 Experimental Data

In this project I have looked at recordings of extracellular potentials from the barrel cortex in rats, an experiment described in detail in Einevoll et al. [11]. Here the researchers used a computer controlled wire loop to flick a whisker on the rat's snout while simultaneously recording the electrical response in the rat's brain. The whisker was flicked using 27 different stimulus conditions, in a combination of 3 different amplitudes (0,40–1,20 mm) and 9 different rise times (20–100 ms). The wire loop delivered these stimulus conditions in a random sequence, repeating each of them 40 times.

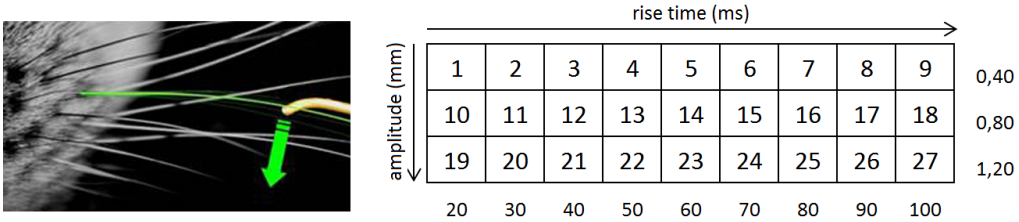

Figure 3.2: Diagram of stimulus conditions. Left: Illustration of whisker flicking. Photograph modified from [**?**]. Right: Stimulus conditions are numbered from 1 to 27, amplitudes are shown in y-direction and rise times are shown in x-direction

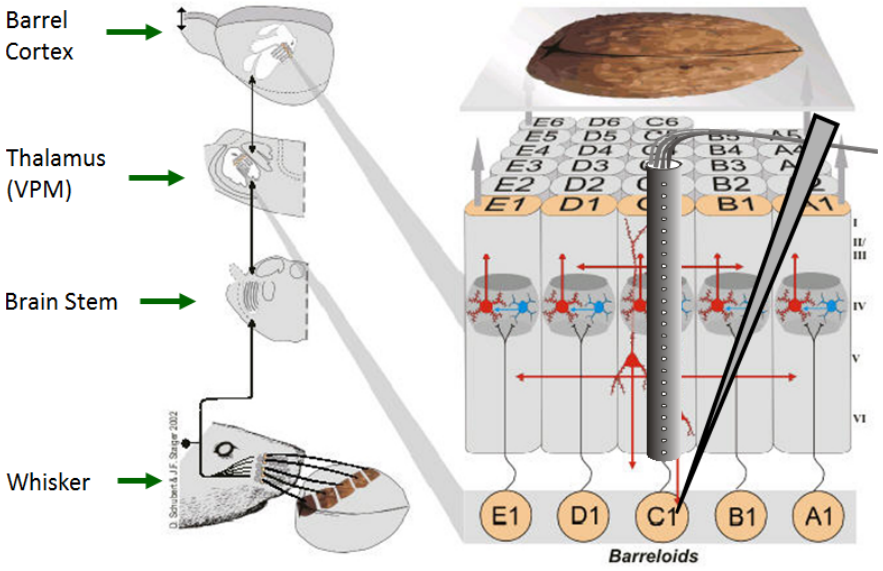

Figure 3.3: Experimental Setup. Figure illustrates the signal's pathway from whisker to barrel cortex and demonstrates how a multielectrode records the electrical signal in a cortical column and how a single electrode records the electrical signal in the corresponding barreloid in thalamus [4].

A *linear array multielectrode* was inserted into the cortical column responding to the stimulated whisker to record the extracellular potential. The 23 electrode contacts were arranged at equidistant intervals of 0.1 mm giving a total electrode length of 2.2 mm, about the same as the average thickness of rat cortex. Thus the electrode records activity from all six cortical layers. In addition, a single electrode was inserted into thalamus to record the activity in the thalamic neuron population projecting signals to the cortical column.

For each trial the electric potential was recorded in time-steps of 0.05 ms for a total duration of 400 ms, 100 ms before whisker stimulation and 300 ms after. The recorded potential was filtered into two signals, a low frequency part (0.1-500 Hz) with a time resolution of 0.5 ms and a high frequency part (150-5000 Hz) with a time resolution of 0.05 ms. These two signals are referred to as the LFP, and the MUA.

The MUA and the LFP from each trial are multichannel time series of potentials, and they can be represented as two-dimensional arrays $\mathbf{\Phi}^{ij}_{MUA}$ and $\mathbf{\Phi}^{ij}_{LFP}$ of values, where each value is the MUA (or LFP) at electrode contact i and time step j.

To prepare the data for analysis, the MUA is band-pass filtered (750-10000Hz) using a second-order Butterworth filter and then rectified to provide a positive signal. At last, the MUA signal is decimated by 10 along the time axis to have the same time resolution as the LFP. The following sections are based on the trial-averaged MUA, as shown in Figure 3.5 [4].

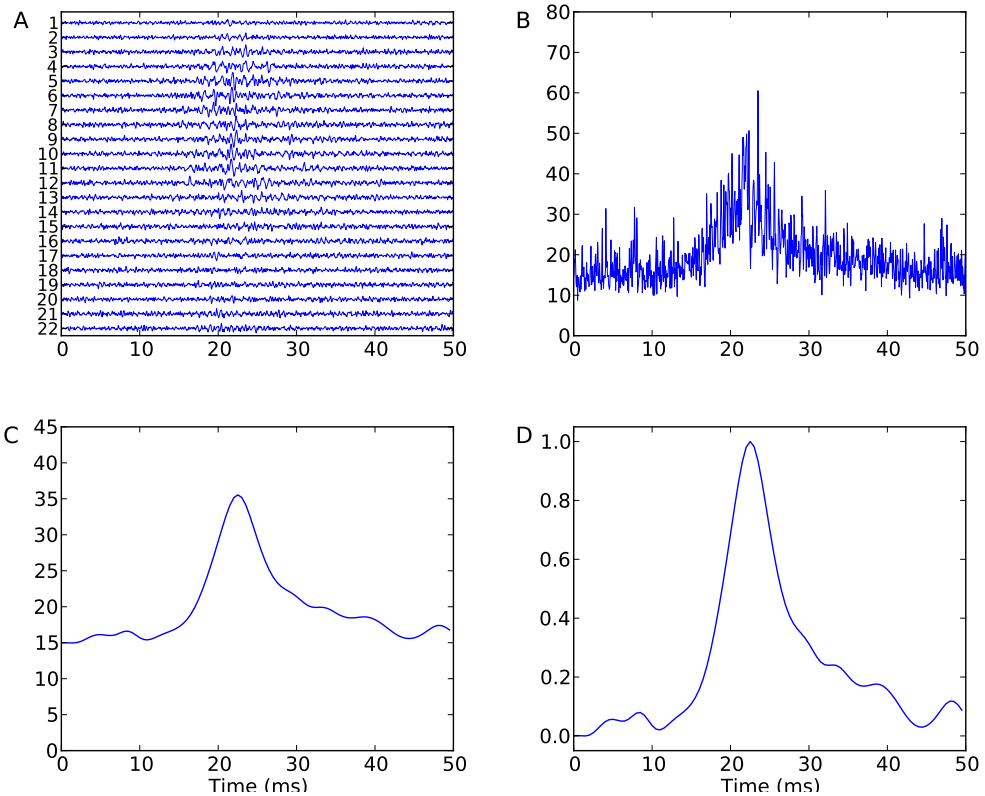

Figure 3.4: Representations of the MUA. A) Raw MUA signal from 22 electrode contacts during 50 ms after stimulus. Contact 1 is positioned close to the cortical surface and the contacts go through all cortical layers. B) Rectified and channel-summed MUA. C) Low-pass filtered version of the MUA in B. D) Normalized version of the MUA in C. The MUA is shifted to make the minimum value zero and then divided by the maximum to make the peak value one.

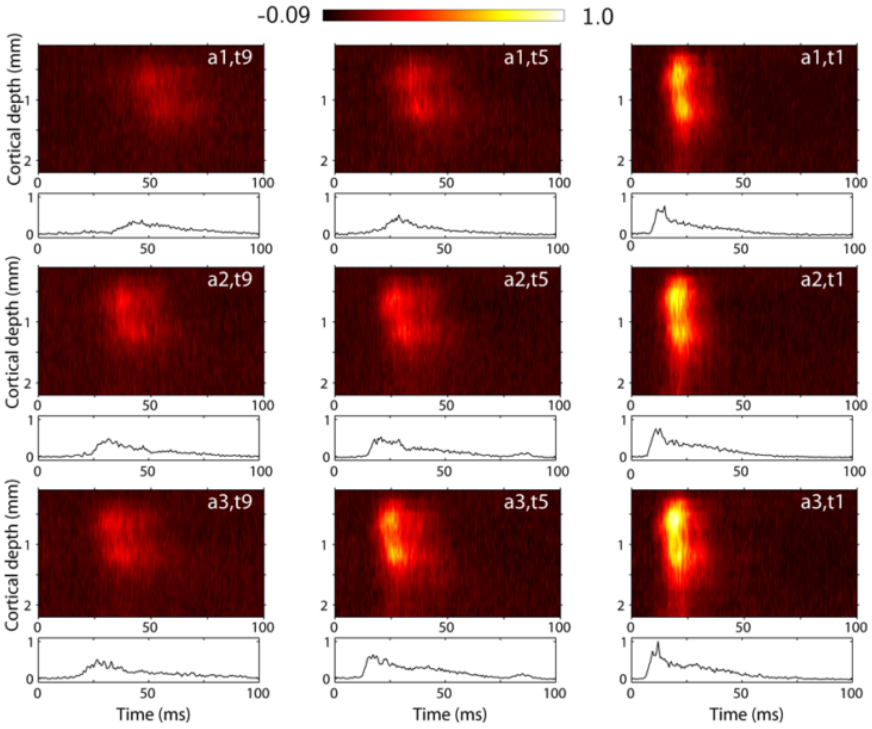

Figure 3.5: Trial-averaged MUA for 9 Stimulus Conditions. Trial-averaged MUA for stimulus conditions 1, 5, 9, 10, 14, 18, 19, 23, 27 from top left corner to bottom right corner. Color image shows the cortical MUA and line plot shows thalamic MUA. Figure is copied from Blomquist et al. [4]

3.3 Laminar Population Analysis

Einevoll et al. [11] developed the method "Laminar Population Analysis" (LPA) to extract information about neural networks from laminar electrode recordings of brain activity. A key step in the analysis is to identify neuron populations and estimate their firing rates based on the high frequency part of the recorded extracellular potentials, the MUA.

Neuron populations and their firing rates can be estimated by reducing redundancy of information present in the MUA recorded by laminar electrodes. If it is assumed that neuron populations are stacked on top of each other within a cortical column being probed by a laminar electrode, and that the number of neuron populations are smaller than the number of electrode contacts, it follows that activity of one population is captured by more than one of the electrode contacts. In barrel cortex, neurons within neuron populations are typically activated at the same time in response to input. Furthermore, neuron populations are activated in a sequential pattern, and it is likely that the MUA will be redundant in the sense that several data channels contain similar MUA patterns. Population MUA can therefore be estimated by combining the signal from mutually

redundant channels into one, and the location of the corresponding population could be inferred based on which channels are combined.

To solve this problem it is useful first to consider the MUA as the matrix product of two matrices:

$$\mathbf{\Phi_{MUA}} = \mathbf{AB} \tag{3.1}$$

The simplest way to find the matrices **A** and **B** would simply be to let **A** be an identity matrix of dimension equal to the channels or rows in $\mathbf{\Phi}_{MUA}$, and **B** be the MUA activity. **A** would be a transformation matrix linking the MUA in **B** to the original MUA. This example is trivial, but there are infinitely many other, more interesting possibilities.

A common statistical tool for reducing redundancy in data is Principal Component Analysis (PCA) [20]. In PCA, the matrices **A** and **B** would be determined in such a way that **B** would be a matrix of uncorrelated MUA time series (temporal scores), arranged in order of decreasing variance, and **A** would contain the information (spatial profiles) linking these temporal scores to the MUA recorded at each contact. These uncorrelated temporal scores could be interpreted as activity of neuron populations, an analysis like this was performed in e.g. Di et al. [10] on CSD data to identify population activities.

However, when applied to MUA, it is hard to interpret the temporal scores as population firing rates. The temporal scores may contribute positively to some of the channels, but negatively to others. Since the population firing rates are positive by definition, they should not contribute negatively to the recorded MUA. In addition, the temporal scores can contribute to the MUA along the entire cortical depth, while evidence suggests that the recorded firing activity of laminar populations will be spatially confined to distinctive regions [33].

Somogyvári et al. [33] recorded extracellular potentials using laminar electrodes and identified the action potential signature of several neurons using spike sorting. Their results indicated that an electrode contact could pick up action potentials from neuronal somas located up to 0.1 mm above or below the contact position. This information provides a way to solve the problem in eq. 3.1 in such a way that population activity can be revealed.

The MUA can be modeled as contributions from several neuron populations, in the form of eq (3.1):

$$\mathbf{\Phi^{mod}_{MUA}} = \mathbf{Mr}^T \tag{3.2}$$

Here $\mathbf{\Phi^{mod}_{MUA}}$ is a $N_{ch} \times N_t$ matrix, **M** is a $N_{ch} \times N_{pop}$ matrix containing weights determining the contribution of each population to the channels and **r** is a $N_t \times N_{pop}$ matrix where each column i correspond to population activities of each population. N_{ch} denotes the number of electrode recording channels, N_t denotes the number of sampling points in time and N_{pop} denotes the number of neuron populations.

Based on the assumption of stacked populations and the results from Somogyvári et al. [33], Einevoll et al. [11] modeled the population distributions as trapezoidal boxes. The box is centered at a point z_0 along the contact positions and has a constant height along a length a. On the border between overlapping populations, the value of the box goes to zero over an interval b.

When the population distributions are defined, an estimate of the population firing rates can be found by solving eq. (3.2) for $\mathbf{r}$ using the pseudoinverse of $\mathbf{M}$:

$$\mathbf{r_{est}} = (\mathbf{M}^{\dagger}\mathbf{\Phi_{MUA}}) \tag{3.3}$$

The model data can now be computed as the matrix product between the population distribution and estimated population firing rates

$$\mathbf{\Phi^{mod}_{MUA}} = \mathbf{M}\mathbf{r_{est}}^{T} \tag{3.4}$$

The best spatial distribution of populations are found by minimizing the difference between the modeled data and the experimental data. So, if it is possible to identify a spatial distribution of populations and a corresponding set of population firing rates which gives a dataset similar to the experimental data, it is reasonable to assume that the identification is plausible. The difference between the modeled data and the experimental data is quantified using a relative mean square error, given as:

$$e_M = \sum_{i=1}^{N_{ch}}\sum_{j=1}^{N_t}[\mathbf{\Phi^{exp}_{MUA}}(z_i,t_j) - \mathbf{\Phi^{mod}_{MUA}}(z_i,t_j)]^2 / \sum_{i=1}^{N_{ch}}\sum_{j=1}^{N_t}[\mathbf{\Phi^{exp}_{MUA}}(z_i,t_j)]^2 \tag{3.5}$$

One way to optimize this equation is to run an iteration based on trial and error. First, start out with some randomly chosen parameters, then compute the error using these parameter. Then, change the parameters a little bit and compute the error again. If the error decreases, these parameters are better. If it increases, the old parameters were better. Now, choose the best parameters, change them a little bit and repeat. This iteration should be run through a few thousand cycles. Also, by running the iteration many times with different initial parameter values, the risk of finding a local minimum in the parameter space is reduced. A guideline could be to repeat the iteration until the same minimum value for e_M is found several times.

Einevoll et al. [11] identified four different neuron populations using the LPA method, and the location of these populations agreed with findings from similar studies as well as known micro-anatomy of the cortex. The identified populations were interpreted as i) a population of layer 2/3 pyramidal cells, ii) a population of layer 4 stellate cells, iii) a population of layer 5 pyramidal cells and iv) a population of layer 6 (or deep layer 5) pyramidal cells.

3.4 Thalamocortical Model

Blomquist et al (2008) derived a thalamocortical firing-rate model which predicts the response in the cortical layer 4 population given a thalamic firing-rate input. The model was constructed using the thalamic firing rate and the layer 4 population firing rate obtained from multielectrode recordings and laminar population analysis.

A firing-rate model is traditionally used to predict the output firing rate of a neuron based on the inputs this neuron receives. Here the same idea is

instead applied to the activity of entire neuron populations. Blomquist et al. [4] assumed that the output firing rate of an entire neuron population could be predicted based on the total firing rate in a previous neuron population. Pettersen et al. [26] modeled the MUA from a population of layer 5 pyramidal neurons and found that the trial-averaged MUA was proportional to underlying firing rates. Therefore, the experimental MUA is suitable for identifying the optimal model structure and model parameters. Trial-averaged MUA was used to construct a model that performed well across all stimulus conditions. As in the single neuron case, the input firing rates to the neuron population are transformed into currents received by the neuronal somas, and the sum of these currents in turn generate an output firing rate.

The layer 4 population starts firing in response to a stimulus because it receives excitatory input from the thalamus population. However, the firing-rate dynamics during a stimulus response also depends on excitatory and inhibitory recurrent inputs from interneurons in the layer 4 population itself. The inhibitory interneurons in layer 4 are driven by excitatory input from both thalamus and other neurons in layer 4. A model structure that captures these contributions must therefore include terms for both feedforward excitation and inhibition from thalamus and recurrent excitation and inhibition from layer 4 interneurons.

A simple way to formulate this model is:

$$r_4(t) = F_4(I_4(t)) \tag{3.6}$$

This is an abstract description where the layer 4 population firing rate $r_4(t)$ is determined by the amount of electrical current I_4 reaching the somas of the neurons in the population. The activation function F_4 determines the conversion of this current into the population firing rate.

The activation function has the form:

$$F(I) = \begin{cases} 0 & \text{if } I_n < I_n^\dagger \\ a_n(I_n - I_n^\dagger) & \text{if } I_n^\dagger \leq I_n \leq I_n^*) \\ a_n(I_n - I_n^\dagger) + b_n(I_n - I_n^*)^2 & \text{if } I_n^* < I_n \end{cases} \tag{3.7}$$

The activation function is linear for current values between $I_n^\dagger$ and I_n^* quadratic above I_n^*. The constants a_n and b_n determine how fast the function grows. For low values ($I_n < I_n^\dagger$) the firing rate will be zero.

The current I_4 depends both on the contributions from feedforward thalamic neurons and recurrent interneurons:

$$I_4(t) = x_{E4}(t) - x_{I4}(t) + f_T(t) \tag{3.8}$$

Here x_{E4} is the current from excitatory interneurons, x_{I4} is the contribution from inhibitory interneurons and $f_T(t)$ is the joint excitatory and inhibitory current received from the thalamic neuron population. These current contributions are defined by:

$$x_{E4}(t) = \beta_{Er}[h_{Er} * r_4](t) \tag{3.9}$$

$$x_{I4}(t) = \beta_{Ir}[h_{Ir} * r_4](t) \tag{3.10}$$

$$f_T(t) = \beta_{Ef}[h_{Ef} * r_T](t) - \beta_{If}[h_{If} * r_T](t) \tag{3.11}$$

In these equations $h(t)$ denotes a temporal coupling kernel, and $[h * r](t)$ is a temporal convolution given by:

$$[h * r](t) = \int_0^\infty h(s)r(t-s)ds \tag{3.12}$$

The presence of this temporal convolution in the definition of the currents in eqs. (3.9-3.11) means that the current at time t results from firing at previous times. Thus, the temporal coupling kernel, $h(t)$ describes the effect on the present current by past firing activity. The temporal coupling kernels are modeled as delayed decaying exponentials:

$$h(t) = \frac{1}{\tau} e^{-(t-\Delta)/\tau} H(t-\Delta) \tag{3.13}$$

Here Δ is the time delay between firing activity in presynaptic neurons and the resulting current received in neuronal somas, τ is a time constant describing the timescale on which the firing of a population will contribute to the current and H() is the Heaviside unit step function.

Eqs. (3.9) and (3.10) are Volterra integral equations, and since $h(t)$ is an exponentially decaying function, these equations can be differentiated and formulated as two coupled differential equations using the "linear chain trick":

$$\begin{aligned}
\frac{dx_{E4}}{dt} &= \beta_{Er}\frac{d}{dt}\int_{-\infty}^{t} h_{Er}(t-s)F_4(I_4(s))ds \\
&= \frac{\beta_{Er}}{\tau_{Er}}F_4(I_4(t)) - \frac{\beta_{Er}}{\tau_{Er}}\int_{-\infty}^{t} h_{Er}(t-s)F_4(I_4(s))ds \\
&= \frac{\beta_{Er}}{\tau_{Er}}F_4(I_4(t)) - \frac{x_{E4}}{\tau_{Er}}
\end{aligned} \tag{3.14}$$

$$\begin{aligned}
\frac{dx_{I4}}{dt} &= \beta_{Ir}\frac{d}{dt}\int_{-\infty}^{t} h_{Ir}(t-s)F_4(I_4(s))ds \\
&= \frac{\beta_{Ir}}{\tau_{Ir}}F_4(I_4(t)) - \frac{\beta_{Ir}}{\tau_{Ir}}\int_{-\infty}^{t} h_{Ir}(t-s)F_4(I_4(s))ds \\
&= \frac{\beta_{Ir}}{\tau_{Ir}}F_4(I_4(t)) - \frac{x_{I4}}{\tau_{Ir}}
\end{aligned} \tag{3.15}$$

Chapter 4

Exploring the Thalamocortical Model

In this chapter, the thalamocortical model is explored with a triangular, a trapezoidal and a gaussian function to get a better understanding of how the model transforms different inputs.

4.1 Input functions for the thalamocortical model

The triangular shaped input function is chosen because it is simple and because it resembles the shape of the experimental thalamic stimulus response. In the experiment, the amplitude and the rise time of the applied stimulus were varied. This variation also shows up in the trial-averaged thalamic stimulus response; the amplitude, the duration and the steepness of the thalamic stimulus response varies over different stimulus conditions. Therefore the triangular input was chosen to vary according to four features: 1) amplitude, 2) duration, 3) steepness of rising slope and 4) steepness of falling slope. Three features of the model response, 1) amplitude, 2) area and 3) *sharpness* (ratio of amplitude to half width of signal) are compared relative to the same features for the triangular input to characterize how the model transforms these features of the triangular input.

The triangular function is constructed by combining two straight lines, one with a positive slope, a, and one with a negative slope, b:

$$f(t) = \begin{cases} at & t_s < t \leq t_{max} \\ C - bt & t_{max} < t \leq t_e \\ 0 & \text{otherwise} \end{cases} \tag{4.1}$$

The constant C is calculated to connect the two lines at the peak point. All the triangles are constructed by varying a, b and t_{max}. t_{max} corresponds to the time when $f(t)$ has its peak. t_s and t_e denotes triangular start and triangular end time. The triangular function is easily expanded into a trapezoidal function by adding a constant term (A) for a time interval ($t_{max,s}$ - $t_{max,e}$) after peak time:

$$f(t) = \begin{cases} at & t_s < t \leq t_{max} \\ A & t_{max,s} < t \leq t_{max,e} \\ C - bt & t_{max,e} < t \leq t_e \\ 0 & \text{otherwise} \end{cases} \tag{4.2}$$

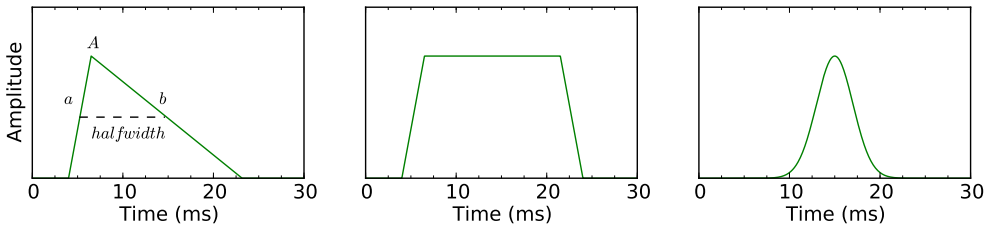

Figure 4.1: Input functions. Left: The triangular input is described by a rising slope a, a falling slope b and an amplitude A. The half width of the triangle is used for comparison with the half width of the model response. Center: Trapezoidal input is made by adding a straight horizontal line between the rising and falling slopes, and the length of this line can be adjusted to make an input signal of different durations. Right: Gaussian bell shape described by eq. (4.3).

Because of the kink at its, maximum the triangular function is not fully realistic. Each sensory stimulus is filtered during sensory processing before it arrives at the thalamocortical interface. Thus, the input from thalamus is smoother than the abstract triangular functional form. To see how the model responds to input with a rounded peak, a gaussian function is used:

$$f(t) = A \exp^{-((t-15)/w)^n} \tag{4.3}$$

Here A denotes the amplitude and w and n are used to set the width and the shape of the signal.

4.2 Varying Amplitude and Duration of the Triangular Input

To see how the size of the triangular input affects the model response, first the amplitude and the duration of the triangular input is varied, while rise and decay times are kept equal. Two different scenarios are presented: i) the duration is kept constant while the amplitude is varied and ii) the amplitude is kept constant while the duration is varied. The results are shown in Figures 4.2 and 4.3.

Figure 4.2A shows the model responses to triangular inputs with different amplitudes and equal durations. Figure 4.2B illustrates how the amplitude, the area and the sharpness of the model response all get fractionally smaller relative to the same features of triangular input when increasing the amplitude of the triangular input from 0.1 to 0.4. However, when the amplitude of the triangular input exceeds 0.4, this relationship is reversed. In the model responses

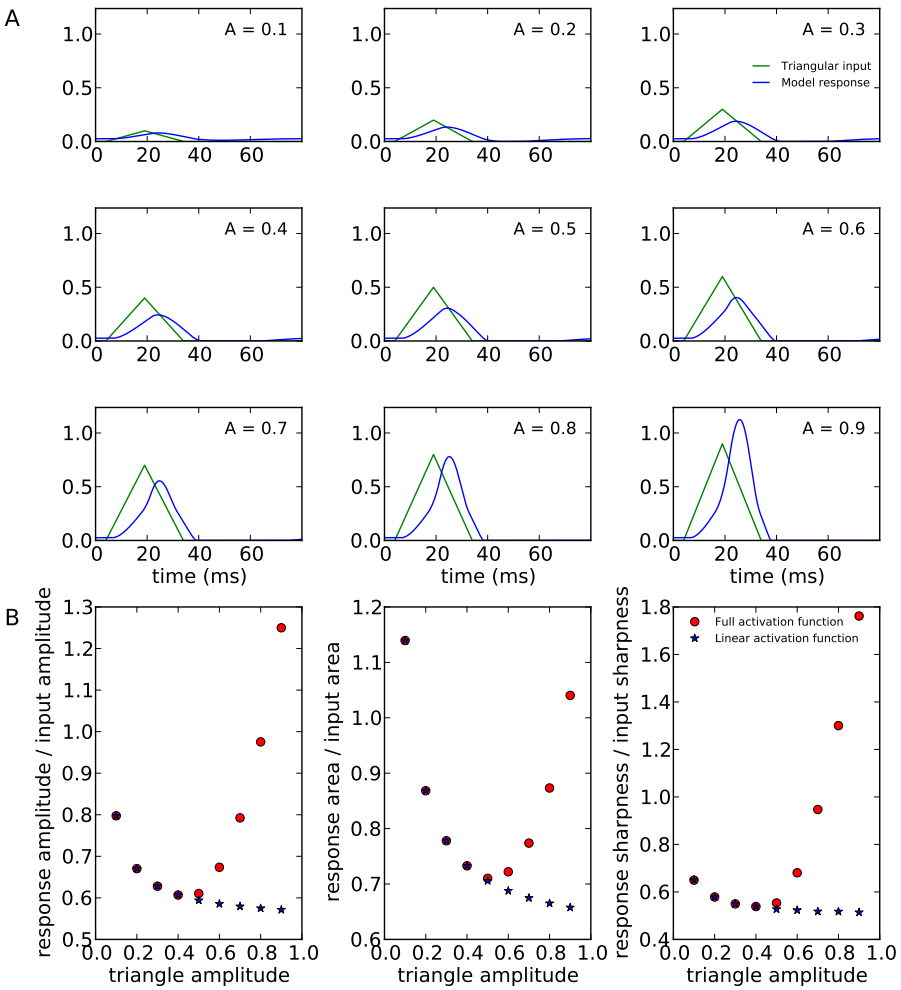

Figure 4.2: Varying the amplitude of triangular input. A) The triangular input (green line) and the model response (blue line) for 9 different amplitudes of the triangular input. The amplitude is varied from 0.1 to 0.9 in steps of 0.1 from top left plot to bottom right plot. Amplitudes are shown in the top right corner of each plot. B) Scatter plots showing the ratio between different features (amplitude (left), area (middle) and sharpness (right)) of the model response and the triangular input as a function of the amplitude of the triangular input. Red dots show the results for the complete thalamocortical model and blue stars show the result when the quadratic term in the activation function (eq. (3.7)) is deactivated.

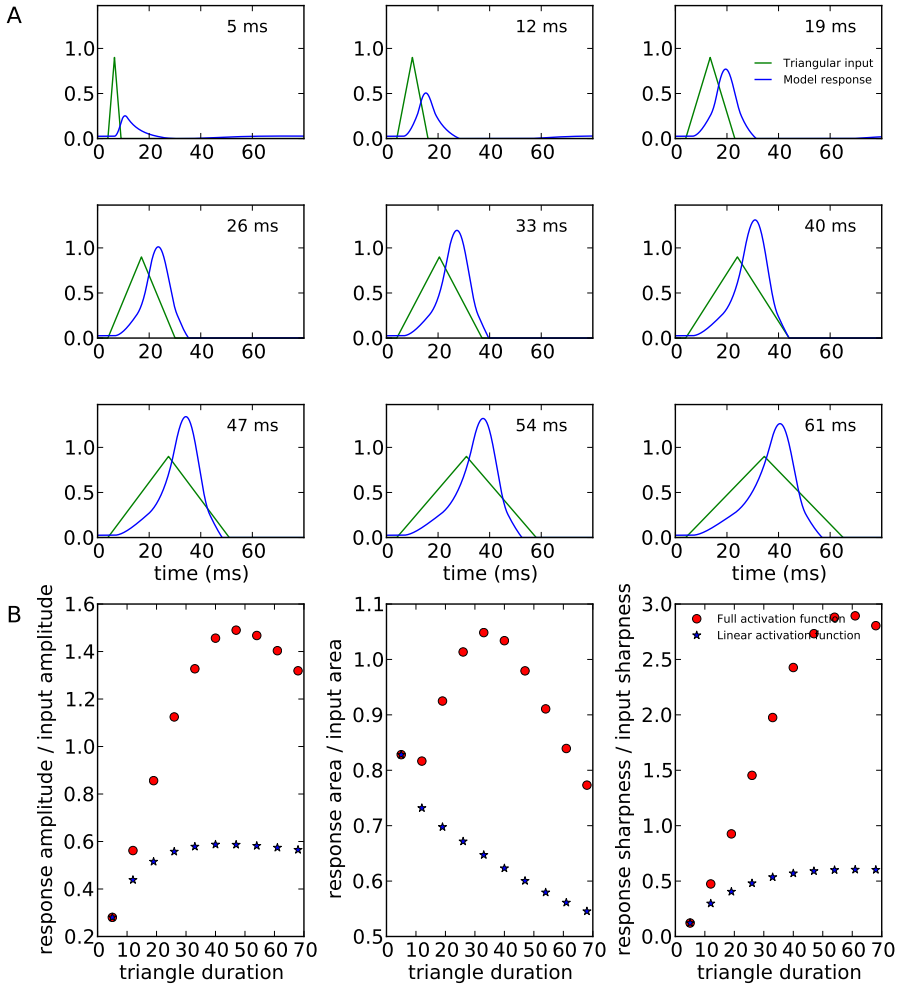

Figure 4.3: Varying duration of triangular input. A) The triangular input (green line) and the model response (blue line) for 9 different durations of the triangular input. The duration is varied from 5 to 61 ms in steps of 7ms from top left plot to bottom right plot. B) Scatter plots showing the ratio between different features (amplitude (left), area (middle) and sharpness (right)) of the model response and the triangular input as a function of the duration of the triangular input. Red dots show the results for the complete thalamocortical model and blue stars show the result when the quadratic term in the activation function (eq. (3.7)) is deactivated.

for $A = 0.6$ to $A = 0.9$ there is a location on the graph where the rise turns visibly steeper, suggesting that the quadratic term from eq. (3.7) is activated. Indeed, turning this quadratic term off before running the simulation results in exponentially decaying relationships in all the scatter plots of Figure 4.2B.

While initially the output amplitude and area decrease non-linearly with the input amplitude and area, above a certain threshold this relationship is reversed due to activation of the quadratic term in eq. (3.7). In terms of firing activity, this would mean that the firing activity grows non-linearly with input for strong input. For the last trial in Figure 4.2A ($A = 0.9$) the amplitude and area is even amplified relative to the triangular input. The rightmost plot in Figure 4.2B also shows that the model response is sharpened for the highest input amplitudes ($A = 0.8$, $A = 0.9$).

In Figure 4.3 the duration of the triangle is varied at a constant amplitude of $A = 0.9$. Here we see an opposite effect; initially the output amplitude grows with increasing duration, but when the duration exceeds 47 ms the output amplitude starts decreasing for increasing durations. The same effect is seen for the area and sharpness of the output, but the turning points come at 45 ms and 55 ms, respectively. As will be shown in a later section, this dampening effect at longer durations is due to increasing influence of the inhibitory neuron population. Another interesting thing we observe here is that the amplitude of the response is amplified for durations above 25 ms, but not below. Also, the area is amplified for durations between 25-45 ms. It is obvious that a combination of high amplitude and duration is needed to give a strong response. However, if the duration is too long, both amplitude and area decrease relative to input for increasing durations.

4.3 Varying Slopes of Triangular and Trapezoidal Input

Figure 4.4 shows how the model response changes with varying slopes. The rising slope a and the falling slope b are adjusted so that the amplitude and the duration of the triangle is kept constant. The peak thus moves from the beginning to the end of the signal during the nine trials. The amplitude of the response is seen to be highest when the peak is at a relative position of around 0.4. The same effect is also seen for triangles of different areas (longer durations), but the peak position giving the highest response amplitude moves. When the area increases, the peak position giving the highest response amplitude moves towards the beginning of the signal, in other words toward a steeper upward slope and a shallower decreasing slope.

It therefore seems that there is an "optimal" ratio between the slopes of the rising and falling flank of the signal, but this ratio changes with the intensity of the signal. For strong signals, the rising slope is relatively steeper, and for weak signals the rising slope is relatively shallower.

To see how the rising slope alone influences the model response, the triangle in eq. (4.2) is expanded into a trapezoid. In Figure 4.5, the duration of a trapezoidal

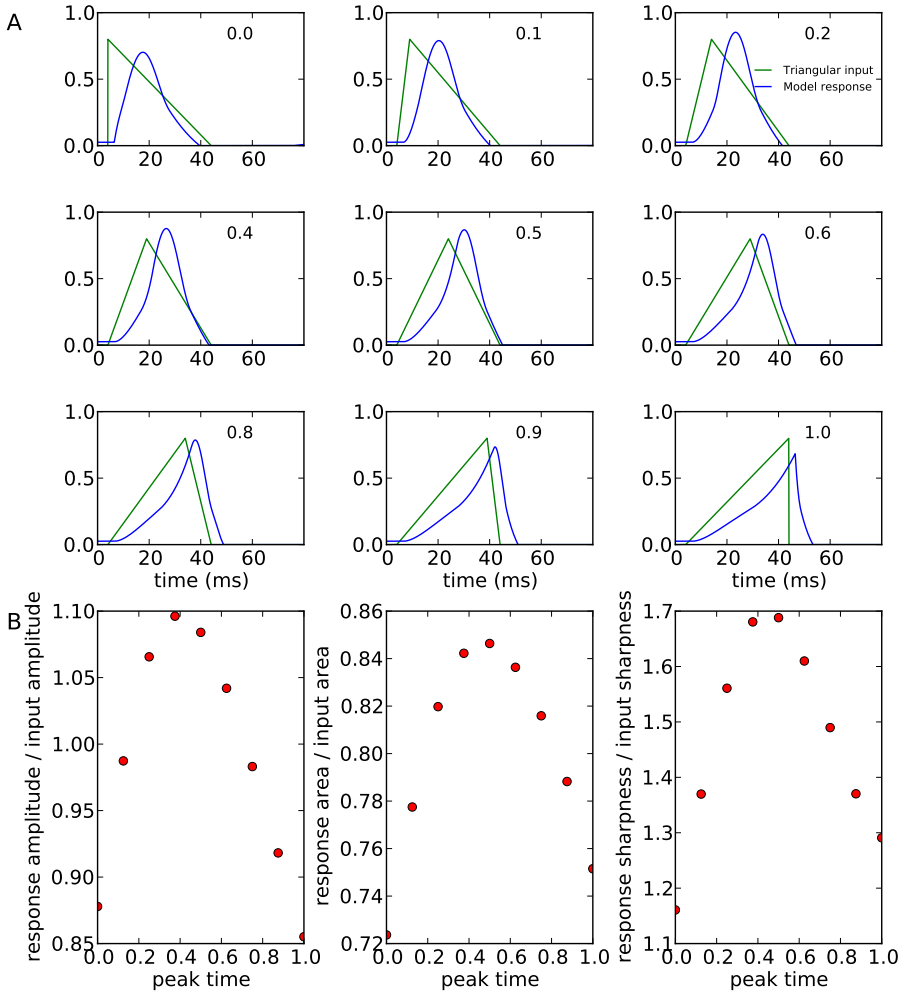

Figure 4.4: Varying peak position of triangular input. A) The triangular input (green line) and the model response (blue line) for 9 different amplitudes of the triangular input. The peak position is varied from start to end of the signal in steps of 0.125 from top left plot to bottom right plot. Relative peak positions are shown in the top right corner of each plot. B) Scatter plots showing the ratio between different features (amplitude (left), area (middle) and sharpness (right)) of the model response and the triangular input as a function of the peak position of the triangular input.

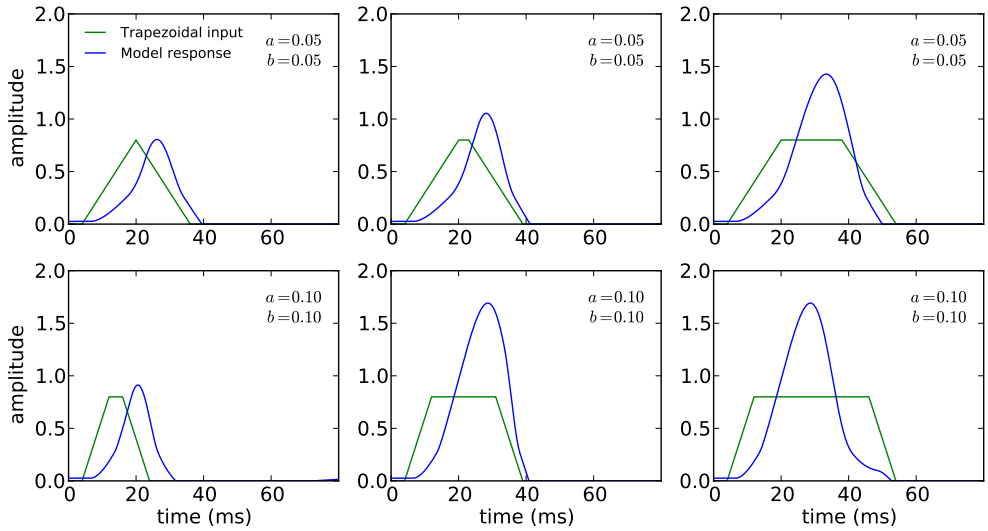

Figure 4.5: Different slopes in trapezoidal input. Trapezoidal inputs with slope $a = 0.05$ (top) and $a = 0.1$ (bottom) and model responses are plotted from left to right. In the top left plot the trapezoid is reduced to a triangle.

input of two different slopes, $a = 0.05$ (top) and $a = 0.10$ (bottom), is varied. As we saw in Figure 4.3 increasing the duration over a certain point does not give a larger response. Whereas increasing the duration for the triangle further led to a decline in amplitude, here the peak amplitude remains the same for increasing duration (see bottom center and bottom right). We also see that the steepness of the rising slope determines how large the amplitude of the response gets, and how fast it peaks.

4.4 Gaussian Input

Figure 4.6 shows four different gaussian inputs and illustrates that a broader peak gives a larger response. It is worth noting that this effect is smaller for low amplitude inputs (left) than for high amplitude inputs (right). This reinforces the conclusion from section 4.2; a combination of high amplitude and long duration gives a larger response due to more contributions from the quadratic term of the activation function.

4.5 Looking at Model Terms

In Figure 4.7 the individual current terms of eqs. (3.9)-(3.11) are plotted to see how the individual components of the model response are behaving for different inputs. In all plots, f_T follows the input, while the difference $x_e - x_i$ first grows together with f_T, but after peaking decreases and turns negative. The early dominance of x_e is overtaken by a dominating x_i. In Figure 4.7A we see that the largest difference between x_e and x_i increases for increasing input amplitudes.

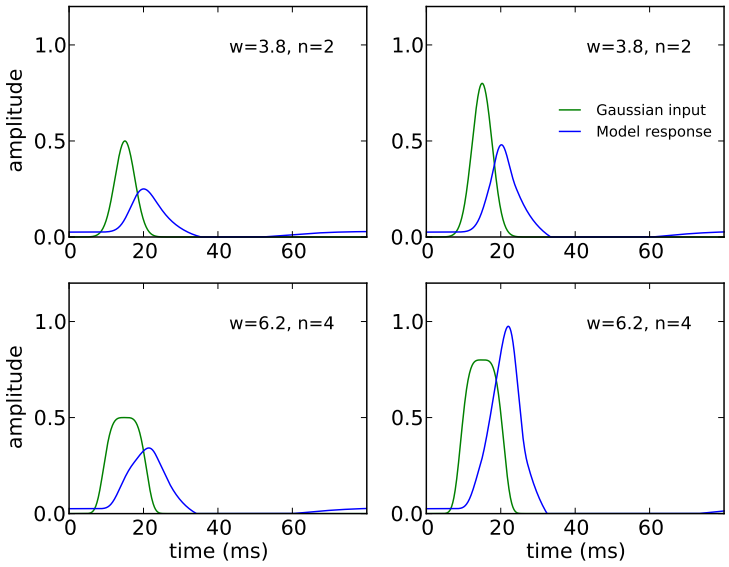

Figure 4.6: Gaussian model input. Model responses (blue line) for four different gaussian inputs (green line). The amplitudes (A) are 0.5 in the two left plots and 0.8 in the two right plots and the values of w and n are shown in each plot.

In Figure 4.7B, the largest difference $x_e - x_i$ first increases, but for the longest duration the largest difference is smaller again. In Figure 4.7C the difference $x_e - x_i$ is very similar, but fine-reading shows that the largest difference increases with a steeper slope.

4.6 Chapter Summary

The model responses depend strongly on the amplitude and duration of the input. A weak response follows an input of either low amplitude or short duration or both, while a strong response follows an input of high amplitude lasting for a sufficiently long time to activate the quadratic term in the activation function. The activation and contribution of this term to the response is governed by the thalamic current term and by the difference between the excitatory and inhibitory current terms. Both these current terms grow with amplitude and duration of the signal.

The position of the peak can also influence the amplitude and size(area) of the model response. For short lasting, low amplitude inputs a shallow rising onset and a sharp offset give the strongest response, while the opposite is true for a longer lasting, higher amplitude input. A steep onset gives a stronger early excitatory dominance, while a quick offset quickly reduces the thalamic contribution, and the balance between these two effects determines the response.

Lastly, the amplitude of the model response is seen to depend strongly on

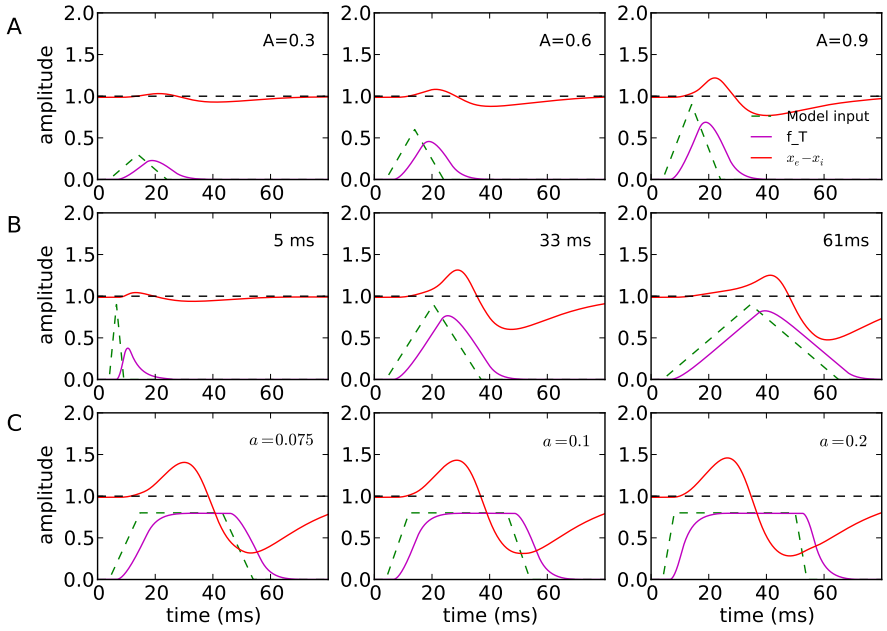

Figure 4.7: Individual current terms for varying input. Thalamic current term (f_T) and the difference between the excitatory and inhibitory current terms ($x_e - x_i$) from eqs. (3.9)-(3.11) are plotted for different model inputs. The line for $x_e - x_i$ is shifted upwards to include the negative part, and the black dotted line marks the zero level. A) Triangular input, amplitude increases from left to right. B) Triangular input, duration increase from left to right. C) Trapezoidal input, slope increases from left to right.

the sharpness of the input when the input has a high amplitude. Only a slight increase in the width of the input gives larger contributions from the quadratic term and therefore a stronger response.

Chapter 5

Explore Single Trial Data

The models from chapter 3 are made from trial-averaged data. This chapter explores single-trial data and the trial-to-trial variability present in the MUA in both thalamus and cortex. Then, the thalamocortical model is tested with thalamic input from single trials to see how well it predicts the single-trial layer 4 population firing rates.

5.1 Trial-to-Trial Variability

The trial-to-trial variability in the MUA is substantial. The examples in Figure 5.1 show some of the typical cases in the cortical response, e.g. no stimulus response, a distinct stimulus response and spontaneous activity prior to stimulus response. The thalamic activity is seen to have both different temporal patterns and variable amplitudes.

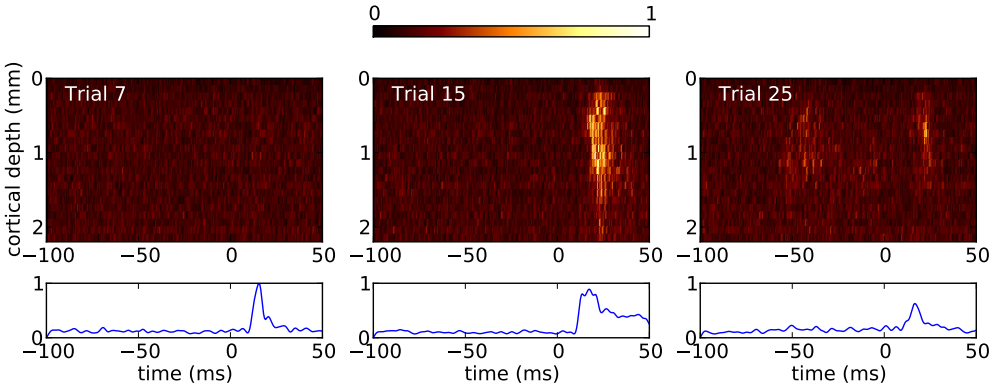

Figure 5.1: Example of trial-to-trial variability in MUA. Experimental data for three trials from stimulus condition 1 is presented from left to right. Color plots show the rectified cortical MUA for a period of 150 ms, spanning the entire cortical depth. The line plots show the MUA in thalamus for the corresponding trial and time duration. Data is normalized to the largest value of the three trials.

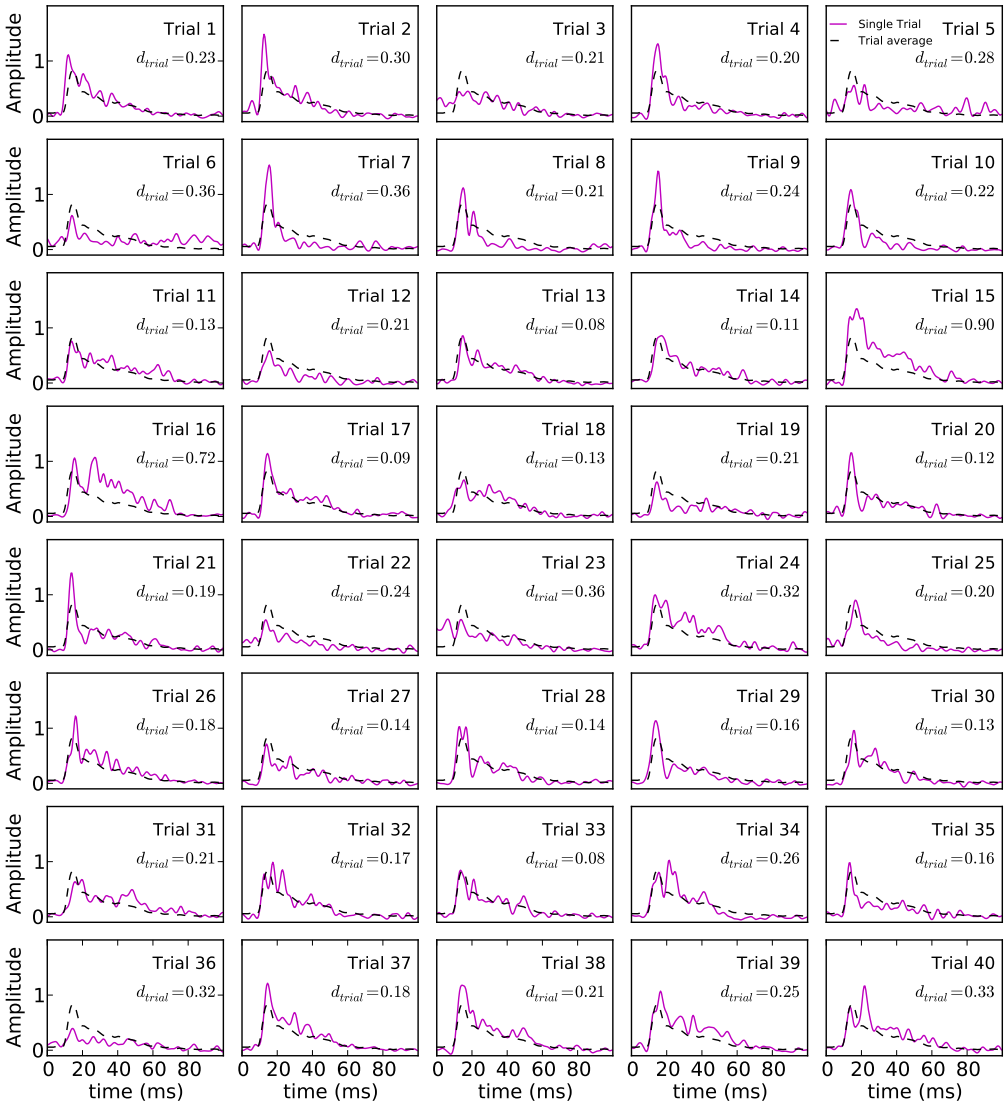

Figure 5.2: Thalamic MUA trialset. Thalamic MUA (purple line) for the 40 trials of stimulus condition 1. The trial-averaged MUA (black dotted line) is also added to each plot for comparison. MUA is rectified, summed over all channels, low-pass filtered and normalized before plotting. Trial deviations d_{trial} are computed from eq. (5.1)

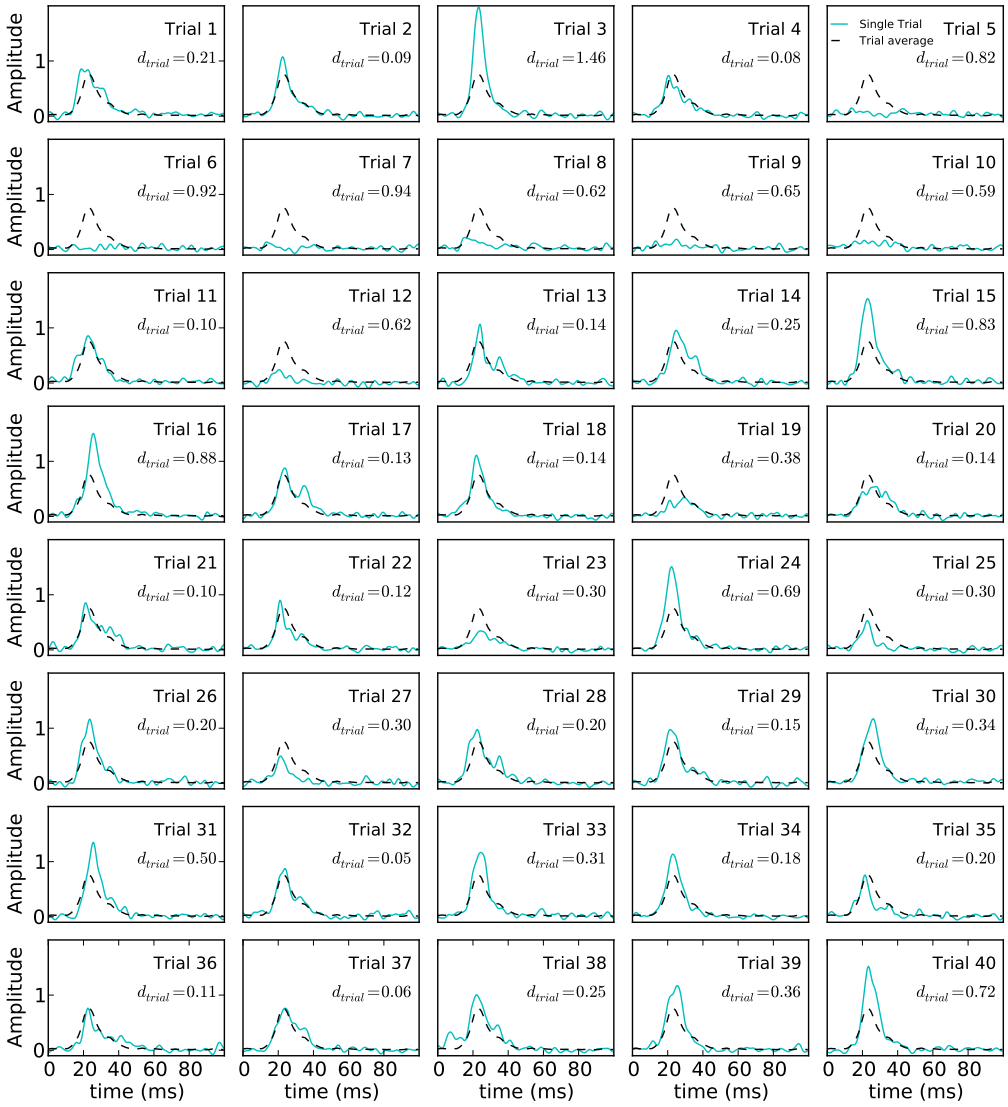

Figure 5.3: Cortical MUA trialset. Cortical MUA (blue line) for the 40 trials of stimulus condition 1. The trial-averaged MUA (black dotted line) is also added to each plot for comparison. MUA is rectified, summed over all channels, low-pass filtered and normalized before plotting. Trial deviations d_{trial} are computed from eq. (5.1)

To get a better picture of the trial-to-trial variability, the MUA in thalamus and cortex for all trials in stimulus condition 1 is plotted on the preceding pages, in Figures 5.2 and 5.3, respectively.

The thalamic trial response follows the shape of the trial-averaged response in most trials, having a steep onset and longer lasting oscillating offset. The amplitude however is seen to vary from a factor two below the trial average (e.g. trial 36) to a factor two above the trial average (e.g. trial 7). The same is true for the cortical trial response: with few exceptions (e.g. trial 13, 14 and 17) the stimulus response follows the bell-shape of the trial-averaged response and the amplitude varies from a factor two below the trial average (e.g. trial 19) to a factor two above the trial average (e.g. trial 3). A key observation is the lack of visible stimulus responses for trials 5–10, even though the thalamic population is activated in all these trials.

To illustrate further, in Figures 5.4 and 5.5, the thalamic and cortical response peak of each trial is compared to the trial-averaged response peak for that stimulus condition. This is done for all trials and stimulus conditions. Here, a green dot corresponds to a higher than average peak value while a pink dot corresponds to a lower than average peak value. The first row (stimulus condition 1) in these figures re-illustrates many of the observations from the last paragraph, e.g., in Figure 5.4, trial 7 is dark green while trial 36 is pink. This corresponds to a high and low amplitude in these thalamic responses. Figure 5.5 shows the cortical responses, here trial 3 is dark green while trials 5–10 are dark pink. Two overall trends emerge: i) there is a band of trials (pink) in the beginning of the experiment with repeated low responses, as seen in Figure 5.5, and ii) relatively high responses (dark green) are more frequent among weaker stimulus conditions, as seen in Figures 5.4 and 5.5 (see stimulus conditions (5–9, 15–18, 25–27)).

The band of pink trials (∼5–12) in Figure 5.5 hints that the cortex has been in an silent state for some period during the experiment. For all stimulus conditions, many of these low response trials succeed each other. However, since stimuli were delivered in a random pattern, they are not located directly on top of each other, instead they are shifted sideways from row to row. This simply means that some stimulus conditions occurred more times than others before this period of cortical silence. Also, both in the thalamic and cortical activities there seems to be a relative high occurrence of strong stimulus responses prior to the low response trials. This indicates that the brain circuits respond well to stimuli in the beginning, but then the cortical circuits are suddenly disconnected or strongly inhibited.

Both in the thalamic and cortical responses, relatively high responses are more frequent for weaker stimulus conditions. Also, the trials falling in the no-response category (e.g. 10–14, stimulus condition 8 (a1,t8)) have light colors, suggesting that the trial-averaged peak is close to the peak in a no-response trial. Most likely the trial-averaged value is underestimated due to higher trial-to-trial variability in the weak stimulus conditions. This is investigated further in section 5.3.

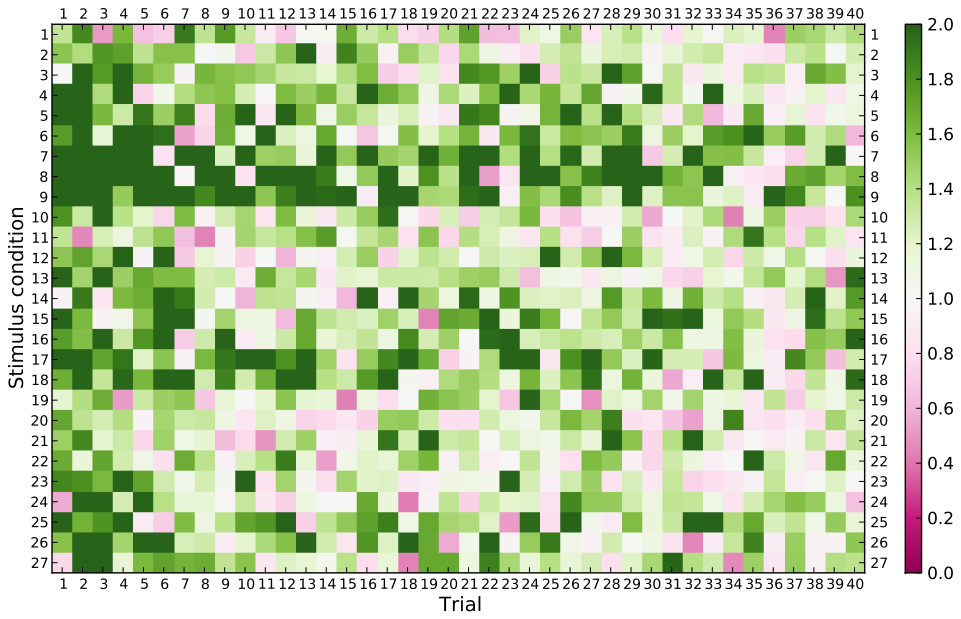

Figure 5.4: Thalamic response peak summary. Image showing the ratio between the peak amplitude of each trial and the peak amplitude of the trial average for that stimulus condition, for all trials and stimulus conditions. Peak values are found from MUA which are prepared in the same way as in Figure 5.2.

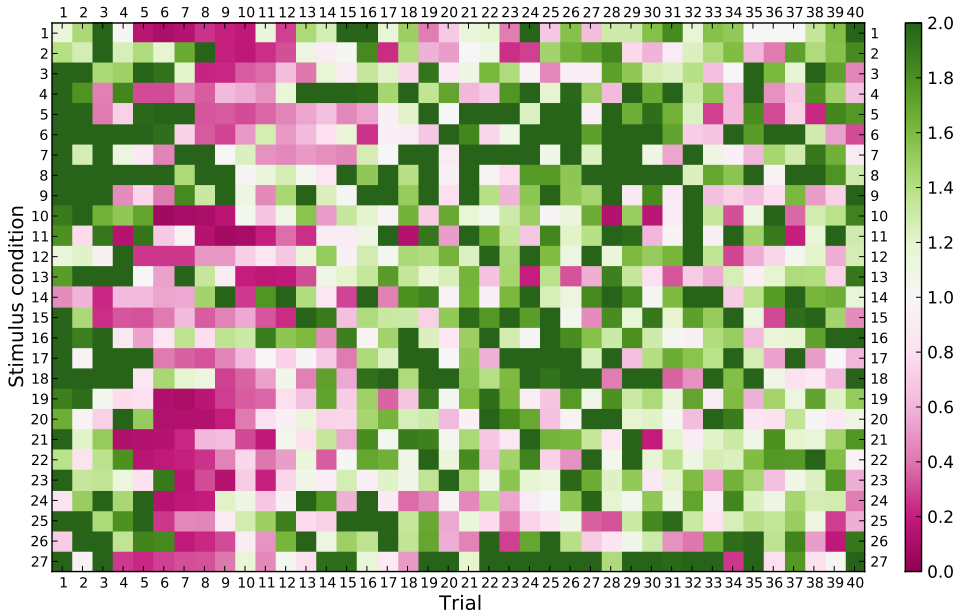

Figure 5.5: Cortical response peak summary. Image showing the ratio between the peak amplitude of each trial and the peak amplitude of the trial average for that stimulus condition, for all trials and stimulus conditions. Peak values are found from MUA which are prepared in the same way as in Figure 5.3.

5.2 Spontaneous Activity Revisited

Spontaneous activity, as illustrated in Figure 5.1, is present in many of the trials, and when it occurs before stimulus the stimulus response is often weak. Figure 5.6 shows how the stimulus response is weaker when spontaneous activity is present in the time window of 100 ms before stimulus for stimulus condition 1 (a1,t1). Trials 12, 25 and 27 all have strong spontaneous activity and a weak stimulus response compared to the other trials. An exception is trial 3, with a high stimulus response despite some spontaneous activity occurring before stimulus. These patterns repeat themselves throughout all stimulus conditions (not shown); more often than not, spontaneous activity dampens the stimulus response.

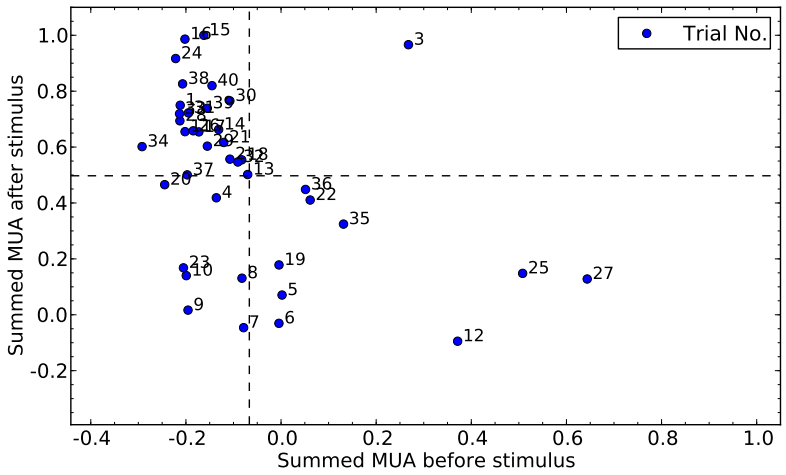

Figure 5.6: Stimulus response vs spontaneous activity. For each trial in stimulus condition 1 (a1,t1), MUA activity in the 100 ms before stimulus is summed and plotted against the sum of the MUA activity in the 100 ms after stimulus. The sums are normalized according to the largest value along each axis and dotted lines show the average sum along each axis. Data is rectified and baseline subtracted and sums are over all channels.

Two factors are likely to best explain this negative relationship, as seen in Figure 5.7: the magnitude of the spontaneous activity and the time interval between spontaneous activity and the stimulus. To investigate this relationship, the stimulus evoked activity is plotted against these two variables. Due to relatively few trials of spontaneous activity per stimulus condition, all trials containing spontaneous activity are pooled together. The results are shown in Figure 5.8. In the left image, high and low stimulus responses are distributed evenly over all time intervals. However, in the right image high stimulus responses occur most often at low spontaneous activities and vice versa. From these plots, it therefore seems that the stimulus response is independent of the time interval between spontaneous activity and stimulus, but decreases with increasing magnitude of spontaneous activity.

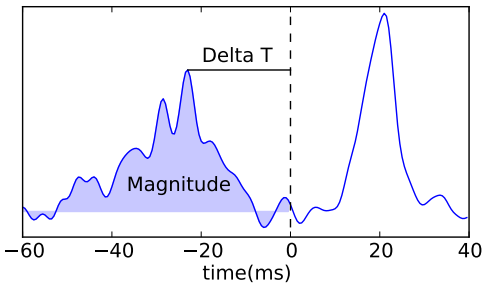

Figure 5.7: Explaining parameters of spontaneous activity. Image shows the cortical MUA for one trial with spontaneous activity preceding stimulus response. Delta T is the time interval between the peak of spontaneous activity and the stimulus onset (dotted line). The magnitude of the spontaneous activity refers to the area under the curve.

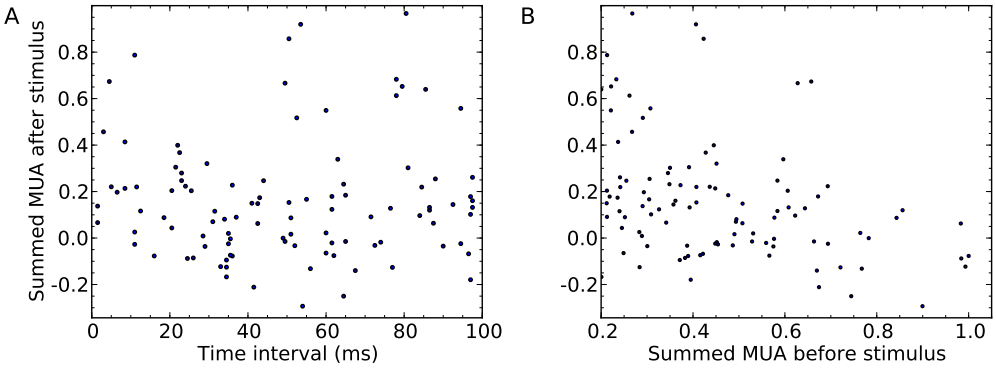

Figure 5.8: Influencing Factor of Spontaneous Activity. A) All trials containing spontaneous activity above a selected threshold are pooled together. Magnitude of stimulus response is plotted against the time interval between peak spontaneous activity and stimulus onset. B) Magnitude of stimulus response plotted against magnitude of spontaneous activity.

5.3 Attempt at Quantifying Variability

The last two sections suggest that the trial-averaged stimulus response is underestimated. A silent state in cortex prevents a stimulus response in many trials and spontaneous activity prior to stimulus dampens the stimulus response. In this section, the trial-averaged stimulus response and the trial-to-trial coefficient of variation (standard deviation divided by mean) are compared among trial sets including all trials and trial sets excluding the aforementioned.

Trials to be removed were selected by visual inspection from low-pass filtered MUA. For the trials during the silent state, consecutive trials of the same stimulus condition in the beginning of the experiment (from approx. trials 1-16) with

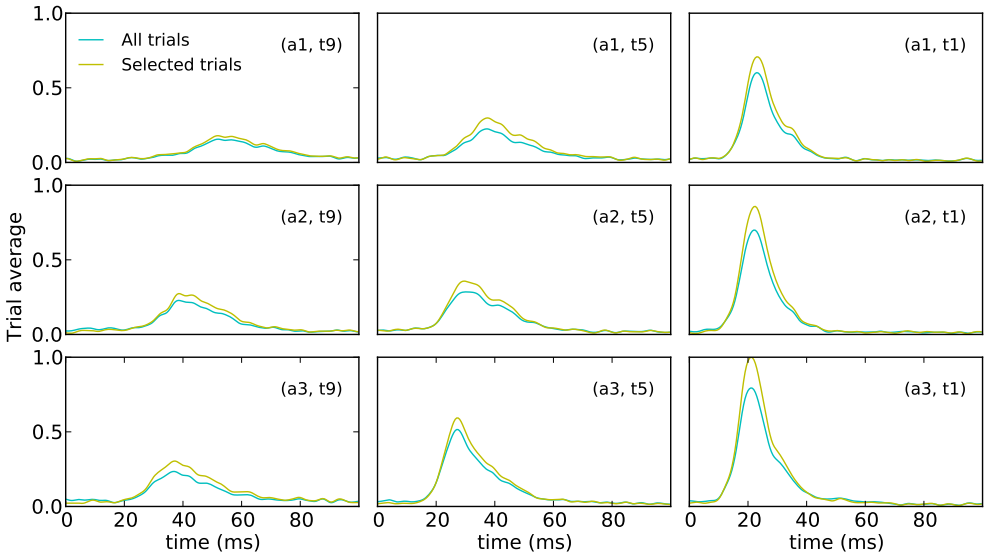

Figure 5.9: Trial average for two trial sets. Figure shows trial averaged data for two different trial sets for 9 out of the 27 stimulus conditions. Trial averages for all trials are shown in blue and trial averages for which trials during the cortical silent state and trials containing spontaneous activity are removed are shown in yellow.

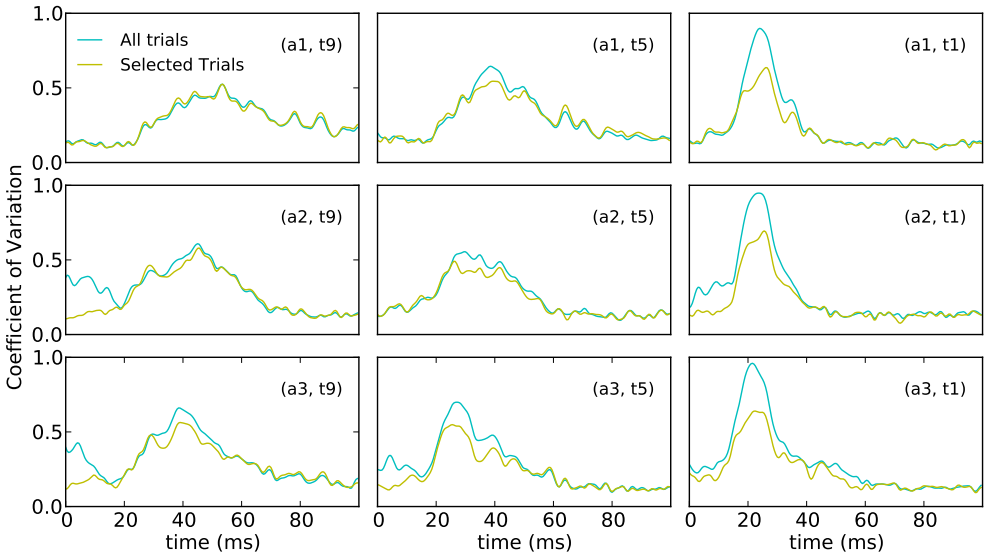

Figure 5.10: Coefficient of variation for two trial sets. CV for two different trial sets for 9 out of the 27 stimulus conditions. Trial average for all trials are shown in blue and trial averages for which trials during the cortical silent state and trials containing spontaneous activity are removed are shown in yellow.

no or strongly inhibited stimulus response were selected. For trials with spontaneous activity, all trials with a magnitude of spontaneous activity above 0.3 were selected. For some trials, spontaneous activity was present during stimulus onset, and these were selected on eyesight.

Figures 5.9 and 5.10 show that trial average goes up and coefficient of variation goes down when some of the trials are removed, and the effect is seen to be strongest for stimulus conditions of fast rise times. A larger mean value is expected since most of the removed trials have low stimulus responses. The coefficient of variance when some trials are removed is seen to have about the same amplitude for all the stimulus conditions shown, but it is more confined at strong stimuli. Therefore, after some trials are removed, the coefficient of variance is largest for weak stimulus conditions.

5.4 Population Trial-to-Trial Variability

The variability in the extracted population firing rates are similar to the variability in the channel summed data. This is illustrated for all populations in Figure 5.11 and for the layer 4 population in Figure 5.12. Both high amplitude responses (e.g. 3, 15, 16, 24 and 40) and low amplitude responses (e.g. 5–10, 12, 19, and 23) are shared in both the layer 4 population and in the channel summed MUA. Trials with high deviations from mean in the total MUA typically also have high deviations from mean in each population as well and vice versa. However, there are also many trials where the deviations in one or two populations are more than twice as big as the deviations in other populations, e.g. 2, 3 ,4, 15, 16, etc.

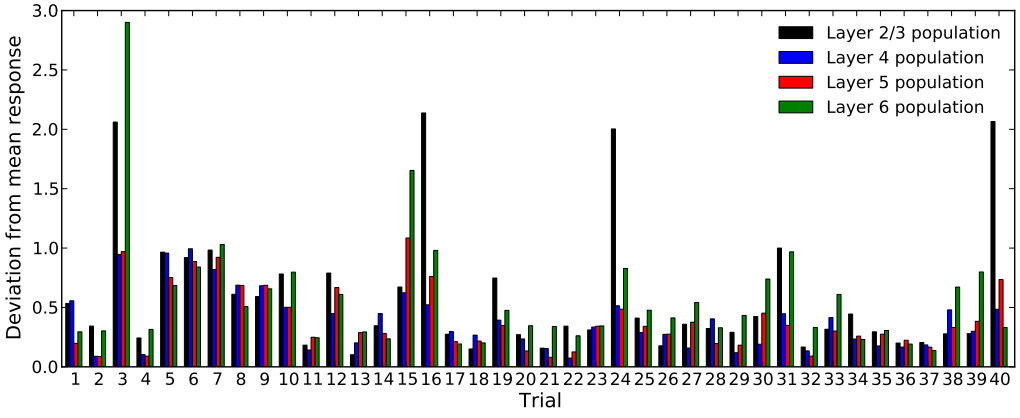

Figure 5.11: Deviation from mean for all populations. This figure show the deviation from mean response of the identified cortical populations for all 40 trials of stimulus conditions 1.

To see whether there is a correlation between thalamic and layer 4 population firing rates, peak amplitudes, total firing and deviation from mean was considered. Figure 5.13 shows that the correlation is best between total firing in thalamus and cortex in stimulus condition 1. There are small differences in

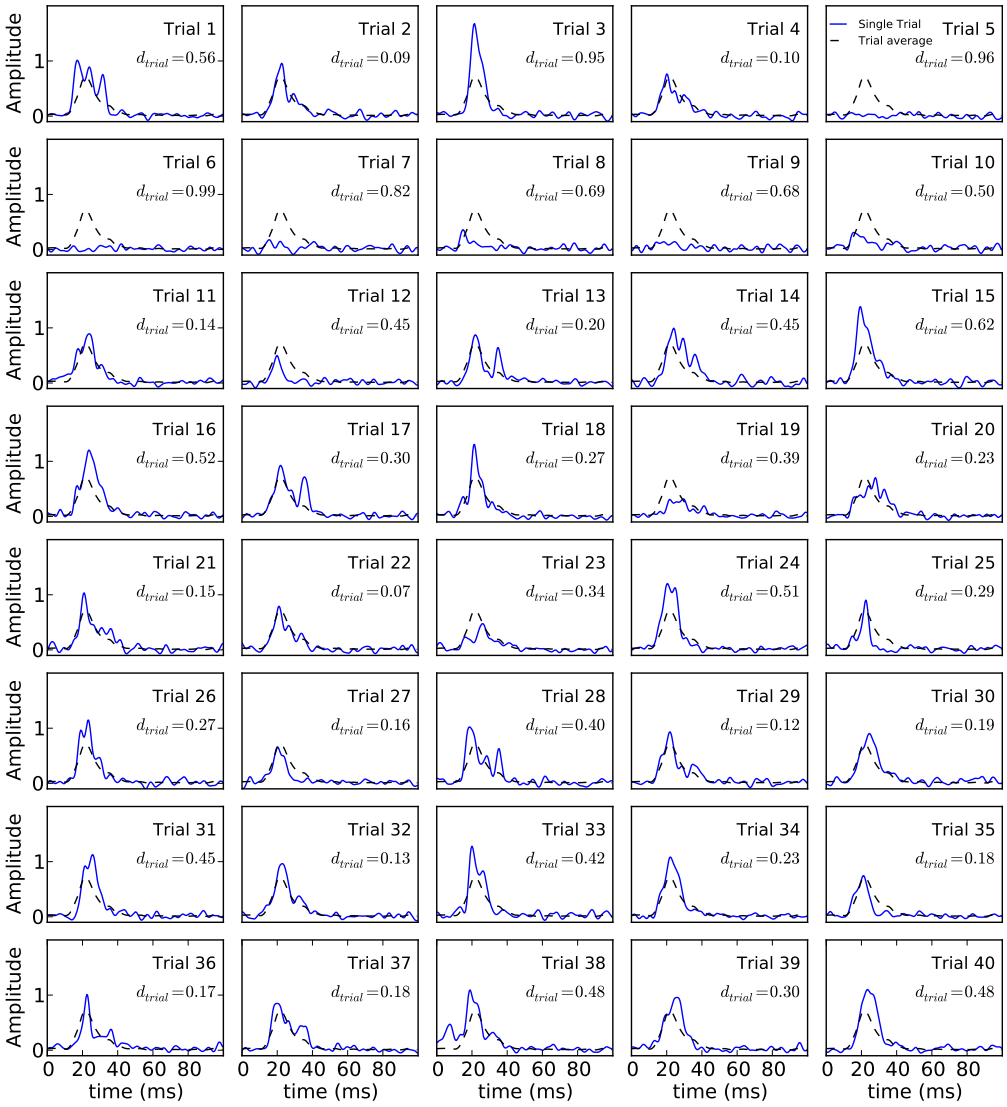

Figure 5.12: Layer 4 population firing-rate trialset. Layer 4 population firing rates (blue line) for the 40 trials of stimulus condition 1 (a1, t1). The trial-averaged firing rates (black dotted line) are also added to each plot for comparison. Firing rates are low-pass filtered before plotting. Trial deviations d_{trial} are computed from eq. (5.1)

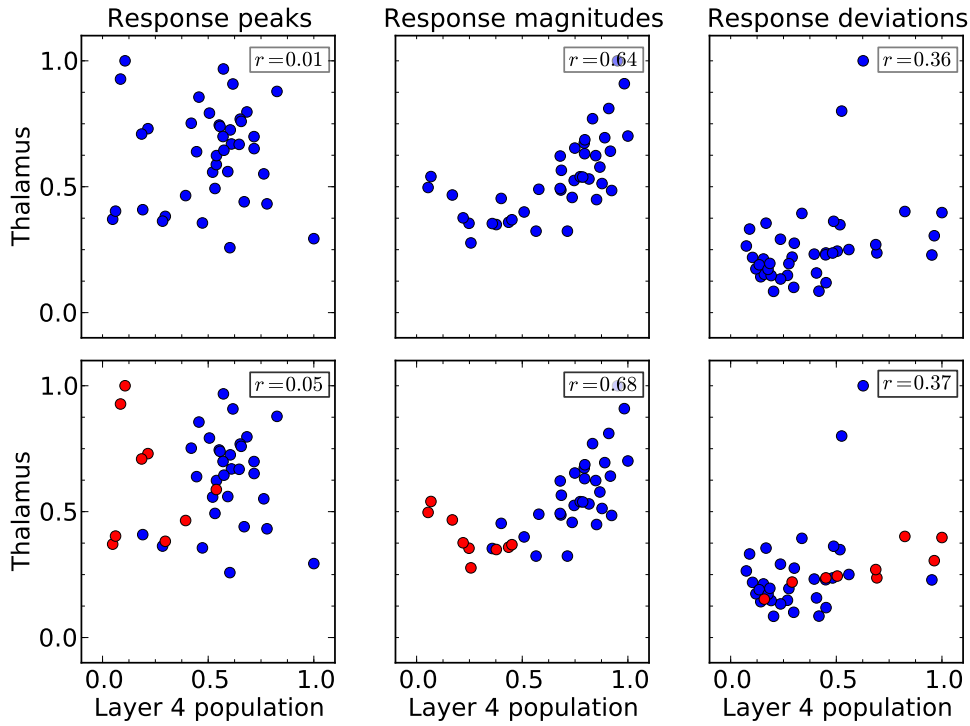

Figure 5.13: Correlation between thalamic and layer 4 population firing rates. Correlation between three features of the thalamic and layer 4 population firing rates: Peak amplitude of response (left), magnitude of response (middle) and deviation from mean (right). Each dot correspond to one trial of stimulus condition 1. In the top plots, correlation is computed from all the trials, while in the bottom plots correlation is computed from selected trials. Red dots correspond to trials not considered in the bottom row.

this correlation when looking at all trials compared to looking at trials in which trials during the period of cortical silence and trials containing spontaneous activity are removed. Correlations of total firing between thalamus and cortex are similar across all stimulus conditions, as seen in table 5.1.

Table 5.1: Correlation between magnitude of thalamic and layer 4 population firing rates. Correlation values are computed for all trials, as demonstrated in the top center plot of Figure 5.13.

Amplitude\Rise time	t1	t2	t3	t4	t5	t6	t7	t8	t9
a1	0.64	0.85	0.56	0.68	0.71	0.59	0.53	0.64	0.51
a2	0.81	0.71	0.79	0.59	0.71	0.74	0.32	0.57	0.52
a3	0.77	0.75	0.77	0.66	0.68	0.67	0.48	0.67	0.55

The trial variation in this section is computed with a relative deviation from

mean response:

$$d_{trial} = \sum_{j=1}^{N_t} \left(\mathbf{\Phi}_{MUA}^{trial-n}(t_j) - \mathbf{\Phi}_{MUA}^{trial-avg}(t_j) \right)^2 \Big/ \sum_{j=1}^{N_t} \left(\mathbf{\Phi}_{MUA}^{trial-avg}(t_j) \right)^2 \quad (5.1)$$

This deviation measures the sum square deviation from the trial-averaged mean response for each time step t_j of a trial, and is scaled by the sum square of the trial-averaged response. The sum is taken over the number of time steps N_t, corresponding to the 100 ms after stimulus.

5.5 Modeling Single Trial Data

As shown in Figure 5.14, the thalamocortical model has some clear limitations when it comes to estimating single-trial responses. While model errors are quite small in many of the trials, they are also large in others. For comparison with thalamic firing rates, Figure 5.14 is plotted together with thalamic firing rates in appendix A. The normalization of the thalamic firing rates and the layer 4 population firing rates are based upon trial averages, to obtain the same scale as that which the model was fitted to. Model errors were computed as square deviations from the mean response:

$$e_m = \sum_{j=1}^{N_t} \left(r_4^x(t_j) - r_4^{mod}(t_j) \right)^2 \Big/ \frac{1}{N_{stim}} \sum_{k=1}^{N_{stim}} \sum_{j=1}^{N_t} \left(r_4^x(s_k, t_j) - \langle r_4^x(s_k) \rangle \right)^2 \quad (5.2)$$

Here, the sum in the denominater also goes over the number of stimulus conditions, N_{stim}, and s_k denotes stimulus condition k. $\langle r_4^x(s_k) \rangle$ denotes the temporal mean of $r_4^x(s_k, t_j)$ [4]. Note that the scaling factor is constant across all stimulus conditions, allowing for comparison of trials of different stimulus conditions.

The most obvious shortcoming of the model when using this normalization scheme is the blow-up in trial 15. The thalamic firing rate in this trial has a higher than average peak amplitude combined with a wide peak. As seen in chapter 4 the model is highly sensitive to high amplitude inputs. More conservative over-estimations are seen in in trials 2, 4, 16 and 38. Because of large trial-to-trial variability, in some trials the thalamic firing rates will have an amplitude higher than 1 when normalized according to trial averages. However, normalizing according to any other scheme seems unsuitable. For example, normalizing to the highest single-trial response would leave most of the trials with very low activity.

Also on lower-than-average thalamic inputs the model predictions can fall short (e.g. trials 22 and 36). These examples demonstrate that weak thalamic firing rates can give strong responses in cortex, and could indicate that other features of the thalamic input than the total firing is important to activate the cortical population. This can also be inferred by comparing trials 13 and 17. Even though experimental firing rates are very similar, the predicted responses based on the thalamic firing rate are somewhat different.

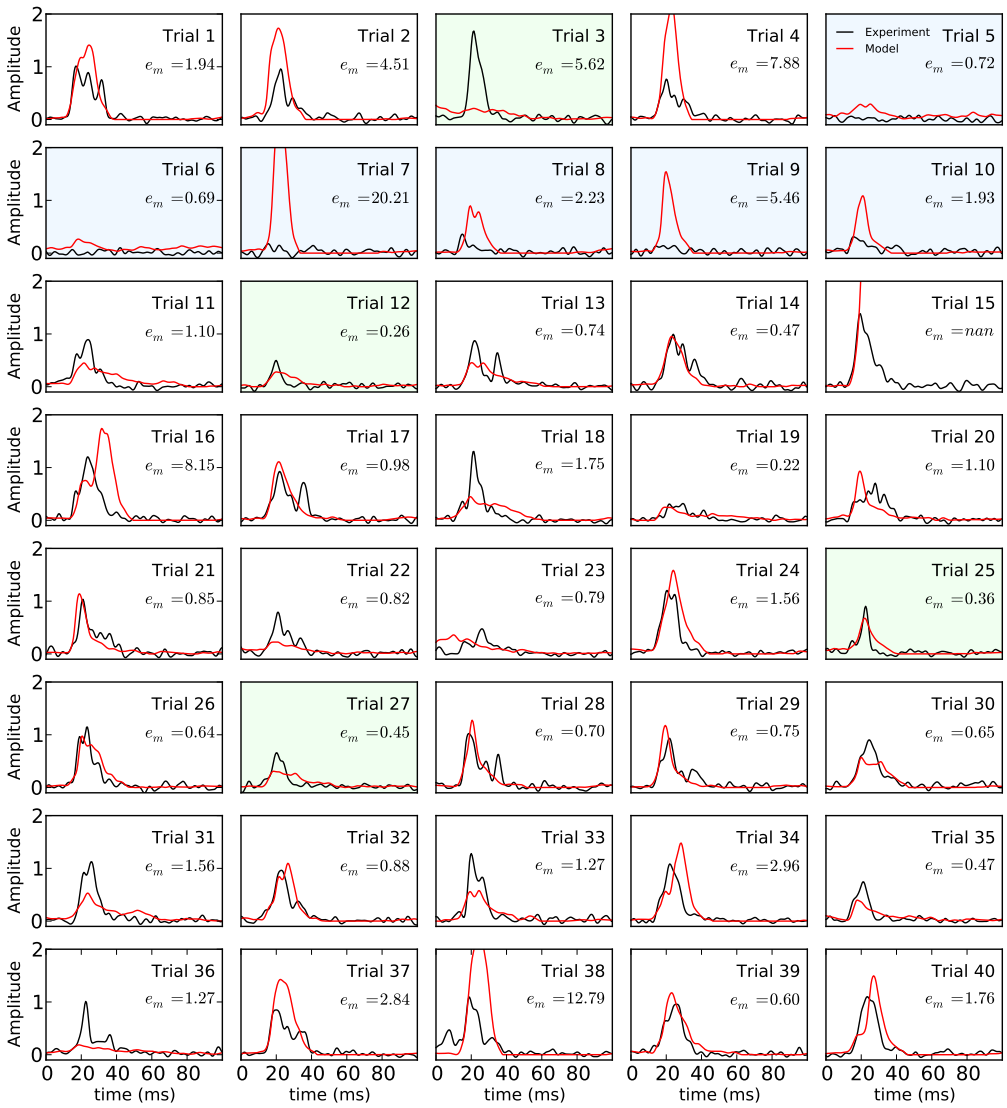

Figure 5.14: Model prediction of single trials. Model prediction (red line) and actual layer 4 population response (black line) for the 40 trials of stimulus condition 1 (a1, t1). Blue plots correspond to trials where cortex is in a silent state and green plots correspond to trials with substantial spontaneous activity prior to stimulus. Error is estimated according to a similar version of eq. (5.1).

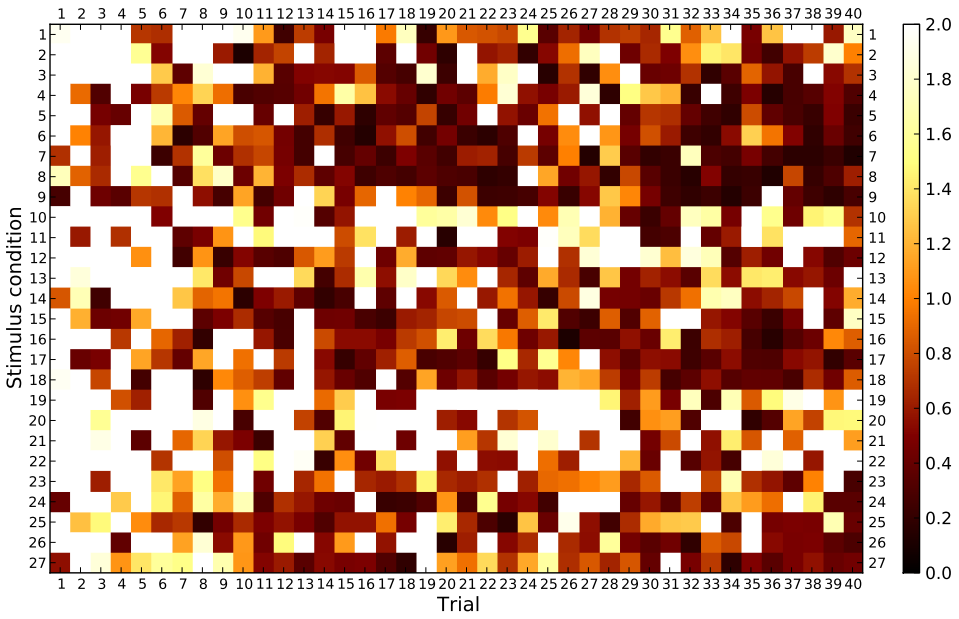

Figure 5.15: Summary of model errors. Image shows the magnitude of model errors for all trials across all stimulus conditions. Colorbar goes from low (black) to high (white) errors.

The dampening of stimulus response due to spontaneous activity seems to be predictable. Except for trial 3, the weak responses in all trials 12, 25 and 27 are relatively well predicted. This imply that thalamic firing rates are also dampened by spontaneous activity

Another weakness of the model is high sensitivity to ups and downs in the thalamic input. Single trial thalamic firing rates show oscillatory patterns and in trials like 16, 34 and 40 the modeled layer 4 population response are affected, by having a delayed peak. Trial averaging smooths the firing rates and hides away these oscillatory patterns. Consequently the average shape of firing-rate responses is predicted quite well, but small temporal deviations are not.

The pattern of model errors for all trials is similar to the pattern of thalamic peak responses in Figure 5.4. During the first trials of the experiment there is a high frequency of large errors, while for later trials large errors occur less often. Also, errors are more frequent among strong stimulus conditions than among weak stimulus conditions, as shown in Figure 5.15. Large model errors at strong stimulus conditions can partly be attributed to the instability of the model. A count shows that model responses blew up in 4% (42 out of 1080) of the trials. For trials with spontaneous activity, a count yields that 78% (111 out of 142) of the trials have an error of less than one, reinforcing the impression from observations in Figure 5.14. Looking at the overall performance of the model; 57% (611 out of 1080) of the trials have an error of less than 1. This result is impressive considering the large trial-to-trial variability, but it also demonstrates that the model is unreliable in predicting model responses.

Chapter 6

Discussion

6.1 Thalamocortical Model

It is unclear how features of the thalamic firing rate shape the cortical response, but since stronger stimuli are represented with stronger responses in firing rates in both thalamus and cortex I compared the relationship between the magnitudes of the input firing rate and the response firing rate. To quantify these magnitudes, I looked at the peak amplitude of the firing rate and the total firing (the area under the firing-rate curve). I also compared half-width of input and response in the same way to see whether firing-rate responses are more or less synchronous than firing-rate input.

The thalamocortical model gives a nonlinear relationship between both the peak and the total firing of the firing-rate input and response. For weak inputs (low amplitude/duration) this relationship is negative and increasing the amplitude or duration of the input leads to a decrease in the relative response magnitude. For stronger inputs, the nonlinear relationship is positive, and there is an increase in relative response magnitude. This reversal from a negative to a positive relationship is due to activation of the quadratic term in the activation function. The same results hold when looking at signal half-widths. In general, a larger response magnitude is also sharper and more synchronous.

A high amplitude input is sufficient to activate the non-linear growth in the response, but if it is a very short-lasting spike-shaped input, the response will not have time to build up, and consequently, it will be weak. On the other hand, a high amplitude input of a longer duration gives a strong response. If the duration is increased even further there will be a saturation in the response (as seen in Figures 4.3 and 4.5). In this case more firing on the input side will not result in more firing on the output side. If instead the amplitude is increased above a certain point, the model breaks down and there is an explosion in the response.

We also see that the model is very sensitive to the shape of the input at high amplitudes. A small difference in the broadness of the input makes large differences in the magnitude of the response. When the activation function receives large current inputs, the firing-rate response grows very fast. So, for stronger inputs the balance between excitatory and inhibitory activity is shifted

towards excitatory dominance. If the input is sufficiently strong, this excitatory dominance can cause a runaway effect where the inhibitory activity does not keep up and consequently the response blows up.

A steep rising slope gives the excitatory activity an initial dominance over the inhibitory activity. If the input grows fast to a high amplitude and decays slowly this effect is large, but if it instead decays fast the effect is much smaller. So, a steep rising slope can give larger responses, but only if the thalamic input does not decay fast.

In another similar study, Pinto et al. [29] explored a simple model of the response of one inhibitory and one excitatory neuron population to thalamic input in the rat whisker system. They found that a model with a strong inhibitory activity gives rise to many of the features observed in thalamocortical transformation of whisker information. A strong inhibitory population will always tend to dampen network activity, and therefore the excitatory population is activated best for strong synchronous input. Their model gave a strong response to triangular input with a steep onset and a weak response to triangular input with shallow onset. A similar effect is seen in the thalamocortical model, although it is less pronounced. Instead, the amplitude and the duration of the input are more important in shaping the response.

6.2 Single Trial Variability

There is substantial trial-to-trial variability in both the total (channel-summed) MUA and in the extracted population firing rates. Both strong and weak responses in the total MUA are in general reflected as strong and weak responses in all the neuron populations. However, as seen in Figure 5.11 there are also large differences in deviations from mean among the populations in many of the trials. We saw that many of the weak responses fell into two groups: they where either i) recorded during a long period of cortical silence (there was no stimulus response for many succeeding trials across all stimulus conditions) or ii) preceded by spontaneous activity.

The silent state in cortex could be due to anaesthesia, and such inactive states has been observed before under arousal or under stimulation of an area in the brain stem called the reticular formation [6]. Petersen et al. [25] showed that ongoing spontaneous activity suppressed action potential generation of neurons in layer 2/3 of barrel cortex.

Another source of variability in stimulus responses are cortical dynamics or alternating cortical states. During waking and REM sleep the cortex is in a *desynchronized* state, characterized by low amplitude, high frequency spontaneous activity patterns. Under sleep and some kinds of anaesthesia (e.g. urethane [7]) the brain is in a *synchronized* state with larger low frequency spontaneous activity patterns. In this synchronized state there are up and down phases, or alternations between network activity and network silence. Sensory-evoked activity can be highly dependent on the dynamics in these states [8, 14]. Sudden stimuli, such as an intense unexpected sound, can generate large stimulus responses in all brain states, whereas background noise would produce activity

only in a desynchronized state [13].

Haslinger et al. [15] investigated variability in stimulus responses during synchronous activity. Using data from the same experiment as considered in this project, they showed that the response magnitude of the MUA under a synchronized state could be predicted from oscillation patterns in the LFP. Further, they argued that this was a cortical phenomenon, unrelated to thalamic input. The MUA patterns in the data considered here are also characterized by spontaneous activity and stimulus responses surrounded by silent periods. Plots in Figure 1 in [8] show synchronous activity in auditory cortex of rats, and these patterns closely resemble the patterns seen in the MUA plotted in Figure 5.1.

Variability within populations was described by Kerr et al. [18]. In a two-photon imaging experiment they studied population responses in layer 2/3 of barrel cortex and they found highly variable responses within small groups of neurons. Watching responses in groups of 5-12 neurons to repeated stimulus, only rarely were the same firing patterns seen twice. Also, the number of activated neurons varied from trial to trial. Pooled results from many of these small groups suggested that responses were sparse, i.e. that only a small number of all the neurons in the population was active at the same time. Sparse coding is thought to be a widespread phenomena in cortex [23], and spontaneous activity also exhibits sparse activation patterns [18, 30].

With these results in mind, it is conceivable that variability could also stem from the recording setup. The electrode picks up signals from surrounding neurons, and if the neurons fire in a sparse and unreliable manner, it is possible that individual neurons (or small groups of neurons) close to the electrode contact can fire during some trials and stay quiet during others. Since the magnitude of the recorded MUA depends strongly on the distance to the active neuron, such close neurons could give large contributions during some trials. However, weak or strong responses are typically shared between populations at different recording positions, and therefore such variable firing patterns are likely a minor source of variability in the recorded data.

6.3 Modeling of Single Trials

The thalamocortical model does not predict single-trial cortical responses very reliably: there are examples of both qualitatively small and large errors. Although there are exceptions, in many of the spontaneous-activity trials weak responses are correctly predicted. This could indicate that spontaneous activity is a phenomenon influencing thalamus, and that the dampening of the cortical response is due to spontaneously dampened thalamic stimulus response. In trials of cortical silence, the model often wrongly predicts a strong stimulus response. Therefore, responses in thalamus seems to be unaffected by the lack of response in cortex. More generally, for trials at both weak and strong stimuli, weak and strong responses are relatively well predicted. However, the predicted magnitude is often far off, giving large errors.

There are two main reasons for the large errors: 1) The model is highly sensitive for strong inputs, best illustrated with the response blow-ups present in

many trial predictions. 2) The model is tuned to trial-averaged data, and these are stereotypical smoothed versions of single-trial stimulus responses. Thalamic firing rates have oscillatory patterns that are averaged out, and single-trial cortical firing rates also vary from the typical trial-averaged bell shape. The oscillations in the thalamic firing rate give unnatural model responses, and similarly, the model often fails to predict the shape of the experimental cortical firing rates.

Many different theories try to explain how the brain codes information. Arabzadeh et al. [2] found that population firing rates could be used to decode information about the vibration speed of a whisker. Panzeri and Diamond [24] argue that it is possible that information in sensory processing is carried by precise spike times, and that much of this information is lost when looking solely at firing rates. There are countless articles taking one view or the other, and even if population firing rates are not the only answer, models like the thalamocortical model can give new insights into the properties of cortical networks [4].

6.4 Conclusion

The thalamocortical model could not predict single-trial responses reliably. Trial averaging hides away features in the single-trial data which might be important for gaining a better understanding of the thalamocortical network. It is therefore important to be aware of this gap between what the brain does on average and what it does in real-time.

An interesting question that arises from this project is whether a population firing rate model can also be fitted to single-trial data. For example, such a model could be tuned to the cortical state using state-dependent parameters.

Bibliography

[1] Arabzadeh, E. [n.d.]. Neural Coding Lab.

[2] Arabzadeh, E., Panzeri, S. and Diamond, M. E. [2004]. Whisker vibration information carried by rat barrel cortex neurons., *The Journal of neuroscience : the official journal of the Society for Neuroscience* **24**(26): 6011–20.

[3] Bialek, W., Rieke, F., de Ruyter van Steveninck, R. R. and Warland, D. [1991]. Reading a neural code., *Science (New York, N.Y.)* **252**(5014): 1854–7.

[4] Blomquist, P., Devor, A., Indahl, U. G., Ulbert, I., Einevoll, G. T. and Dale, A. M. [2009]. Estimation of thalamocortical and intracortical network models from joint thalamic single-electrode and cortical laminar-electrode recordings in the rat barrel system., *PLoS computational biology* **5**(3): e1000328.

[5] Buzsáki, G. [2004]. Large-scale recording of neuronal ensembles., *Nature neuroscience* **7**(5): 446–51.

[6] Castro-Alamancos, M. a. and Oldford, E. [2002]. Cortical sensory suppression during arousal is due to the activity-dependent depression of thalamocortical synapses, *The Journal of Physiology* **541**(1): 319–331.

[7] Clement, E. a., Richard, A., Thwaites, M., Ailon, J., Peters, S. and Dickson, C. T. [2008]. Cyclic and sleep-like spontaneous alternations of brain state under urethane anaesthesia., *PloS one* **3**(4): e2004.

[8] Curto, C., Sakata, S., Marguet, S., Itskov, V. and Harris, K. D. [2009]. A simple model of cortical dynamics explains variability and state dependence of sensory responses in urethane-anesthetized auditory cortex., *The Journal of neuroscience : the official journal of the Society for Neuroscience* **29**(34): 10600–12.

[9] Dayan, P. and Abbott, L. [2001]. *Theoretical Neuroscience - Computational and Mathematical Modeling of Neural Systems*, The MIT Press, Cambridge, Massachusetts.

[10] Di, S., Baumgartner, C. and Barth, D. S. [1990]. Laminar Analysis of Extracellular in Rat Vibrissa / Barrel Cortex Field Potentials, *Journal of Neurophysiology* **63**(4): 832–840.

[11] Einevoll, G. T., Pettersen, K. H., Devor, A., Ulbert, I., Halgren, E. and Dale, A. M. [2007]. Laminar population analysis: estimating firing rates and evoked synaptic activity from multielectrode recordings in rat barrel cortex., *Journal of neurophysiology* **97**(3): 2174–90.

[12] Fox, K. [2008]. *Barrel Cortex*, 1. edn, Cambridge University Press.

[13] Harris, K. D. and Thiele, A. [2011]. Cortical state and attention., *Nature reviews. Neuroscience* **12**(9): 509–23.

[14] Hasenstaub, A., Sachdev, R. N. S. and McCormick, D. a. [2007]. State changes rapidly modulate cortical neuronal responsiveness., *The Journal of neuroscience : the official journal of the Society for Neuroscience* **27**(36): 9607–22.

[15] Haslinger, R., Ulbert, I., Moore, C. I., Brown, E. N. and Devor, A. [2006]. Analysis of LFP phase predicts sensory response of barrel cortex, *Journal of Neurophysiology* **96**(3): 1658–1663.

[16] Hubel, D. H. . [1957]. Tungsten Microelectrode for Recording from Single Units, *American Association for the Advancement of Science* **125**(3247): 549–550.

[17] Kandel, E., Schwartz, J. and Jessell, T. [2000]. *Principles of neural science.*

[18] Kerr, J. N. D., Kock, C. P. J. D., Greenberg, D. S., Bruno, R. M. and Sakmann, B. [2007]. Spatial Organization of Neuronal Population Responses in Layer 2 / 3 of Rat Barrel Cortex, **27**(48): 13316–13328.

[19] Koch, C. [2001]. Neuronal and Synaptic Packing Density, *http://www.dna.caltech.edu/courses/cns187/references/neuronal_densities.pdf* .

[20] Lay, D. [2012]. *Linear Algebra and Its Applications*, 4th ed edn, Pearson.

[21] Lefort, S., Tomm, C., Floyd Sarria, J.-C. and Petersen, C. C. H. [2009]. The excitatory neuronal network of the C2 barrel column in mouse primary somatosensory cortex., *Neuron* **61**(2): 301–16.

[22] Masquelier, T. [2013]. Neural variability, or lack thereof., *Frontiers in computational neuroscience* **7**(February): 7.

[23] Olshausen, B. a. and Field, D. J. [2004]. Sparse coding of sensory inputs., *Current opinion in neurobiology* **14**(4): 481–7.

[24] Panzeri, S. and Diamond, M. E. [2010]. Information Carried by Population Spike Times in the Whisker Sensory Cortex can be Decoded Without Knowledge of Stimulus Time., *Frontiers in synaptic neuroscience* **2**(June): 17.

[25] Petersen, C. C. H., Hahn, T. T. G., Mehta, M., Grinvald, A. and Sakmann, B. [2003]. Interaction of sensory responses with spontaneous depolarization in layer 2/3 barrel cortex., *Proceedings of the National Academy of Sciences of the United States of America* **100**(23): 13638–43.

[26] 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.

[27] Pettersen, K. H., Lindén, H., Dale, A. M. and Einevoll, G. T. [2011]. Extracellular spikes and CSD, *in* R. Brette and A. Destexhe (eds), *Handbook of Neural Activity Measurement*, number April, Cambridge University Press, Cambridge, Uk.

[28] Pinto, D. J., Brumberg, J. C. and Simons, D. J. [2000]. Circuit dynamics and coding strategies in rodent somatosensory cortex., *Journal of neurophysiology* **83**(3): 1158–66.

[29] Pinto, D. J., Hartings, J. a., Brumberg, J. C. and Simons, D. J. [2003]. Cortical damping: analysis of thalamocortical response transformations in rodent barrel cortex., *Cerebral cortex (New York, N.Y. : 1991)* **13**(1): 33–44.

[30] Sakata, S. and Harris, K. D. [2009]. Laminar structure of spontaneous and sensory-evoked population activity in auditory cortex., *Neuron* **64**(3): 404–18.

[31] Shoykhet, M., Doherty, D. and Simons, D. J. [2000]. Coding of deflection velocity and amplitude by whisker primary afferent neurons: implications for higher level processing., *Somatosensory & motor research* **17**(2): 171–80.

[32] Simons, DJ and Carvell, G. [1989]. Thalamocortical response transformation in the rat vibrissa/barrel system, *Journal of Neurophysiology* **61**(2): 311—-330.

[33] Somogyvári, Z., Zalányi, L., Ulbert, I. and Erdi, P. [2005]. Model-based source localization of extracellular action potentials., *Journal of neuroscience methods* **147**(2): 126–37.

[34] Wilent, W. B. and Contreras, D. [2004]. Synaptic responses to whisker deflections in rat barrel cortex as a function of cortical layer and stimulus intensity., *The Journal of neuroscience : the official journal of the Society for Neuroscience* **24**(16): 3985–98.

[35] Wilent, W. B. and Contreras, D. [2005]. Dynamics of excitation and inhibition underlying stimulus selectivity in rat somatosensory cortex., *Nature neuroscience* **8**(10): 1364–70.

[36] Woolsey, Thomas A and Van der Loos, H. [1970]. The structural organization of layer IV in the somatosensory region (SI) of mouse cerebral cortex: the description of a cortical field composed of discrete cytoarchitectonic units, *Brain research* **17**(2): 205—-242.

Appendix A

Thalamocortical Model Predictions for Single Trials

This appendix contains figures of model predictions for all stimulus conditions. For each stimulus condition, there is one figure with thalamic firing rates for all 40 trials and one figure with model predictions for all 40 trials.

In the figure showing thalamic firing rates (purple lines), the trial average (black dotted line) is also added to each plot for comparison. Thalamic firing rates are low-pass filtered and normalized as described in section 3.2.

In the figure showing modeled layer 4 population firing rates (red lines), the layer 4 population firing rates extracted from the experimental data (black line) is added to each plot for comparison. The error e_m is calculated from eq. (5.2). This error is scaled so errors are comparable across stimulus conditions. Plots shaded in blue correspond to trials where cortex is in a silent state and plots shaded in green correspond to trials with spontaneous activity prior to stimulus.

A.1 Stimulus Condition 1

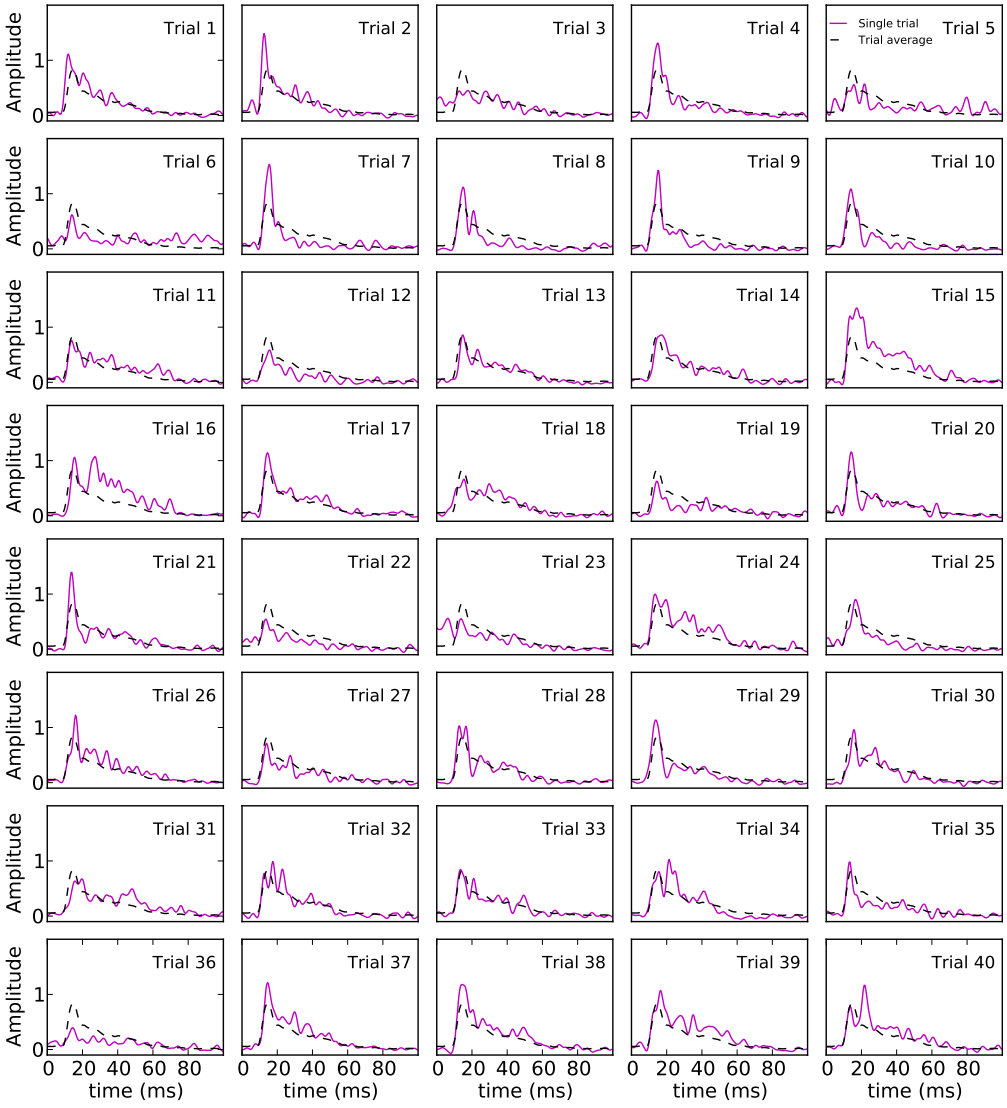

Figure A.1: Thalamic population firing rates.

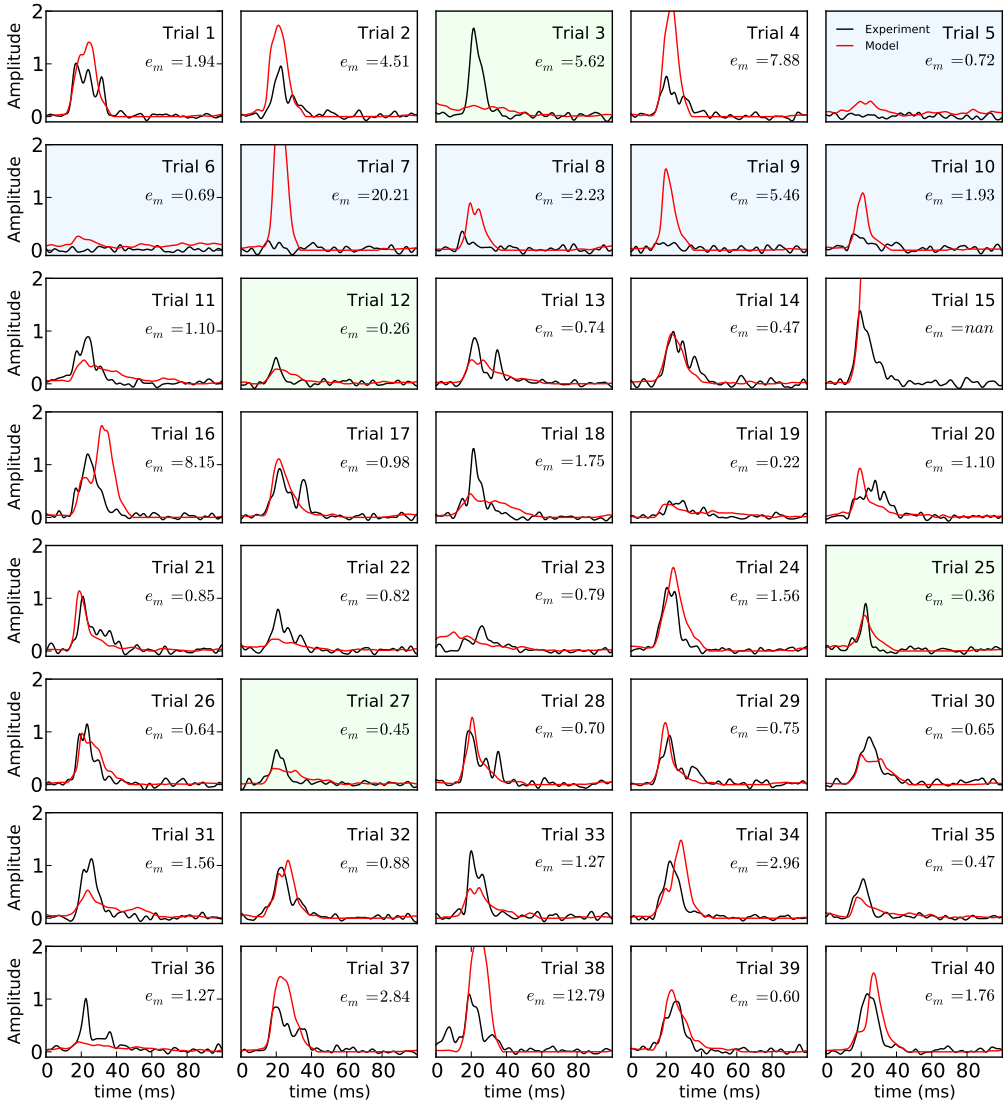

Figure A.2: Modeled layer 4 population firing rates.

A.2 Stimulus Condition 2

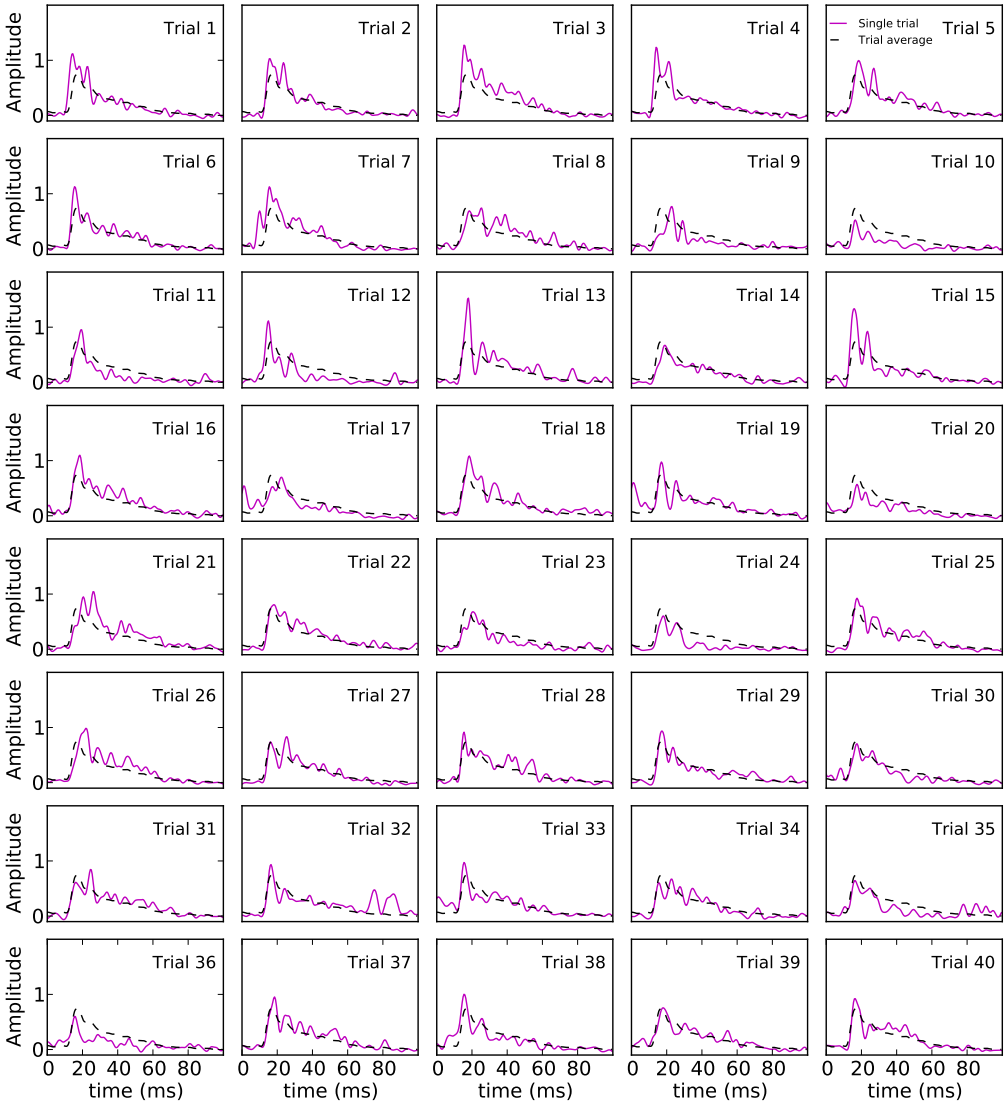

Figure A.3: Thalamic population firing rates.

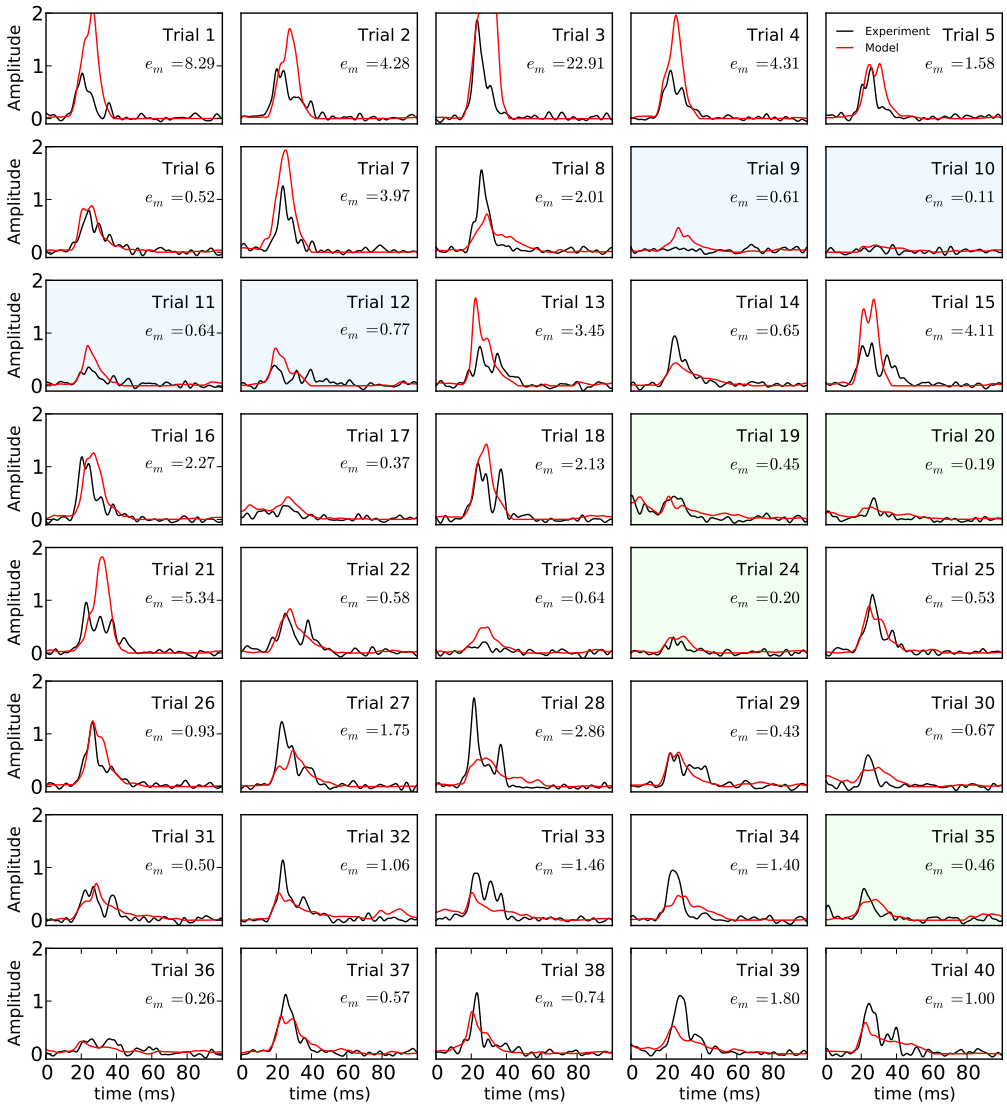

Figure A.4: Modeled layer 4 population firing rates.

A.3 Stimulus Condition 3

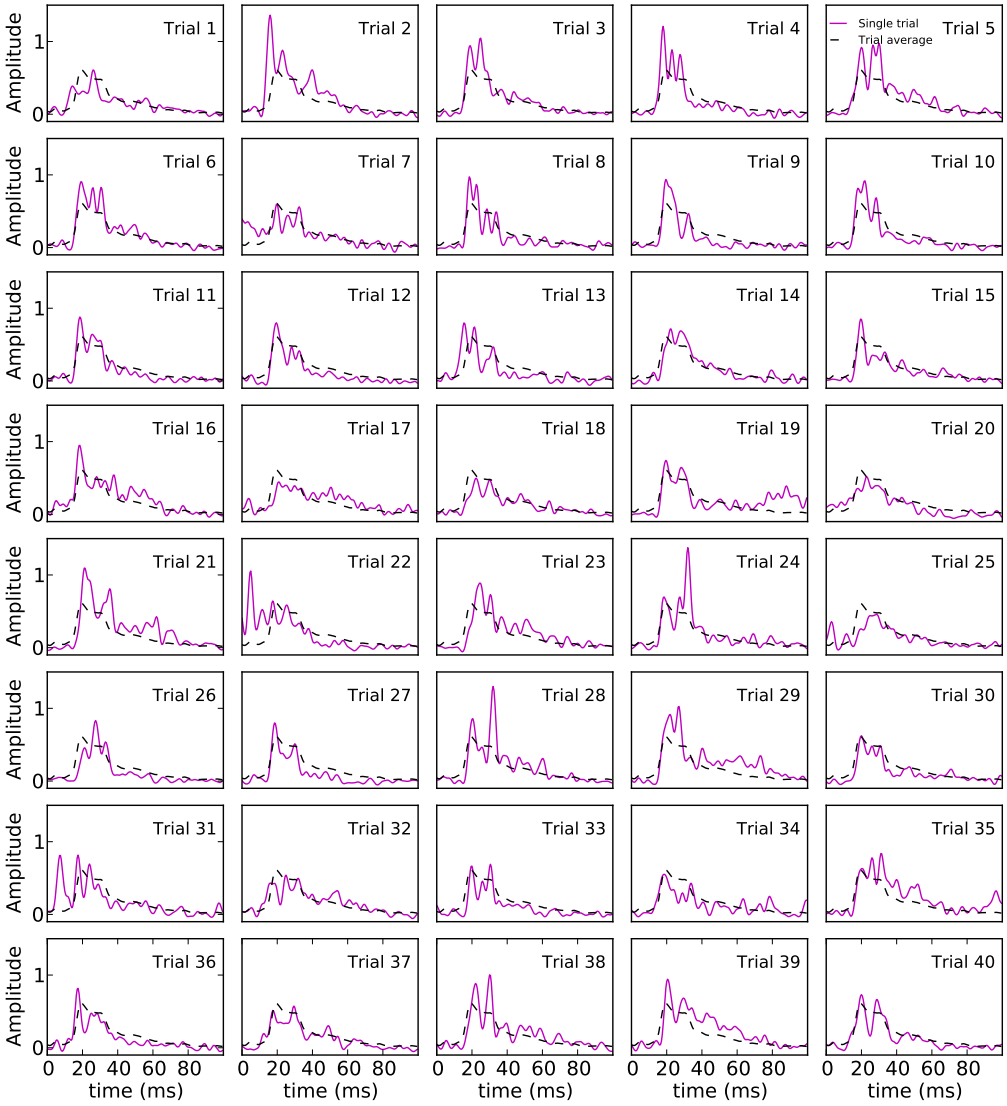

Figure A.5: Thalamic population firing rates.

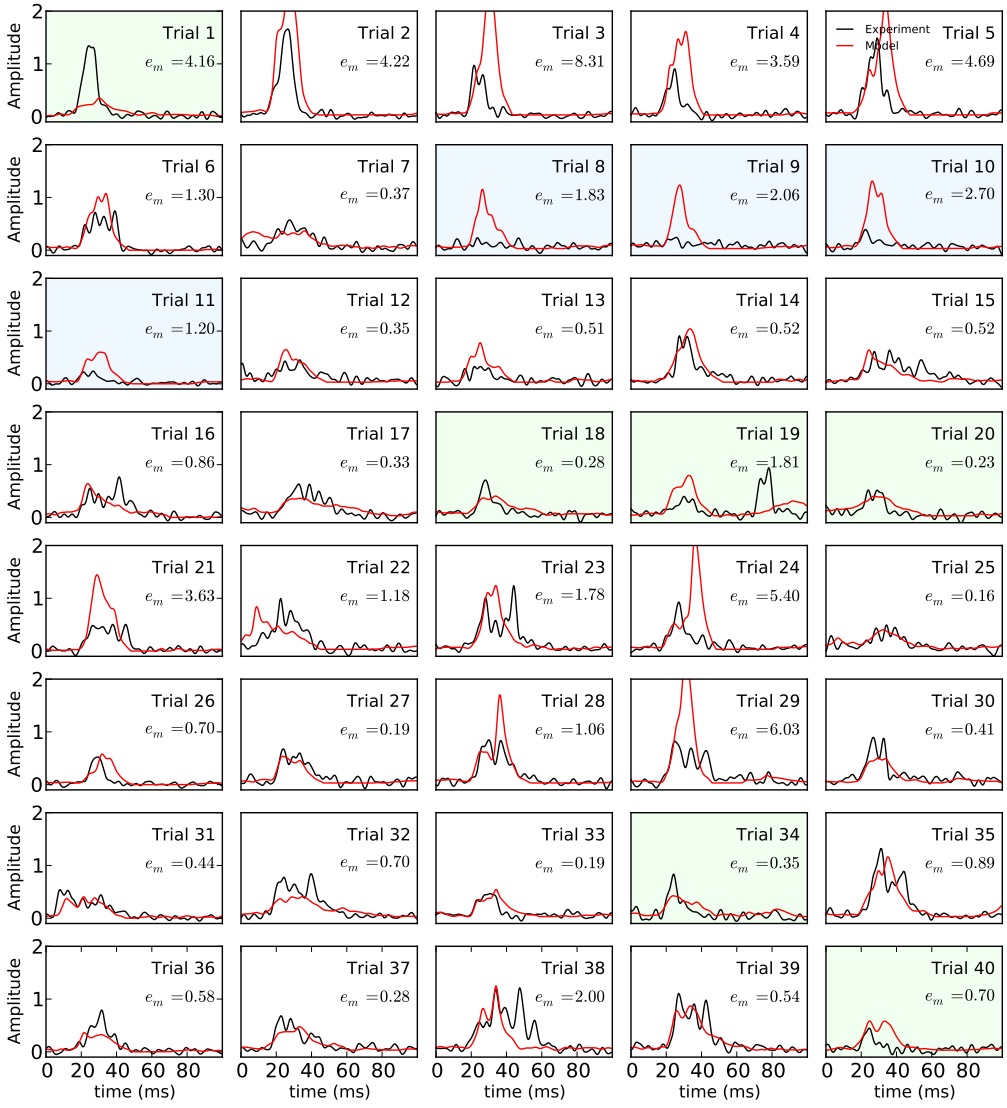

Figure A.6: Modeled layer 4 population firing rates.

A.4 Stimulus Condition 4

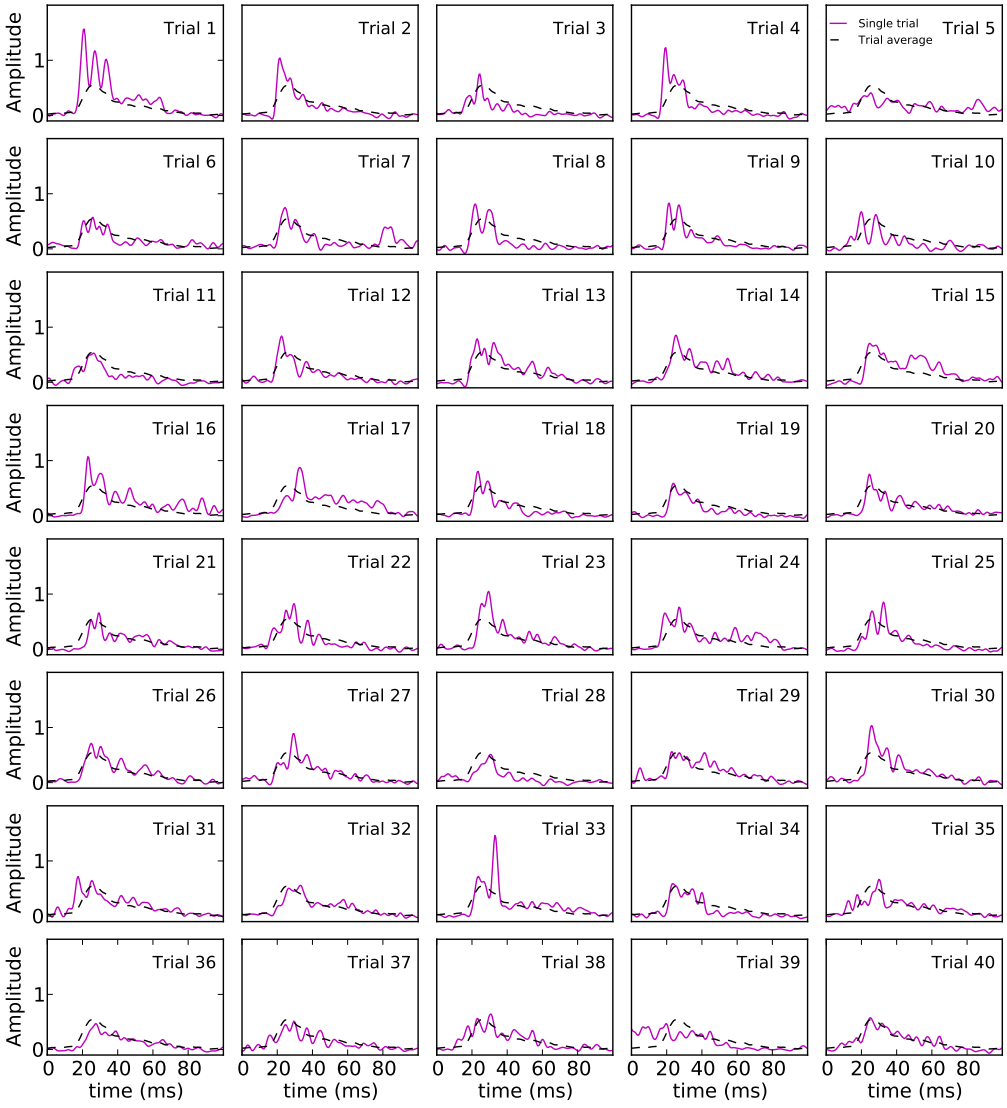

Figure A.7: Thalamic population firing rates.

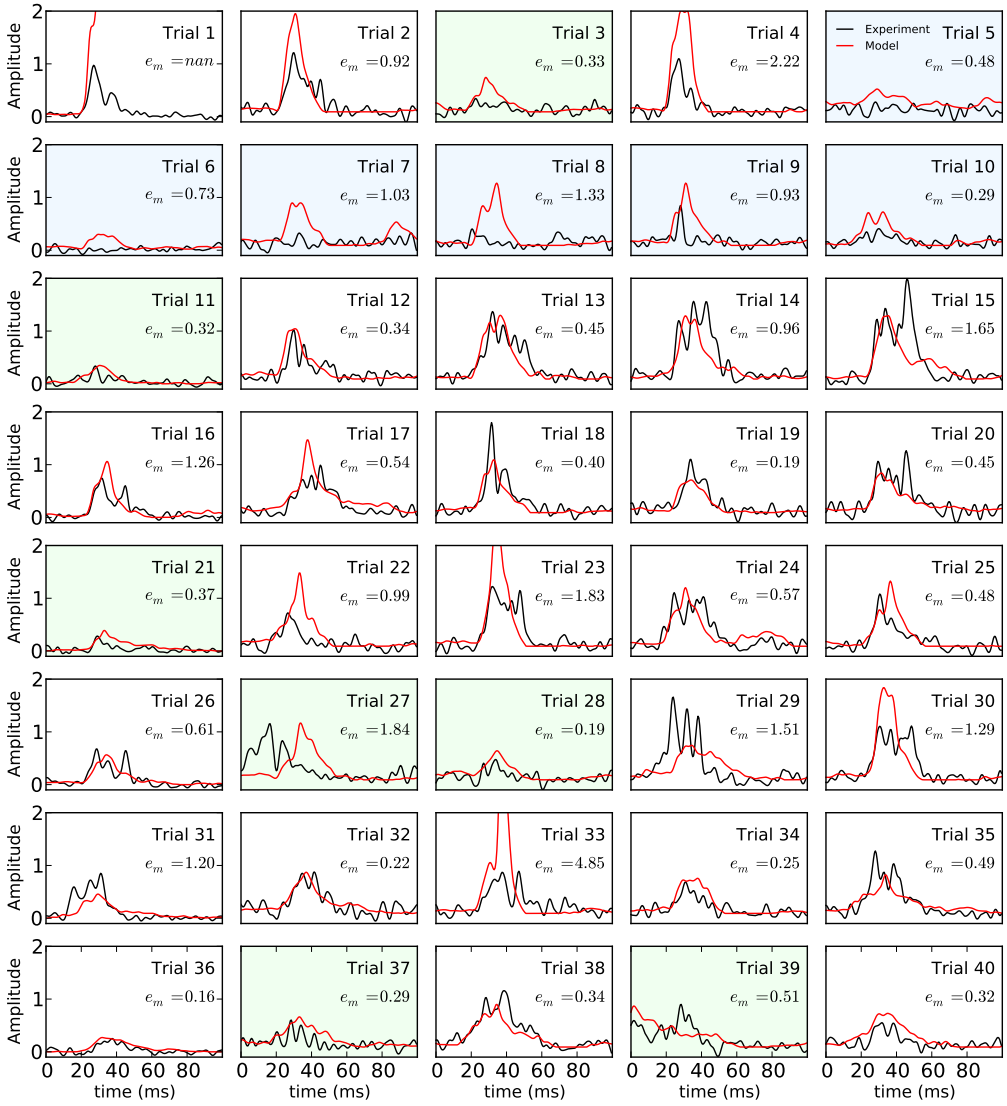

Figure A.8: Modeled layer 4 population firing rates.

A.5 Stimulus Condition 5

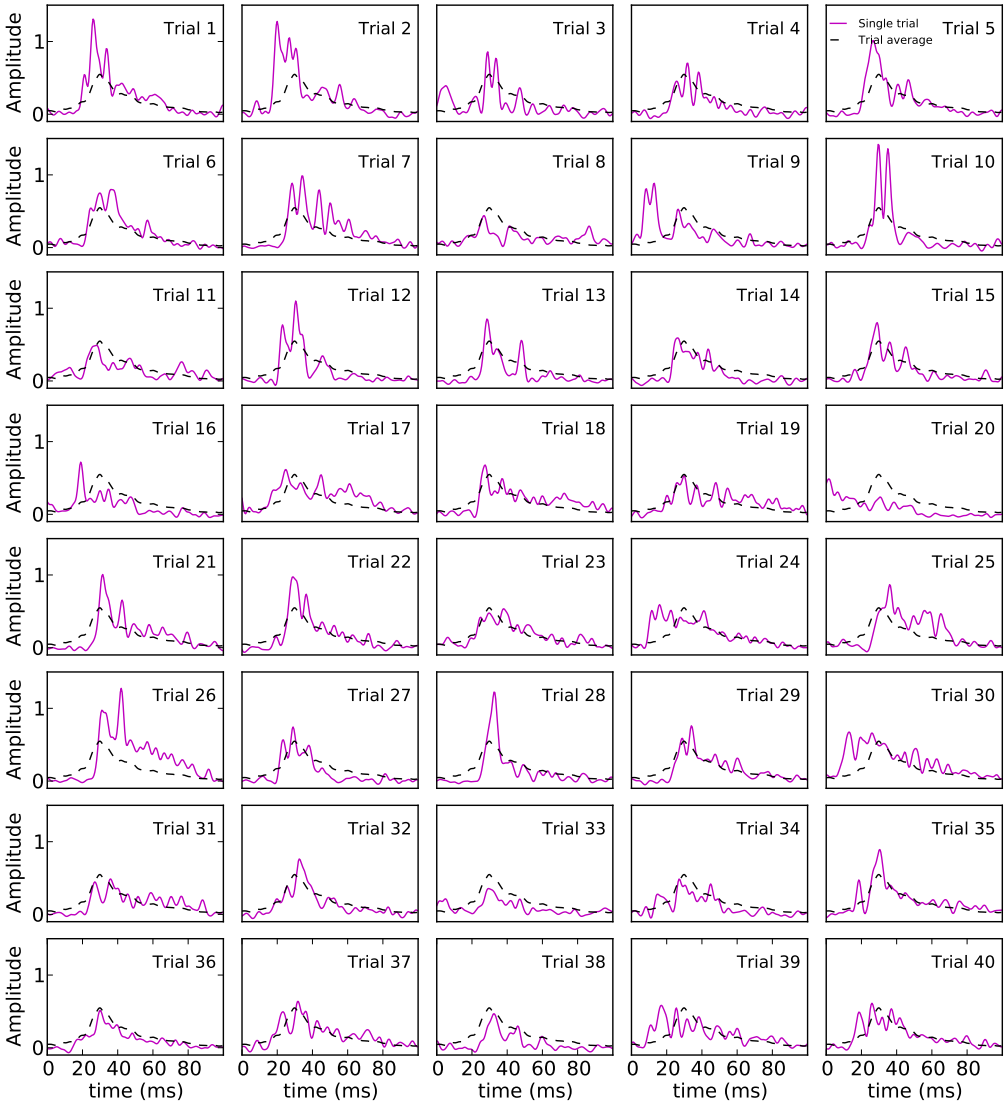

Figure A.9: Thalamic population firing rates.

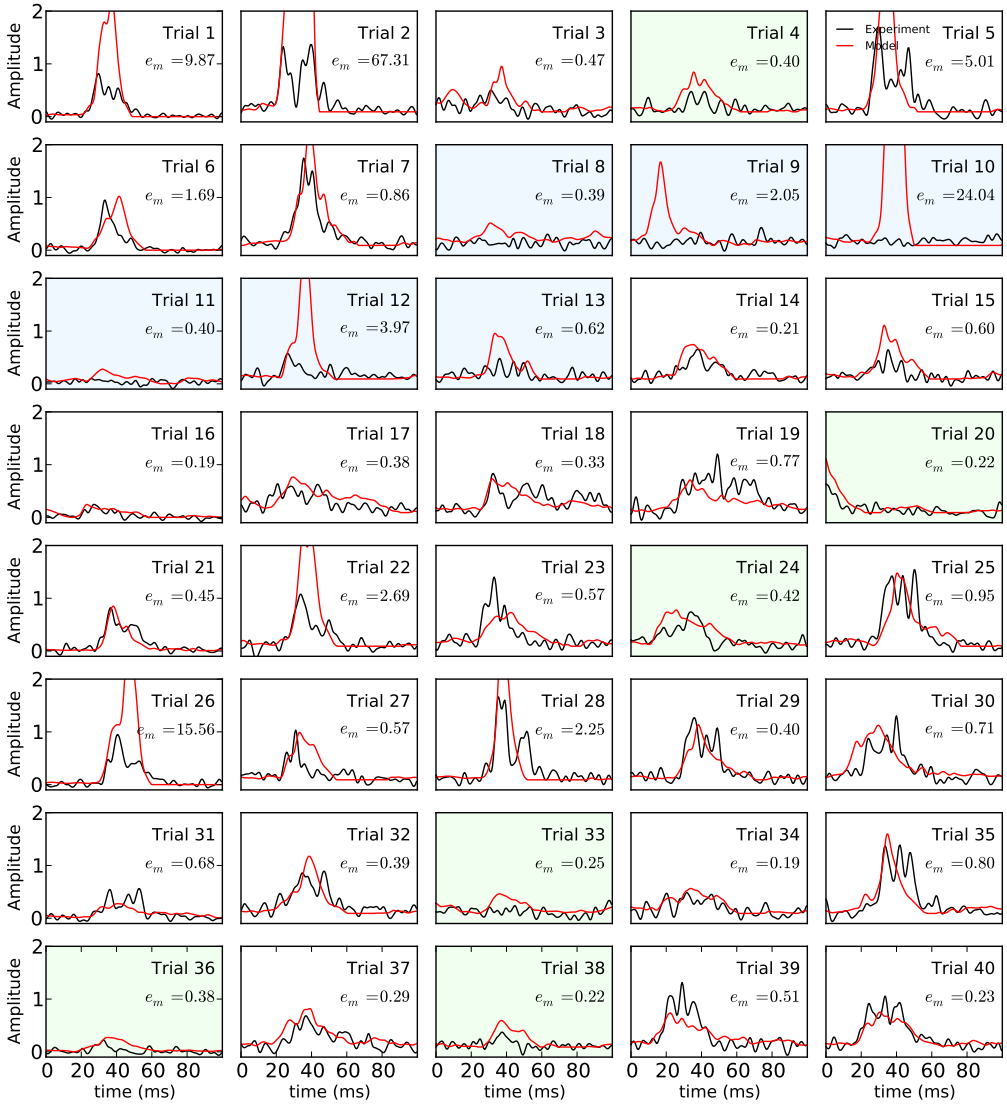

Figure A.10: Modeled layer 4 population firing rates.

A.6 Stimulus Condition 6

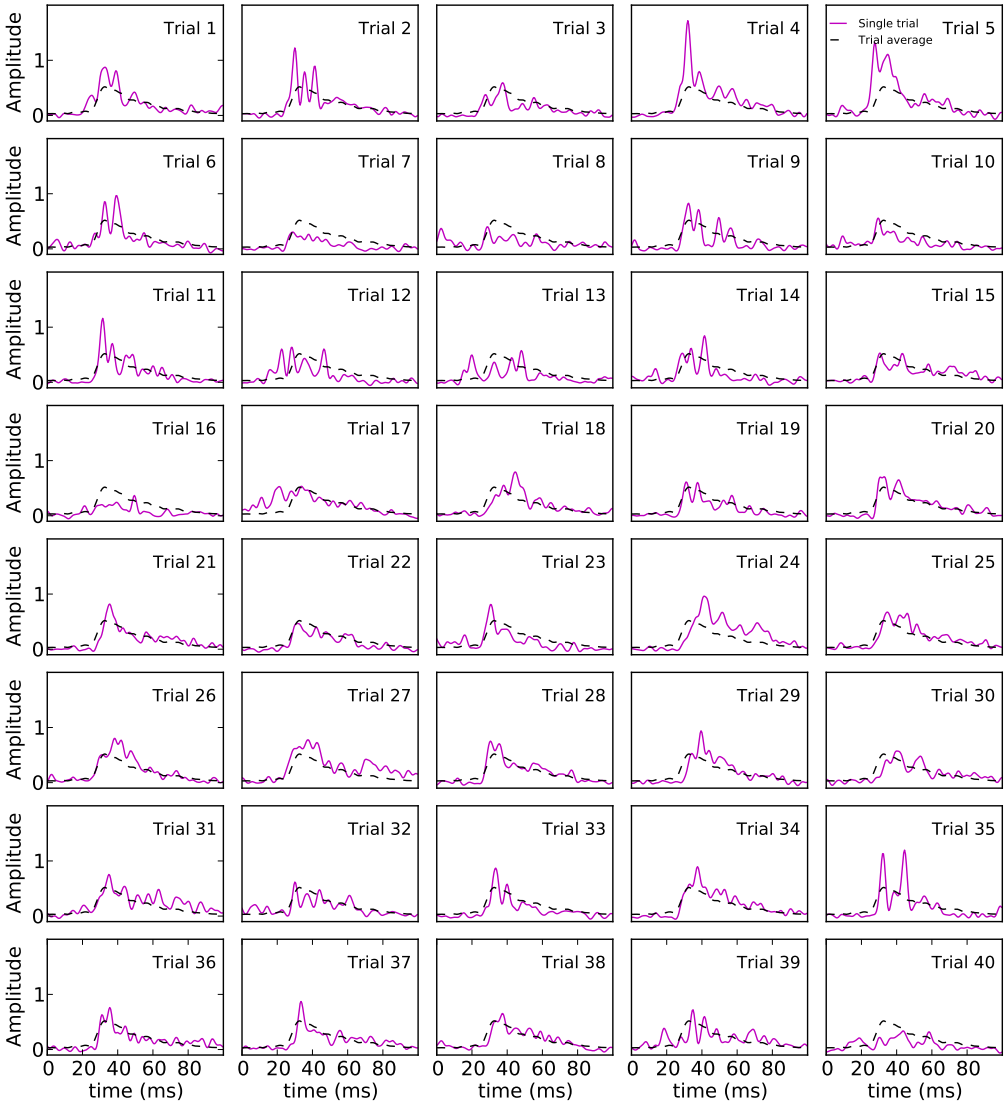

Figure A.11: Thalamic population firing rates.

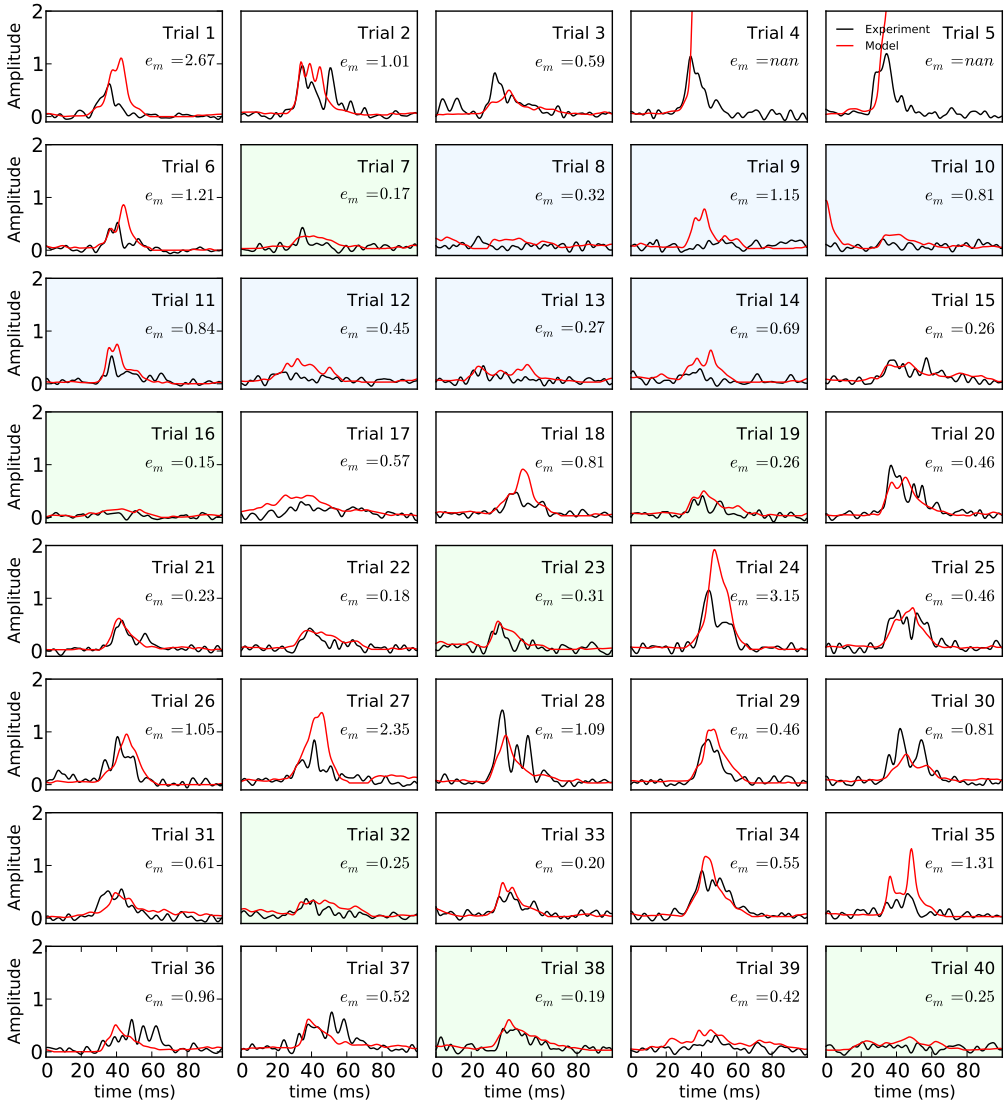

Figure A.12: Modeled layer 4 population firing rates.

A.7 Stimulus Condition 7

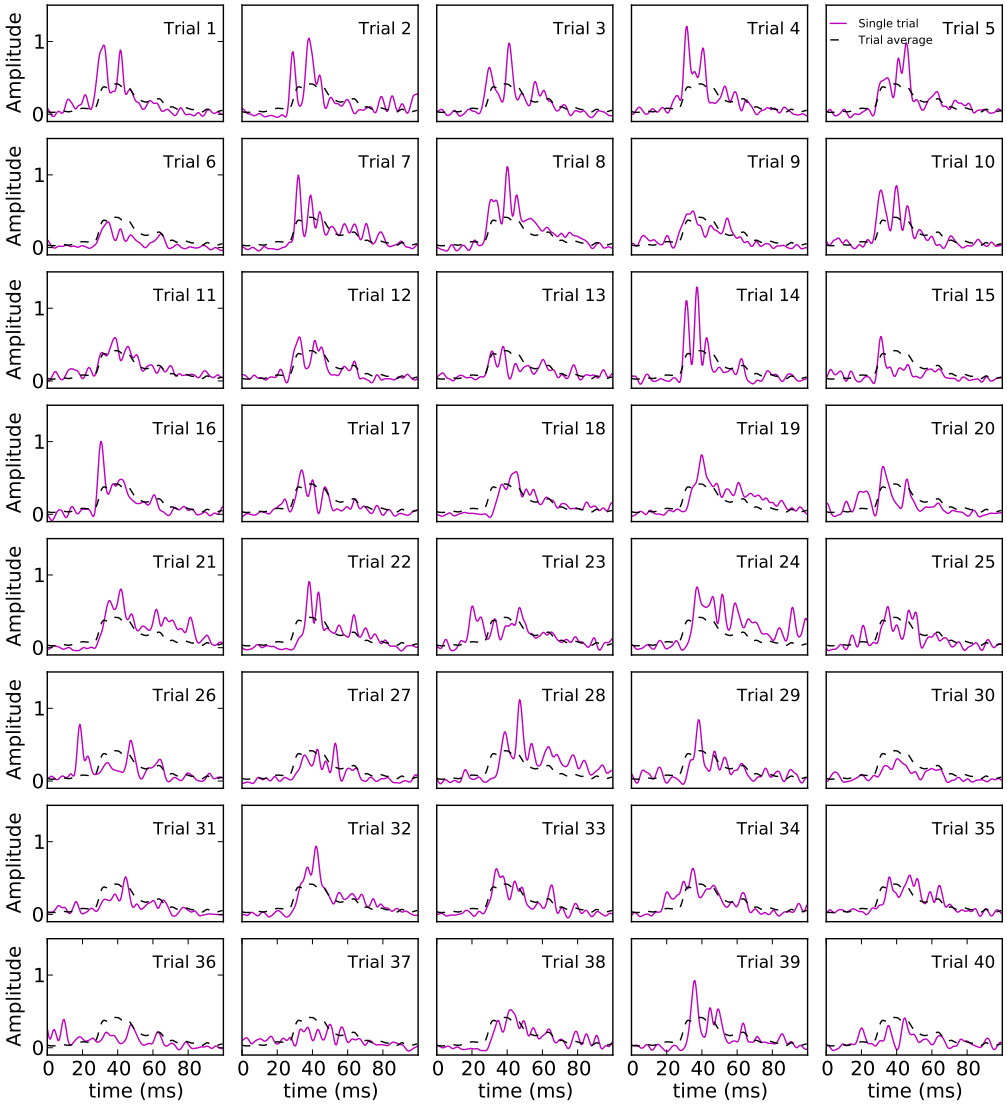

Figure A.13: Thalamic population firing rates.

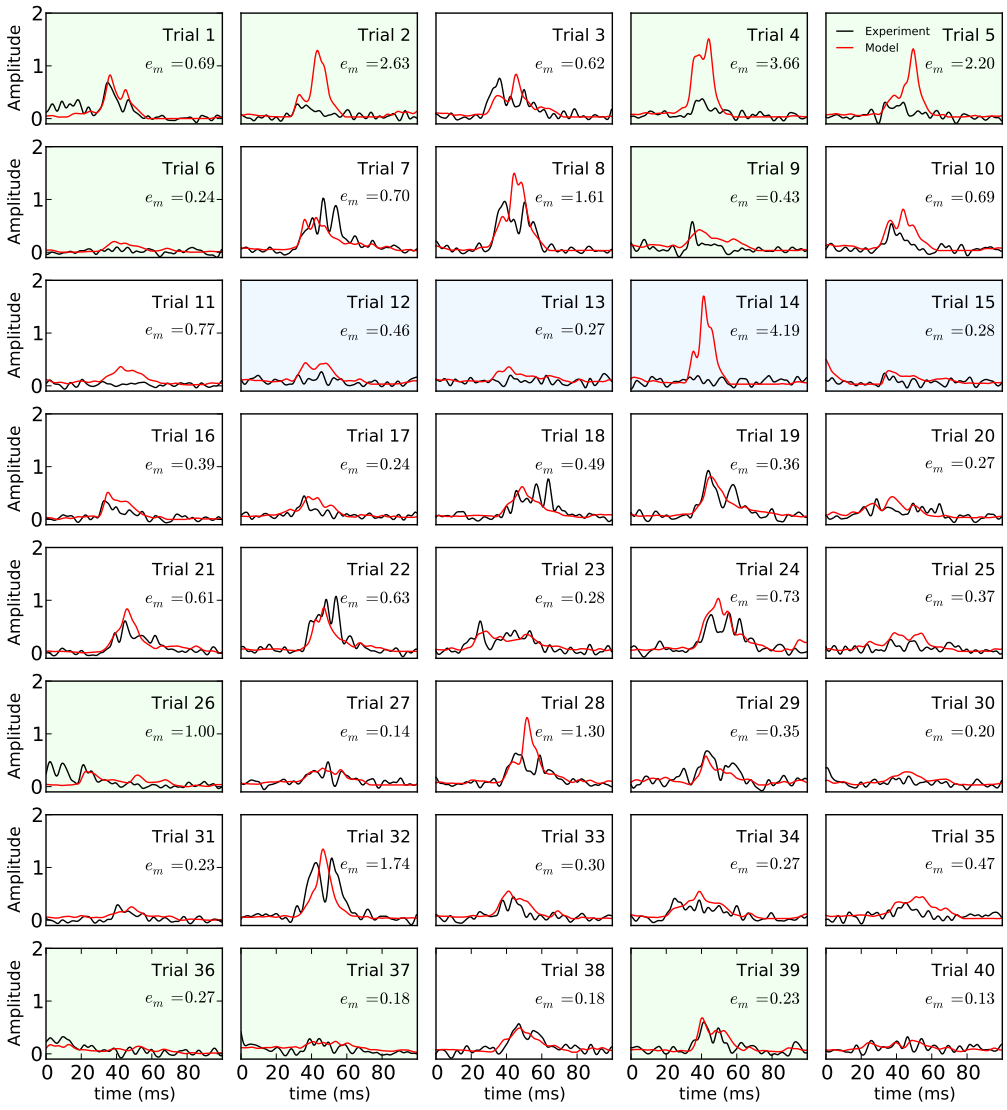

Figure A.14: Modeled layer 4 population firing rates.

A.8 Stimulus Condition 8

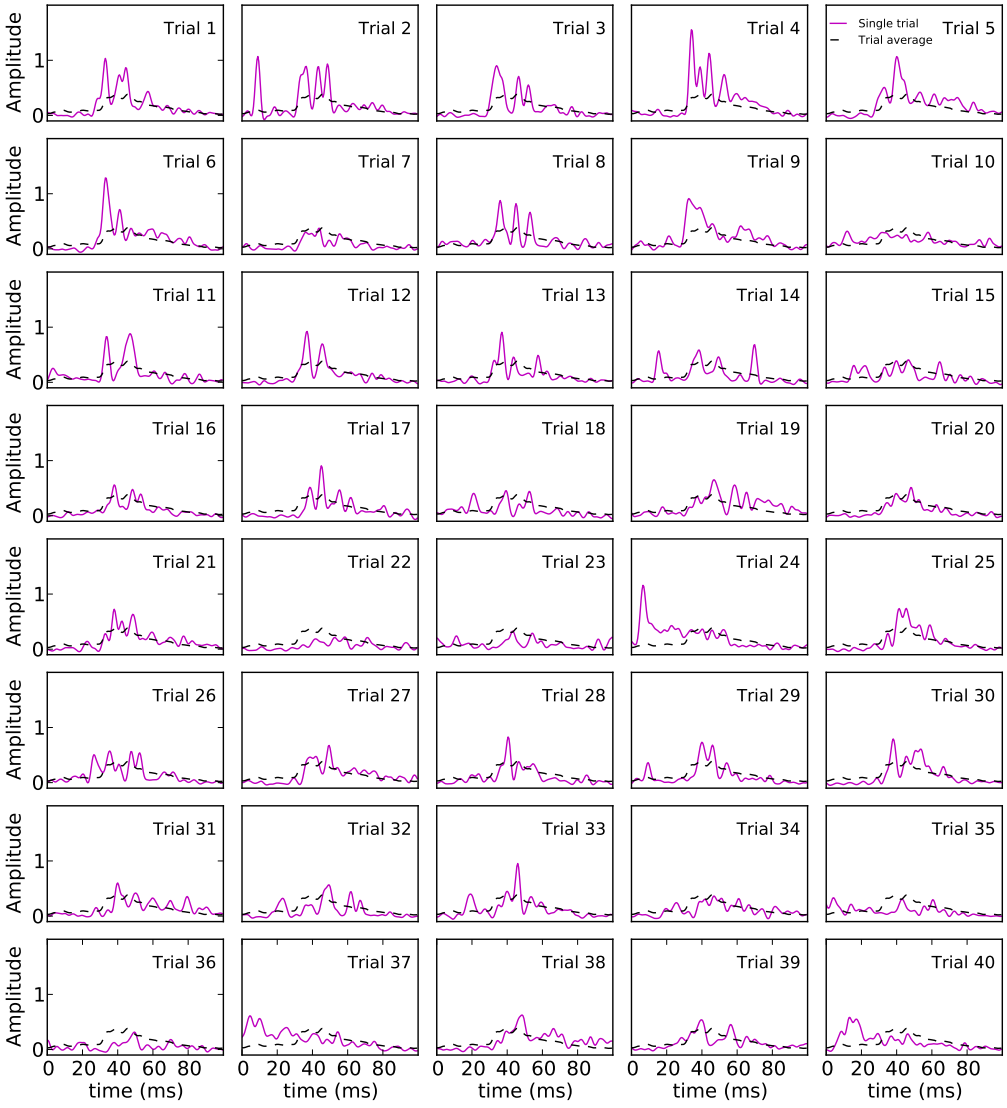

Figure A.15: Thalamic population firing rates.

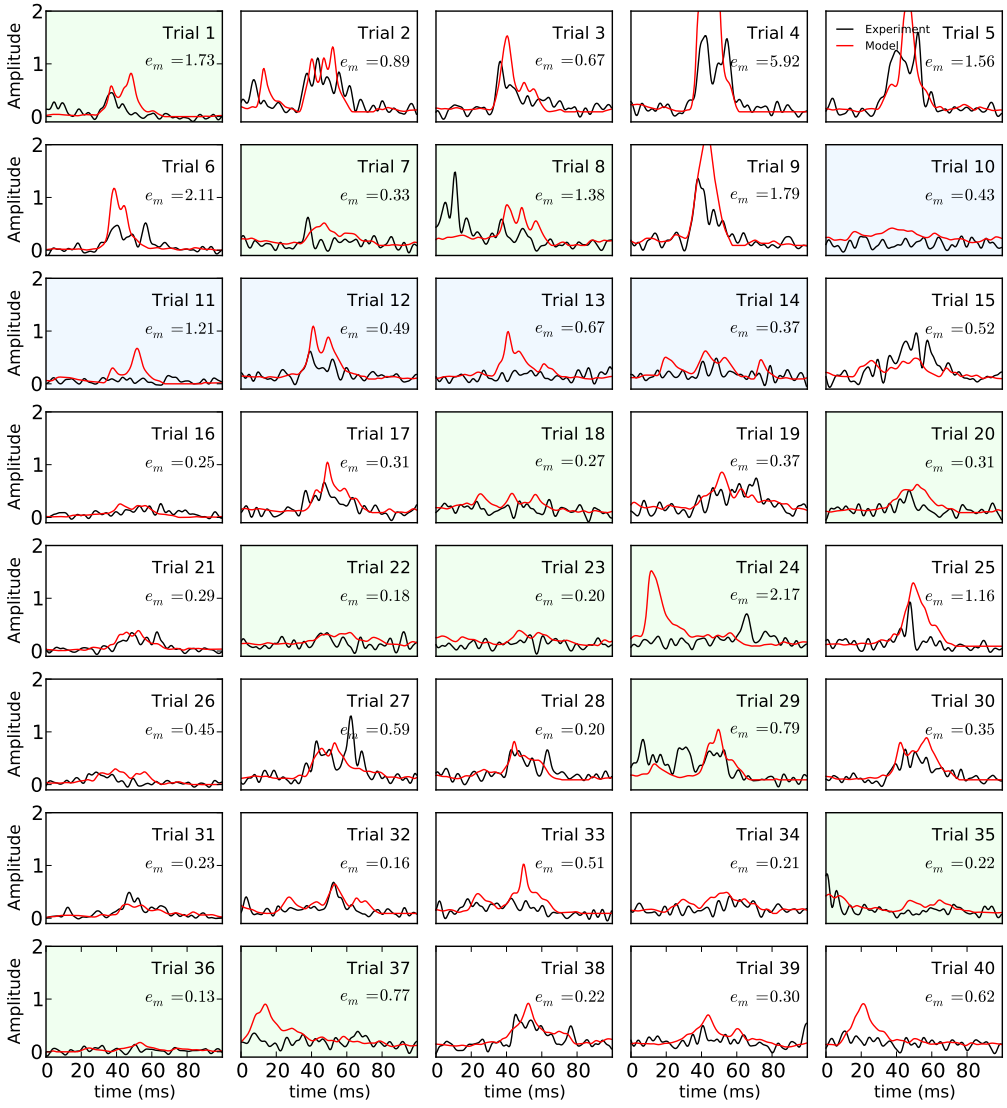

Figure A.16: Modeled layer 4 population firing rates.

A.9 Stimulus Condition 9

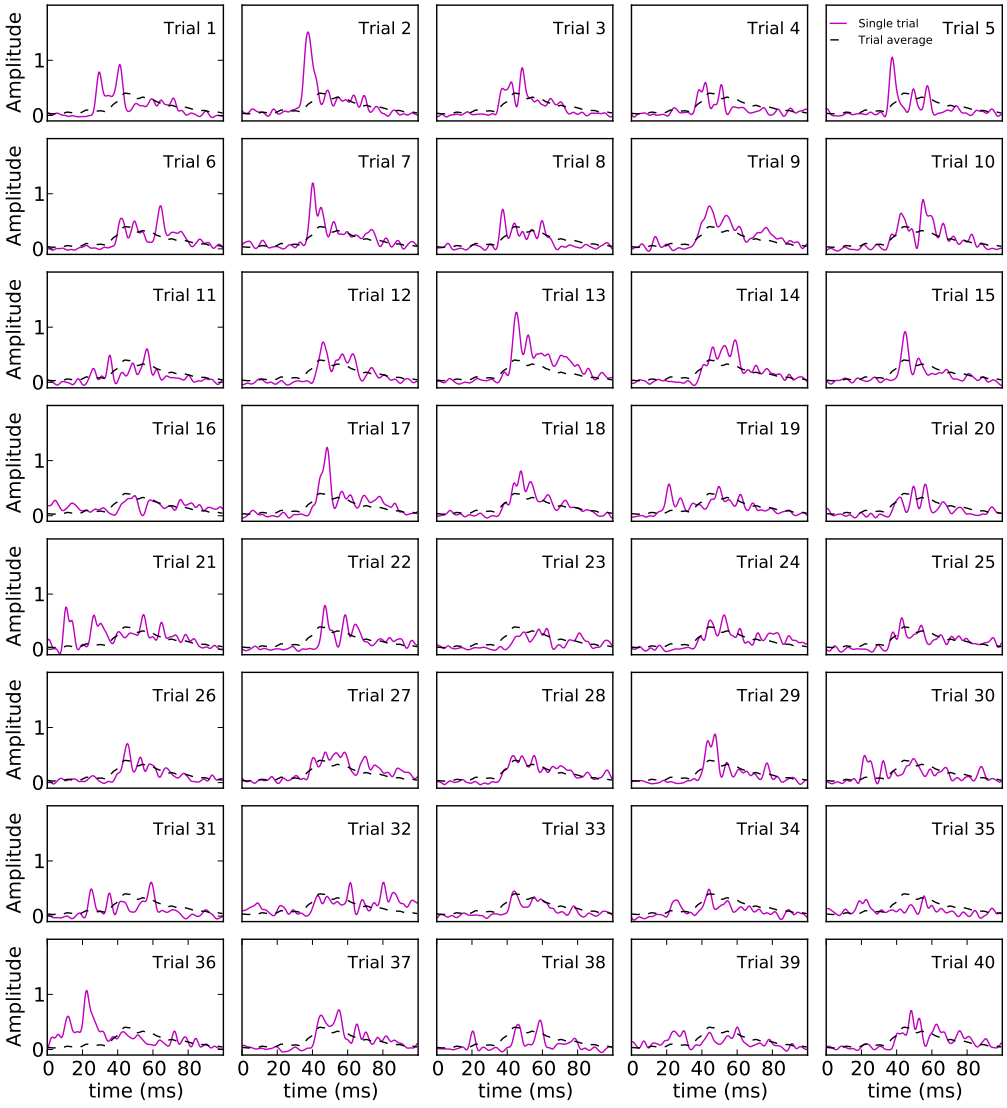

Figure A.17: Thalamic population firing rates.

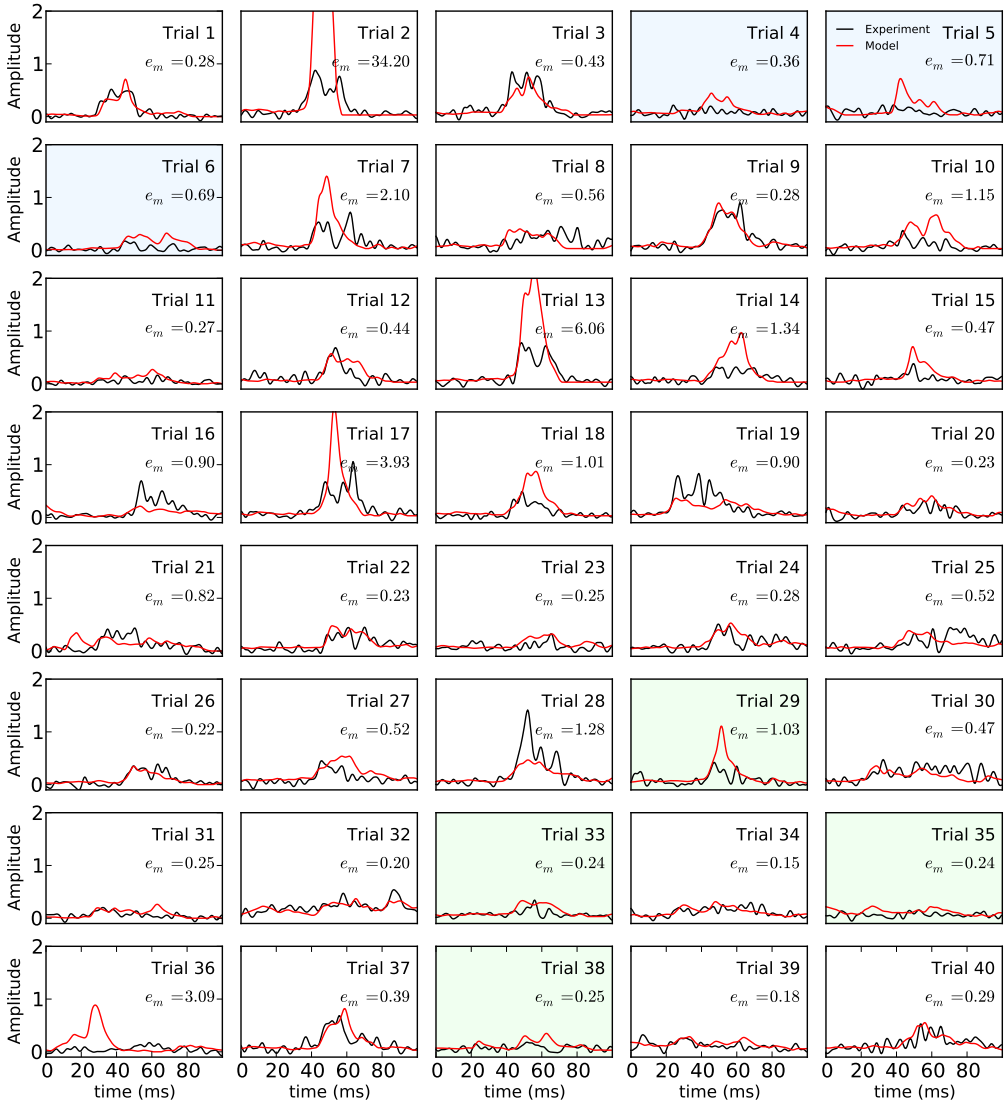

Figure A.18: Modeled layer 4 population firing rates.

A.10 Stimulus Condition 10

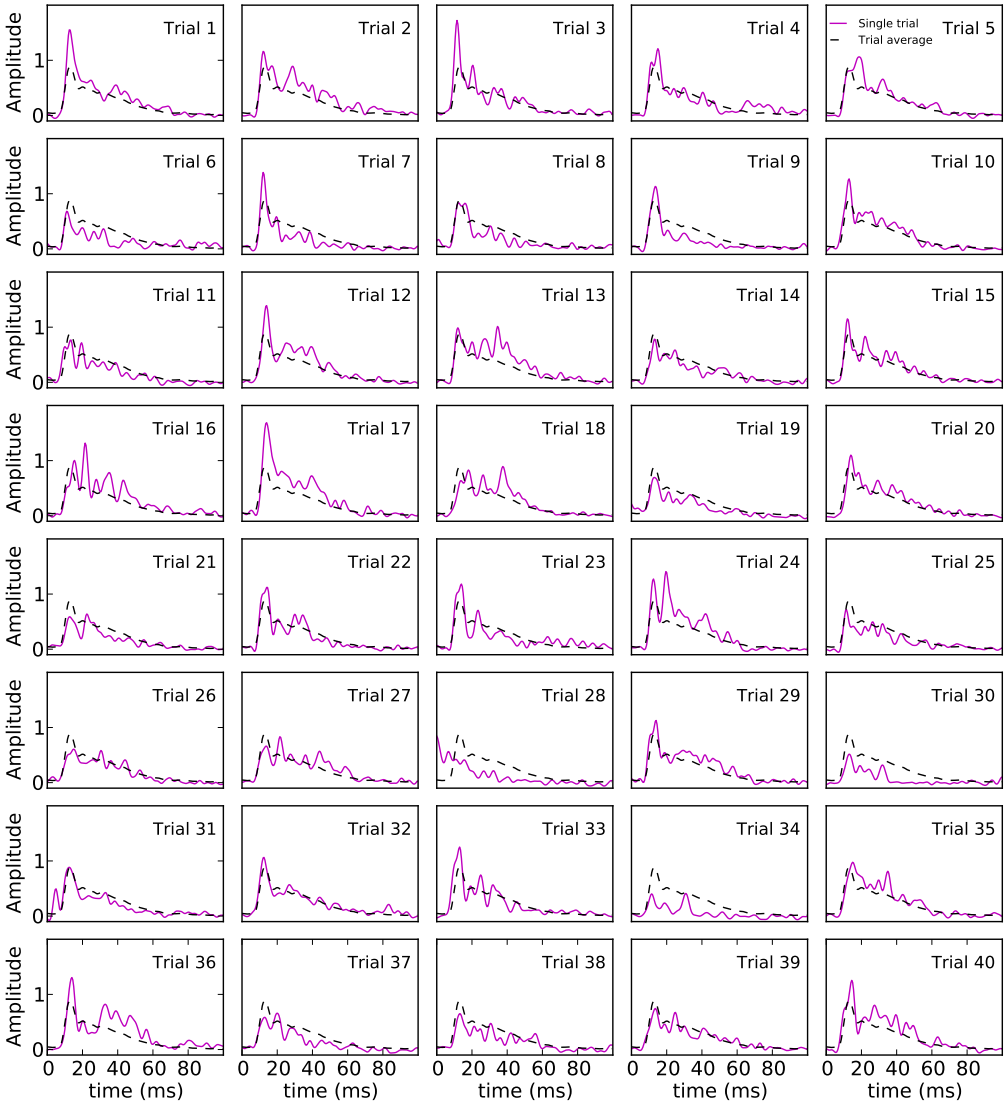

Figure A.19: Thalamic population firing rates.

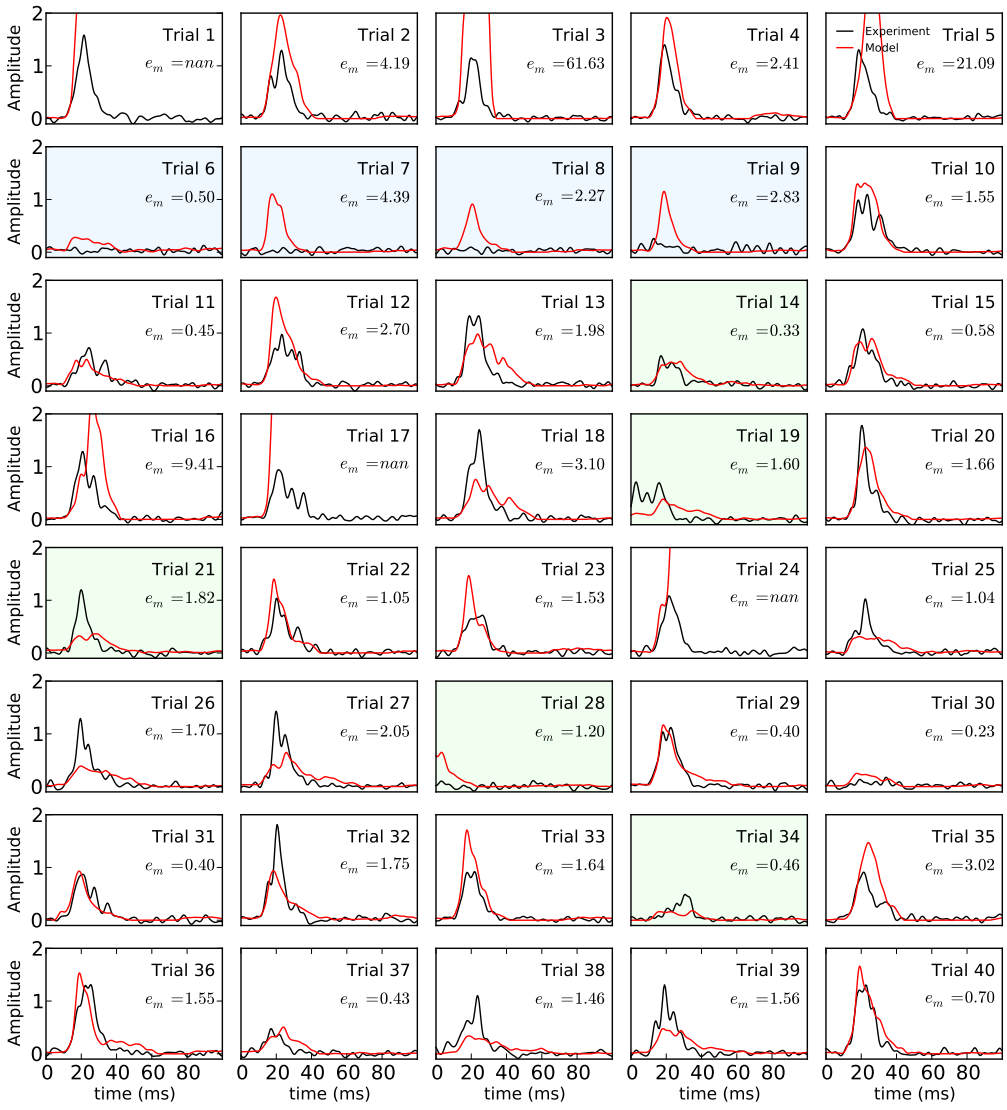

Figure A.20: Modeled layer 4 population firing rates.

A.11 Stimulus Condition 11

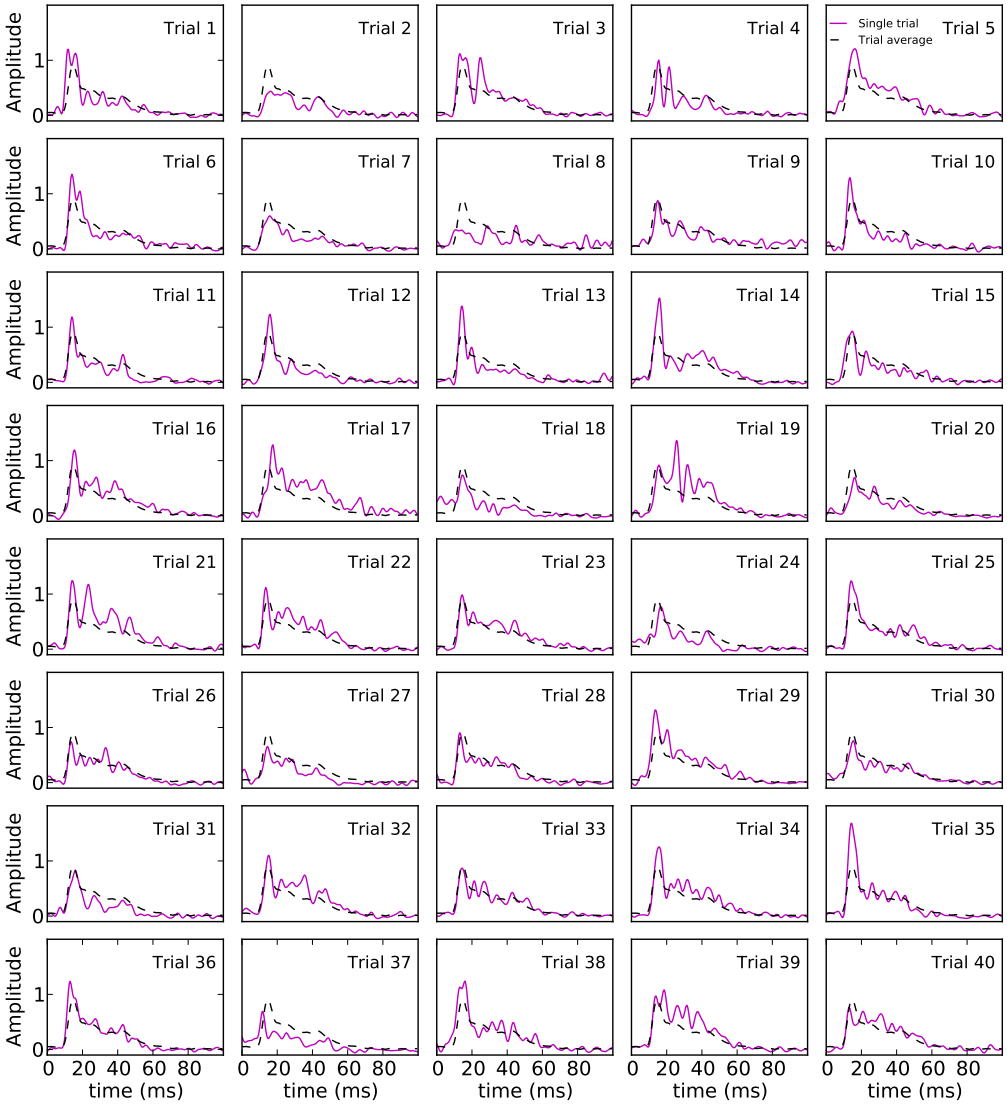

Figure A.21: Thalamic population firing rates.

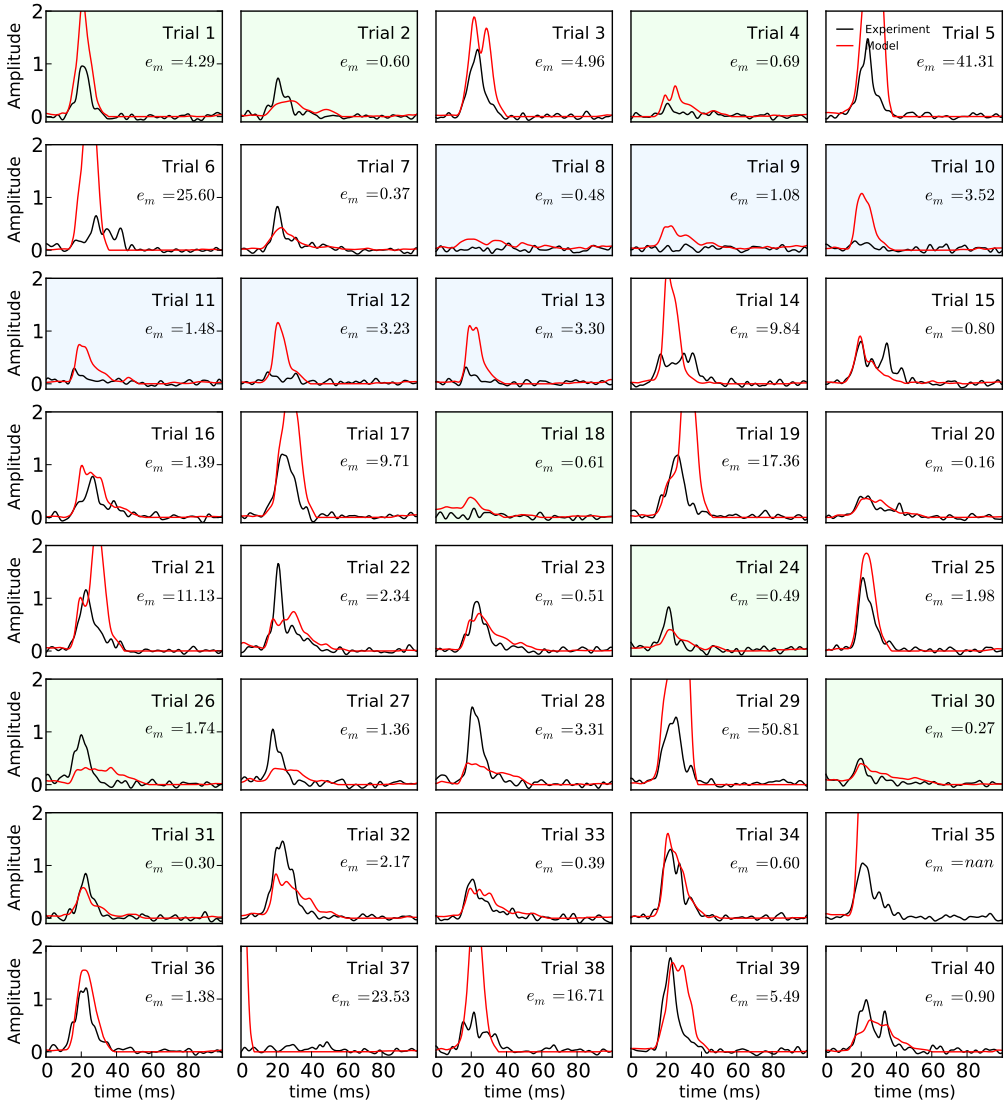

Figure A.22: Modeled layer 4 population firing rates.

A.12 Stimulus Condition 12

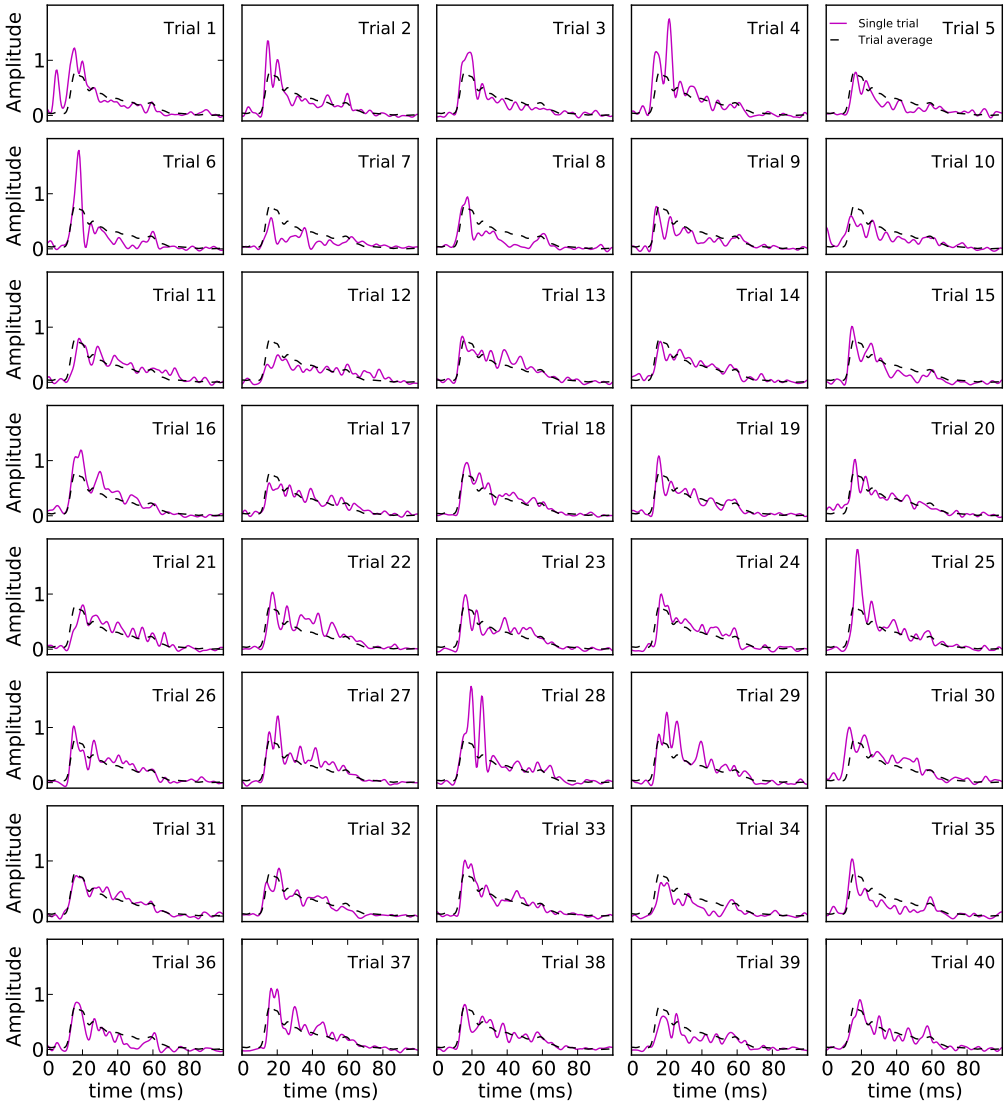

Figure A.23: Thalamic population firing rates.

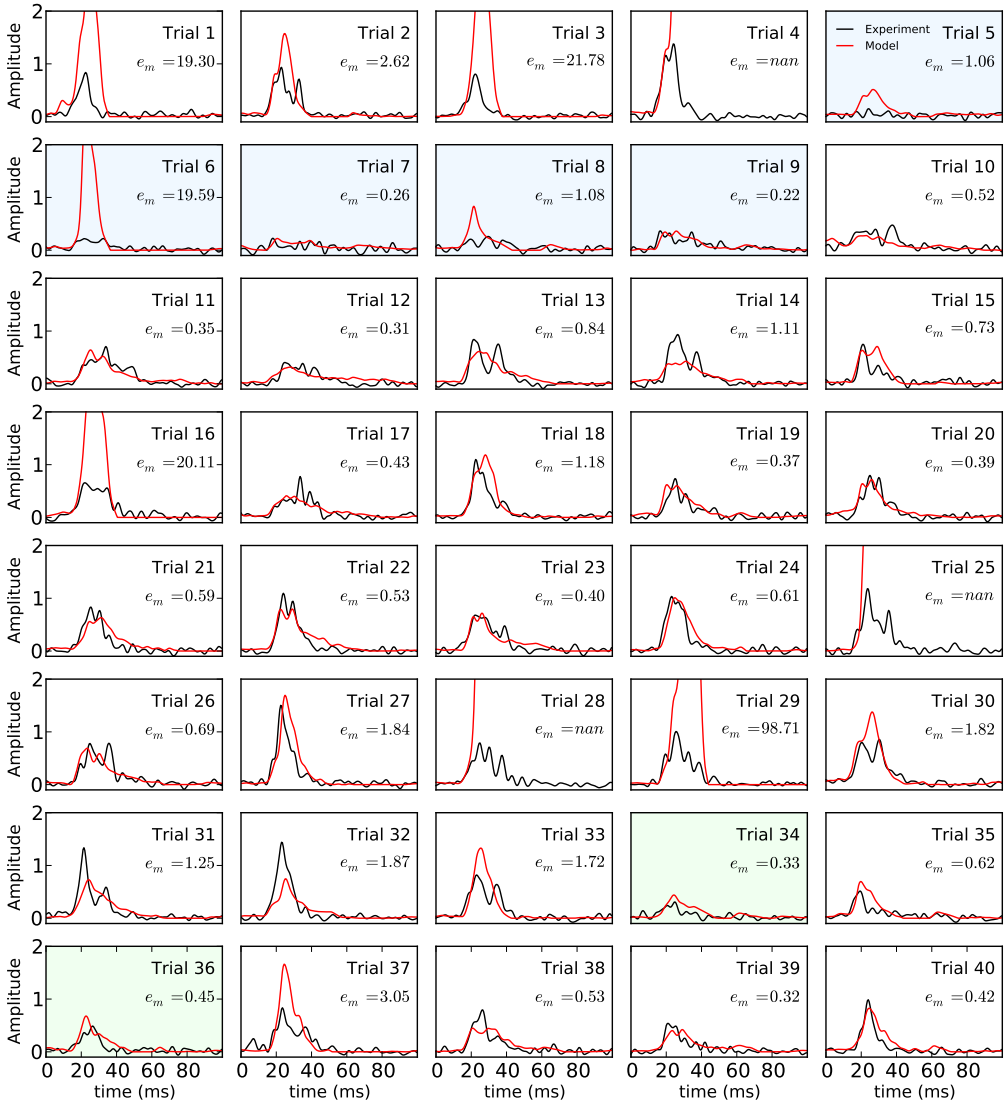

Figure A.24: Modeled layer 4 population firing rates.

A.13 Stimulus Condition 13

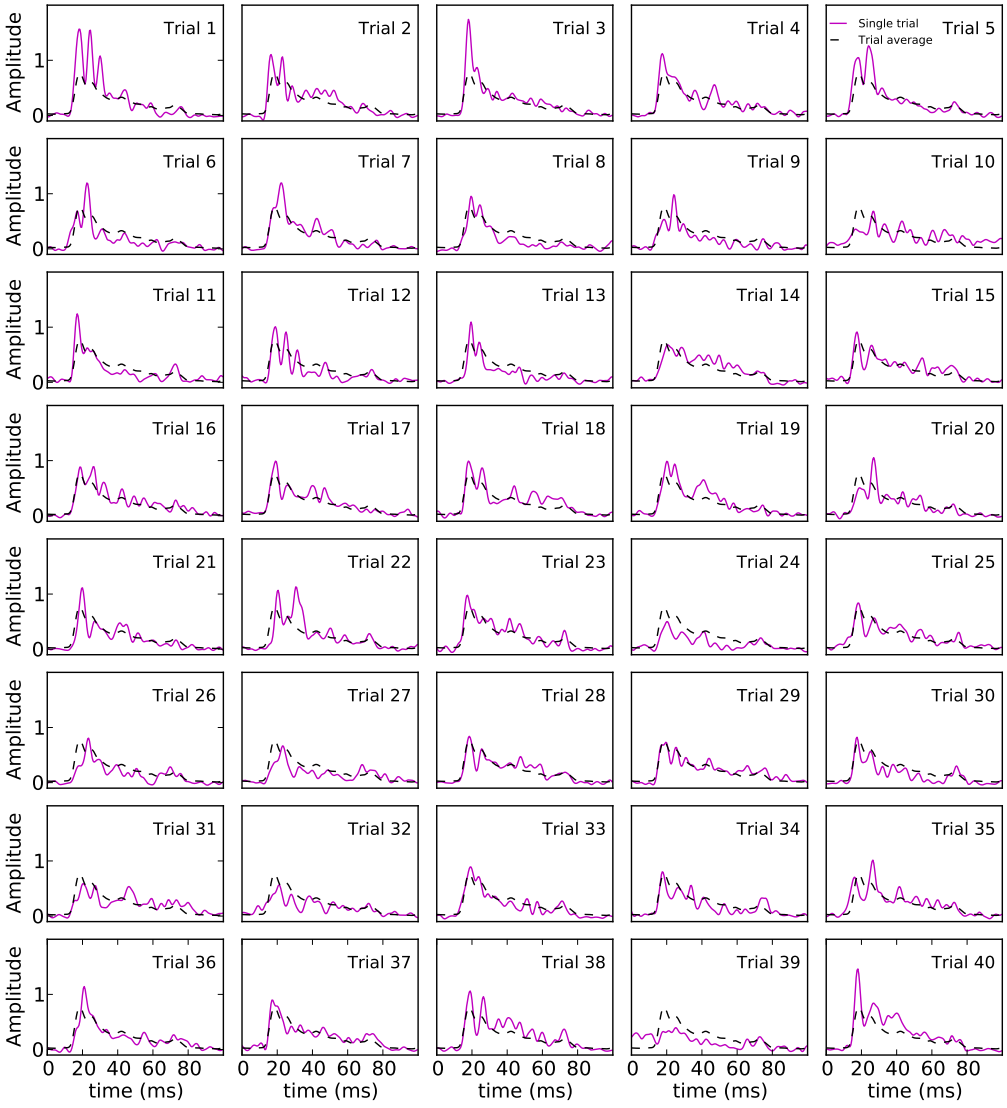

Figure A.25: Thalamic population firing rates.

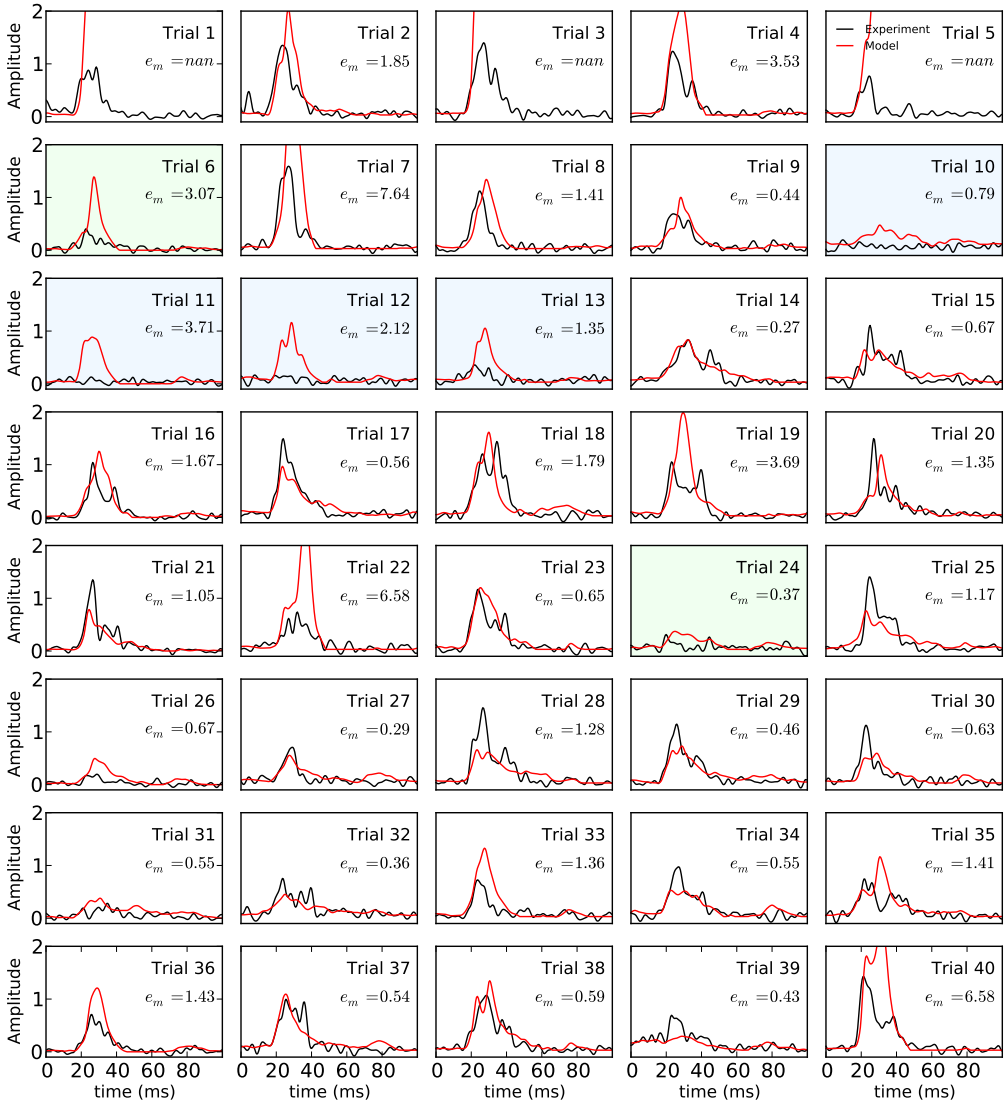

Figure A.26: Modeled layer 4 population firing rates.

A.14 Stimulus Condition 14

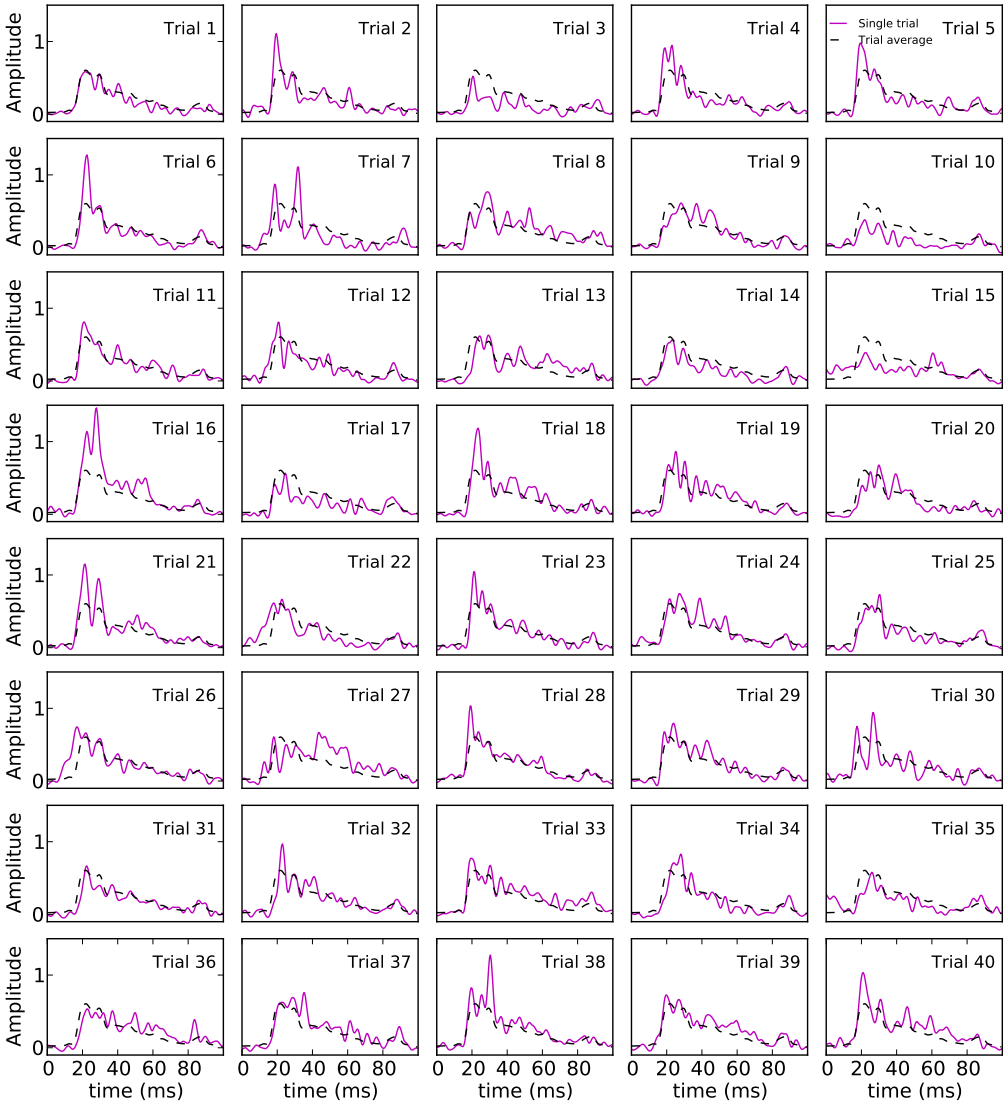

Figure A.27: Thalamic population firing rates.

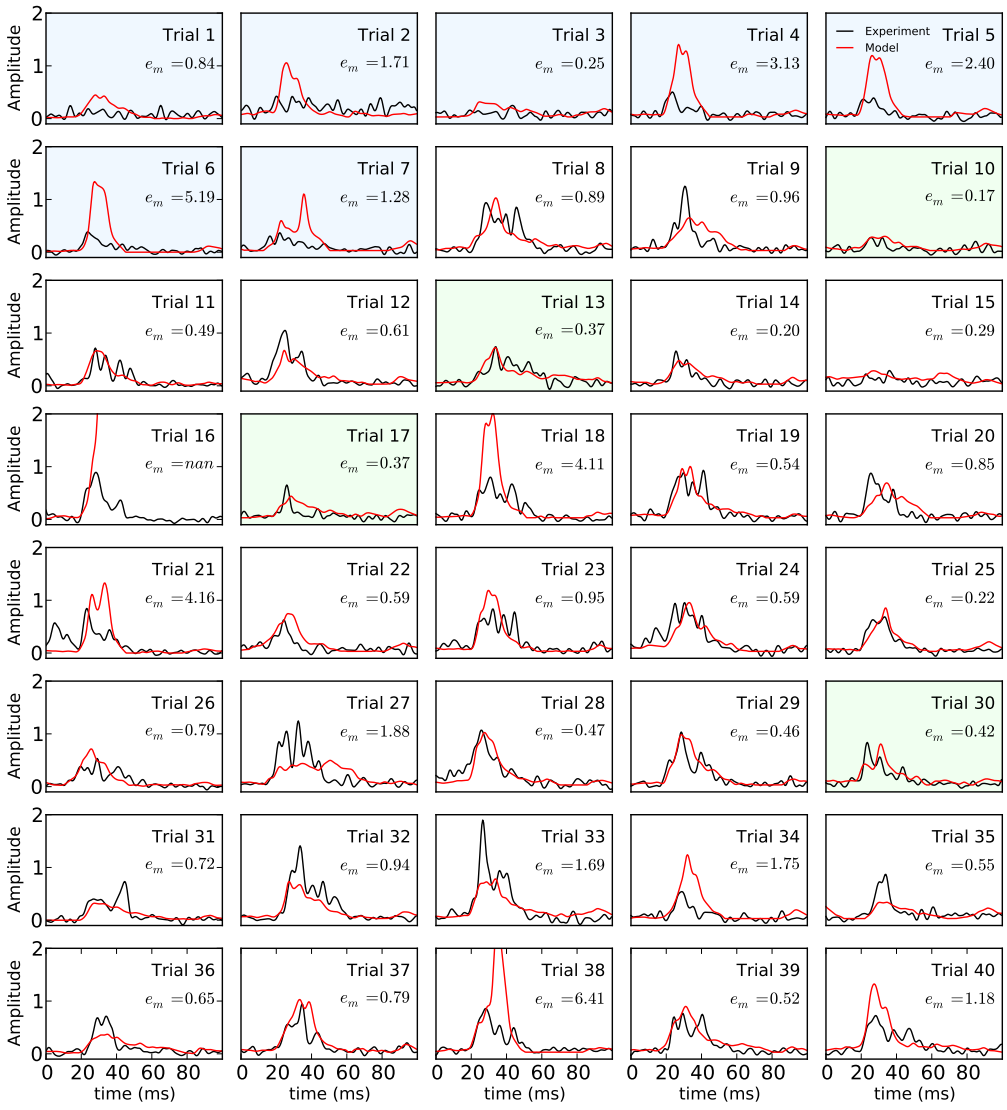

Figure A.28: Modeled layer 4 population firing rates.

A.15 Stimulus Condition 15

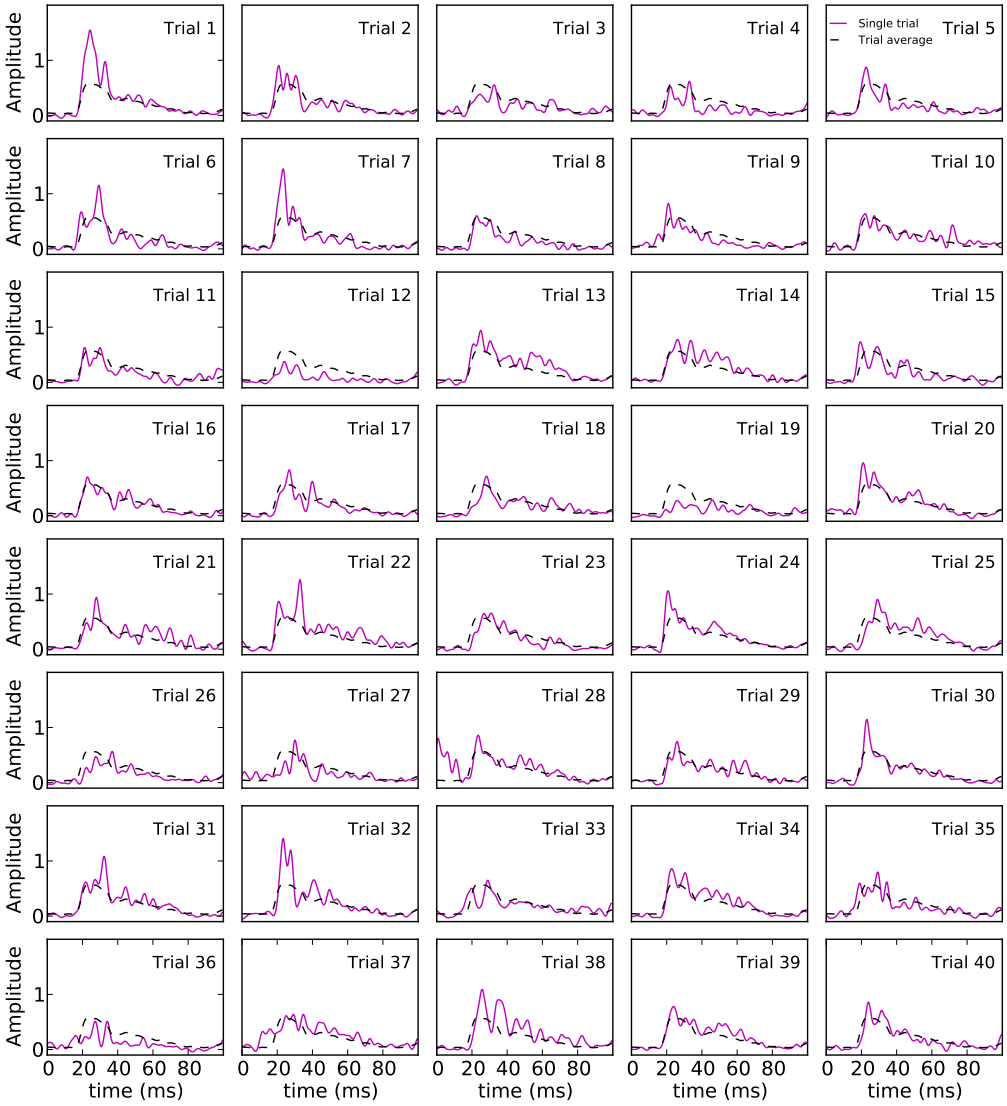

Figure A.29: Thalamic population firing rates.

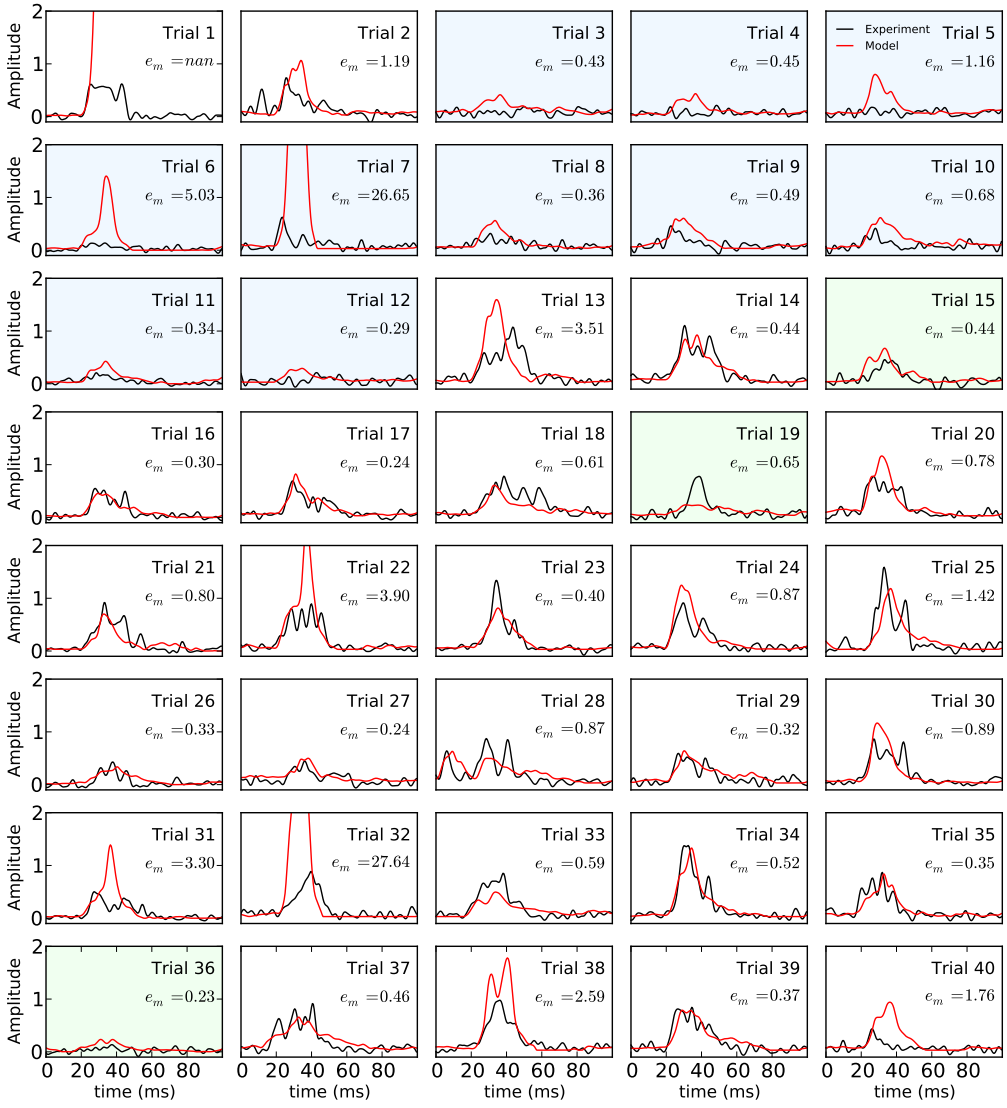

Figure A.30: Modeled layer 4 population firing rates.

A.16 Stimulus Condition 16

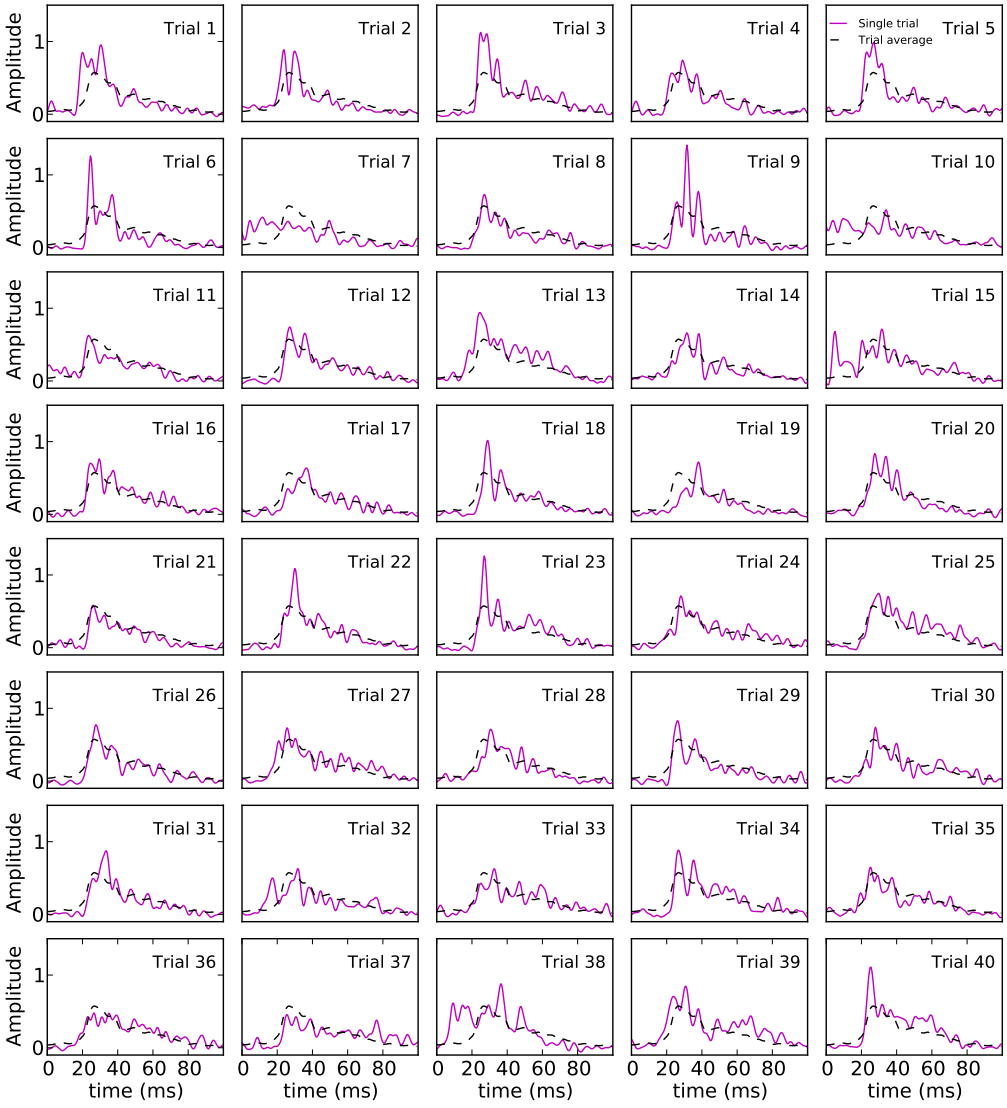

Figure A.31: Thalamic population firing rates.

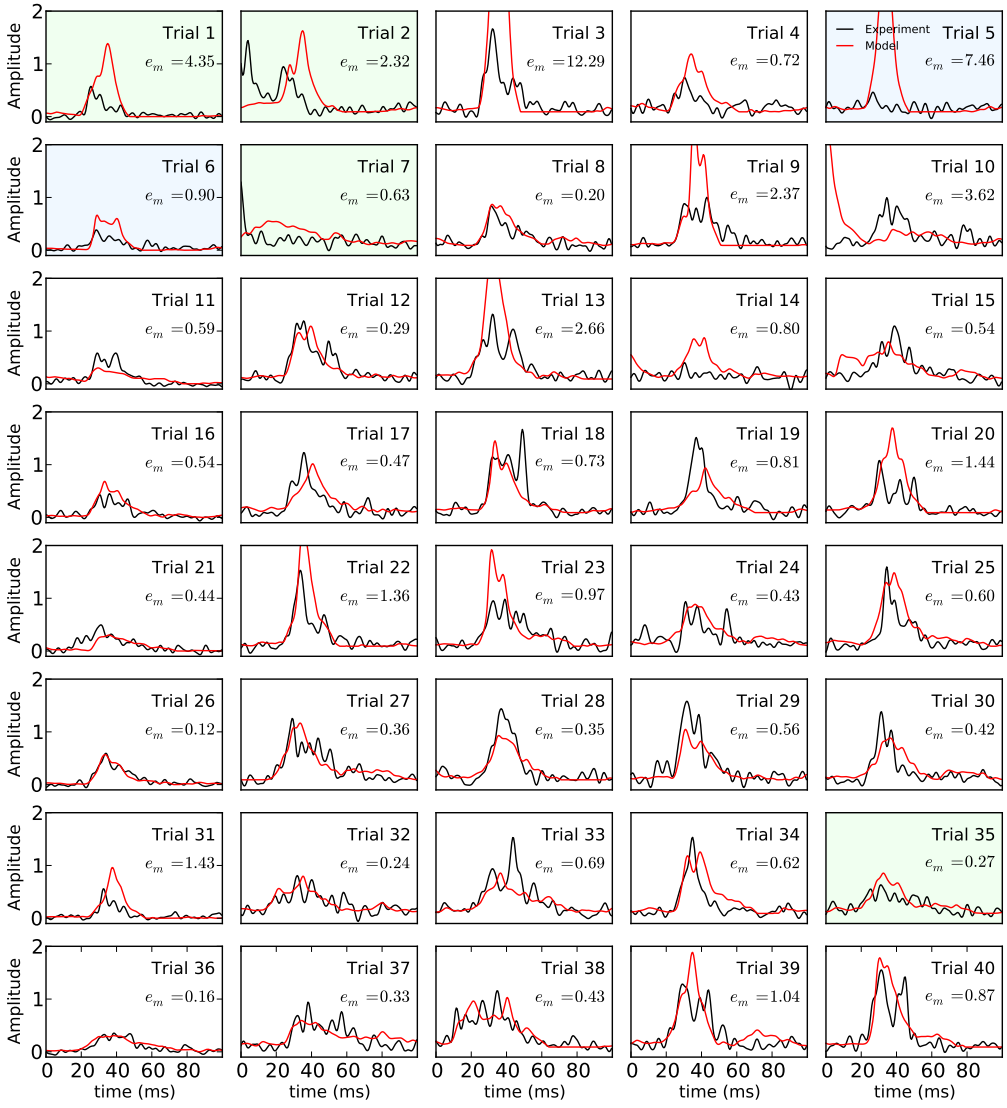

Figure A.32: Modeled layer 4 population firing rates.

A.17 Stimulus Condition 17

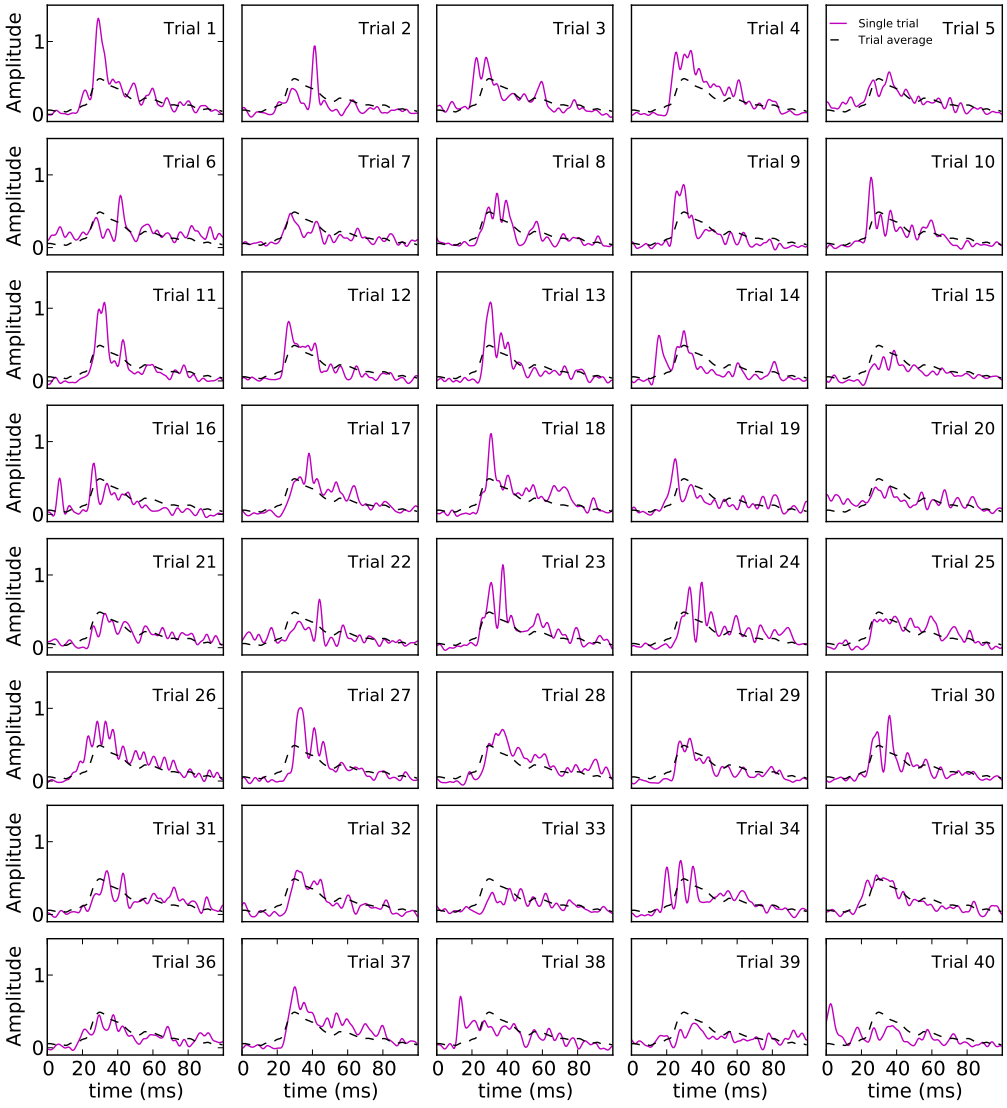

Figure A.33: Thalamic population firing rates.

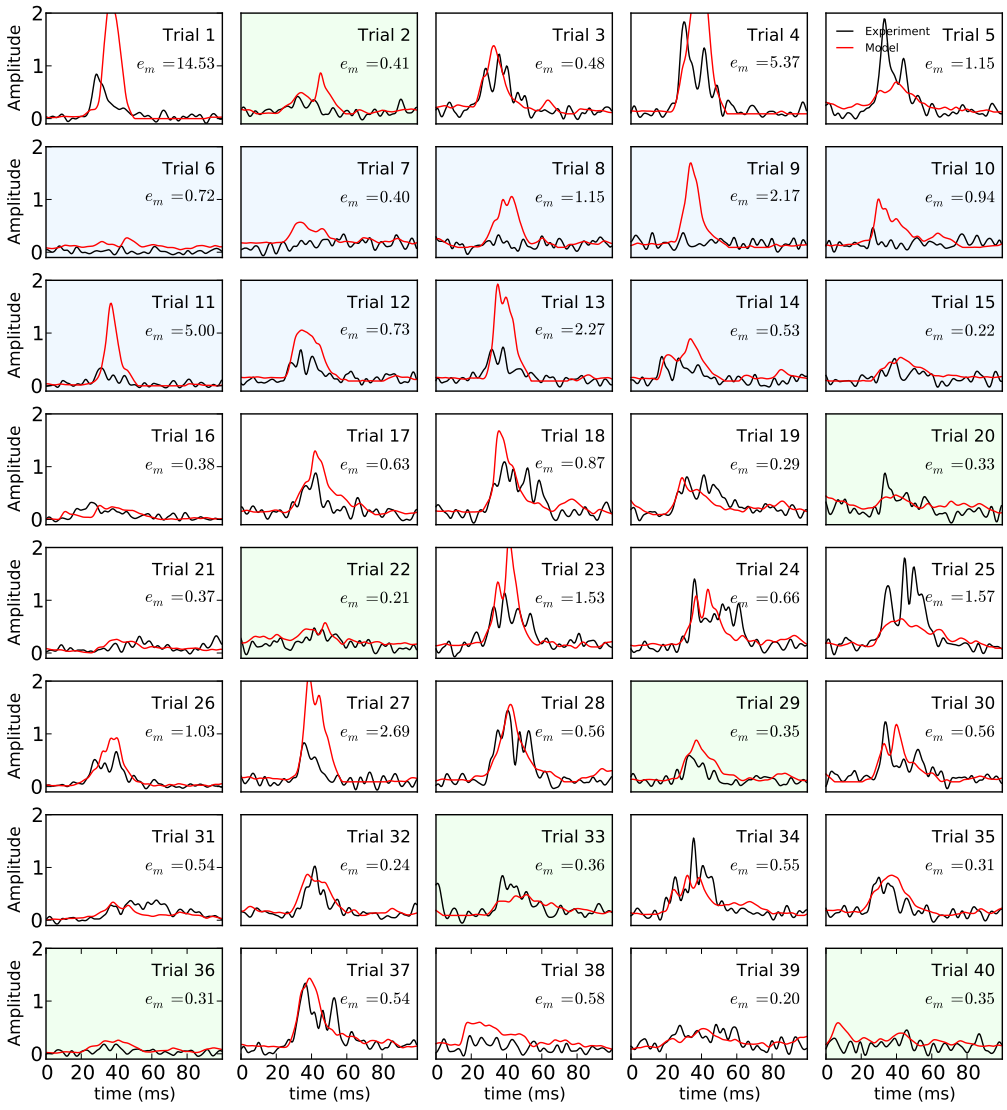

Figure A.34: Modeled layer 4 population firing rates.

A.18 Stimulus Condition 18

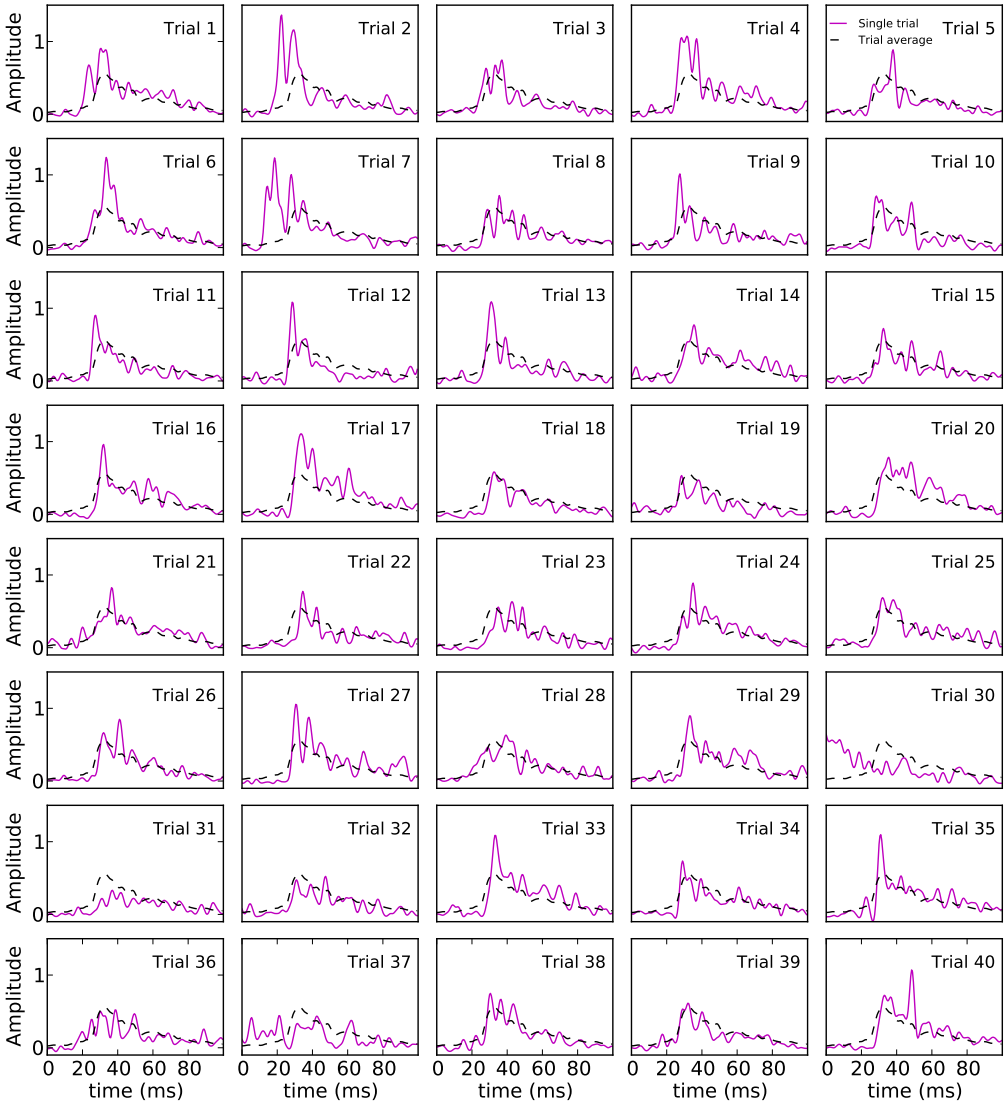

Figure A.35: Thalamic population firing rates.

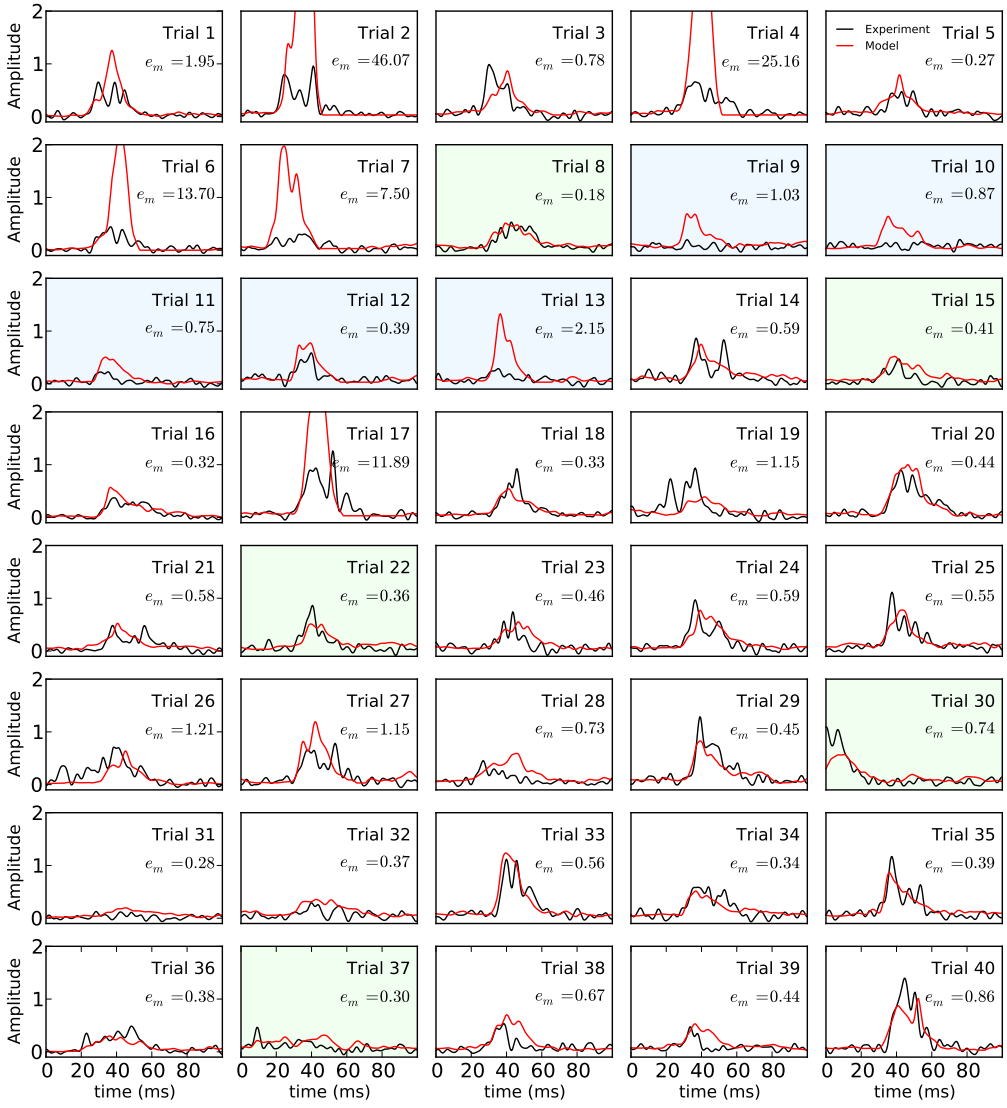

Figure A.36: Modeled layer 4 population firing rates.

A.19 Stimulus Condition 19

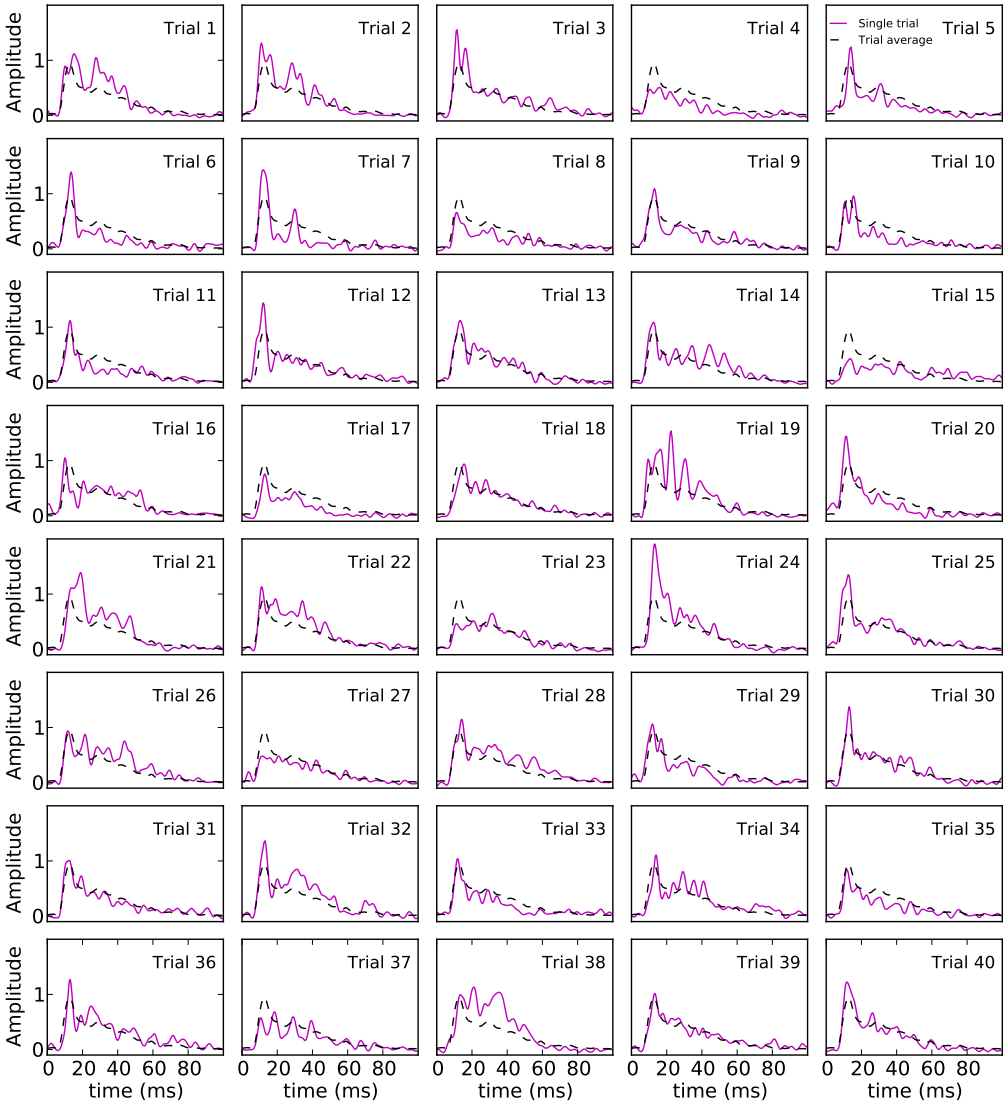

Figure A.37: Thalamic population firing rates.

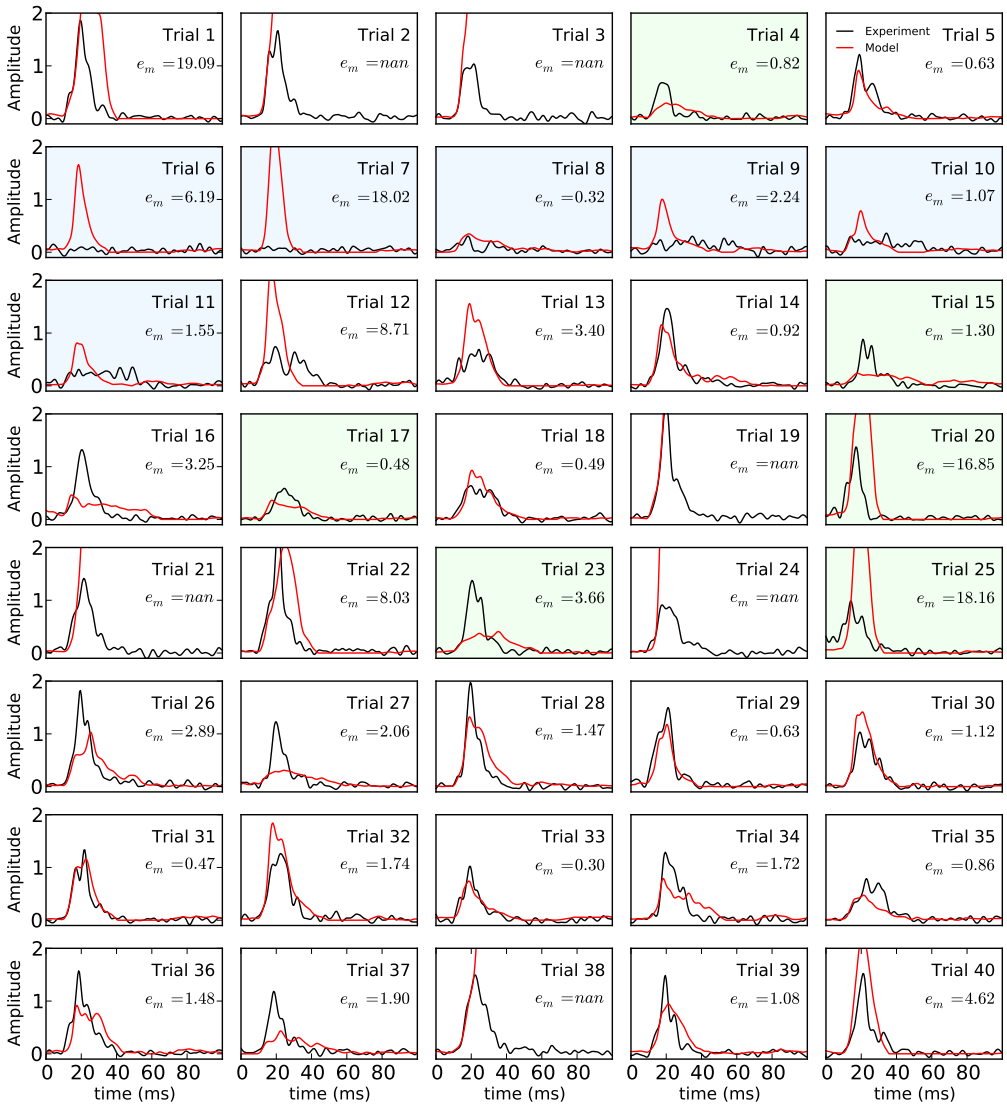

Figure A.38: Modeled layer 4 population firing rates.

A.20 Stimulus Condition 20

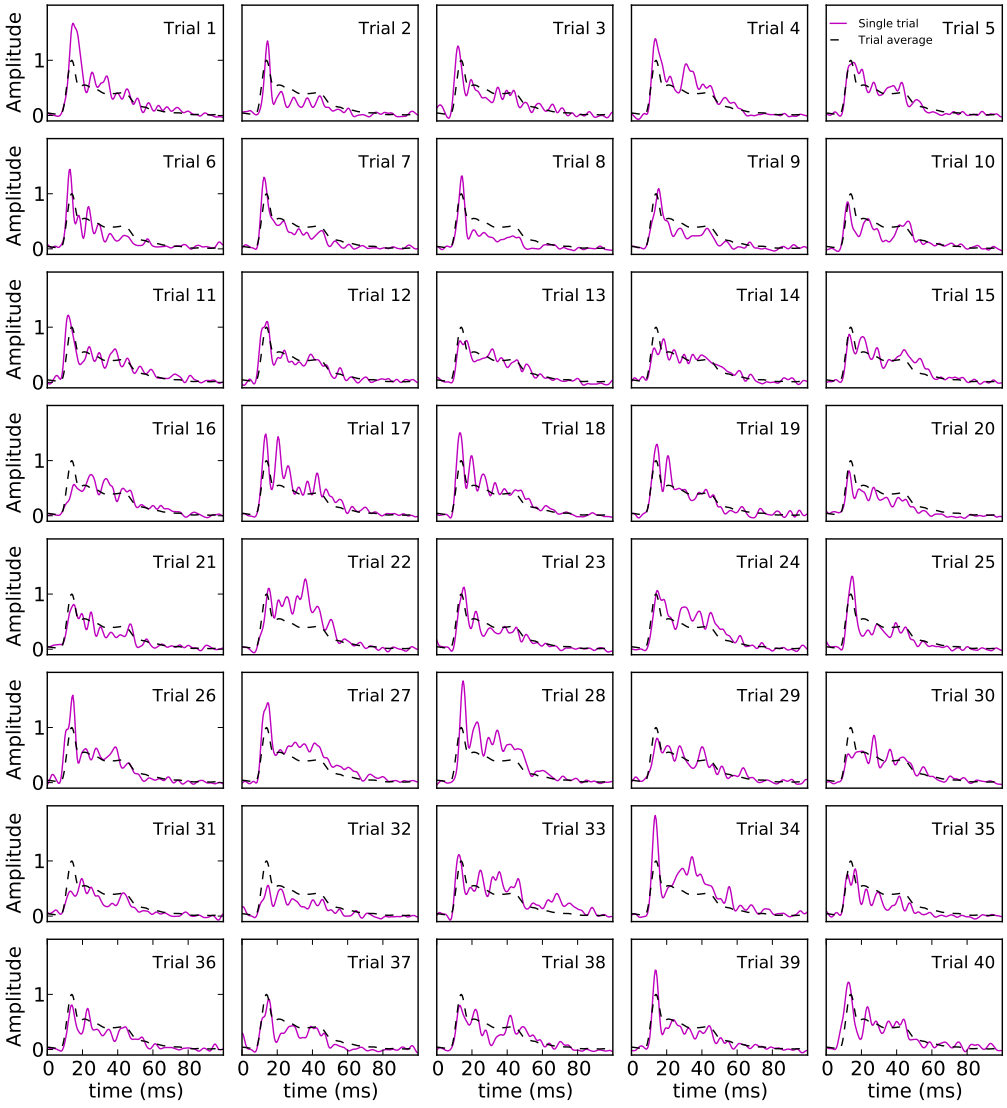

Figure A.39: Thalamic population firing rates.

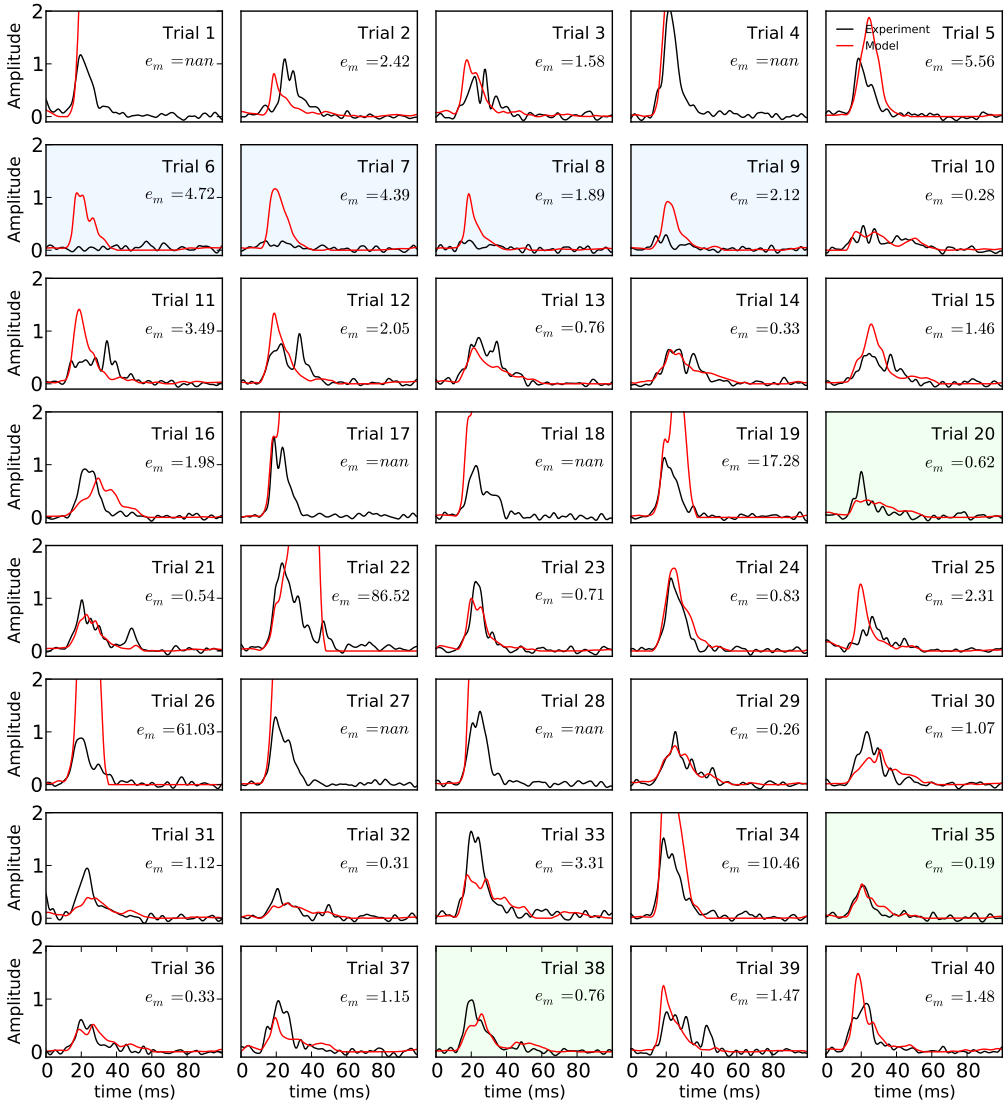

Figure A.40: Modeled layer 4 population firing rates.

A.21 Stimulus Condition 21

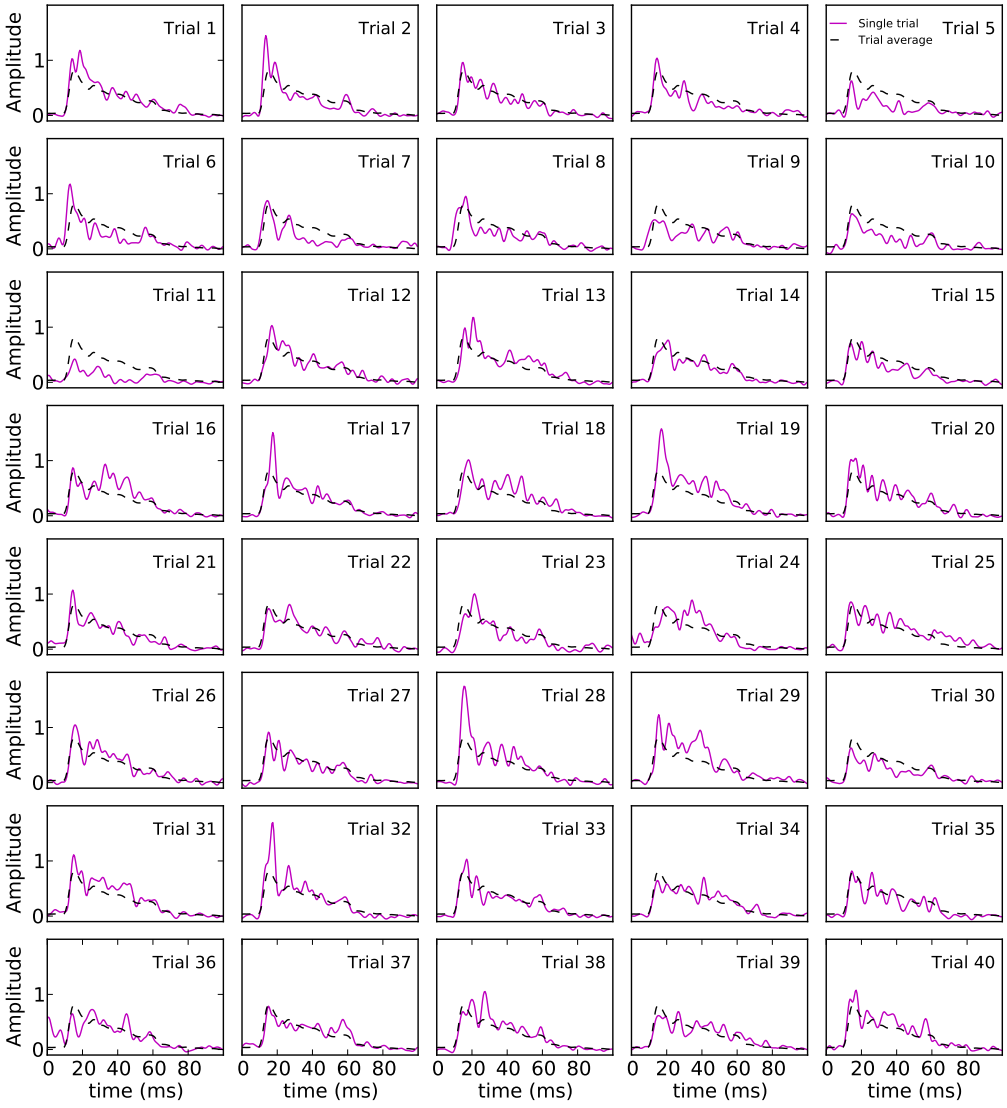

Figure A.41: Thalamic population firing rates.

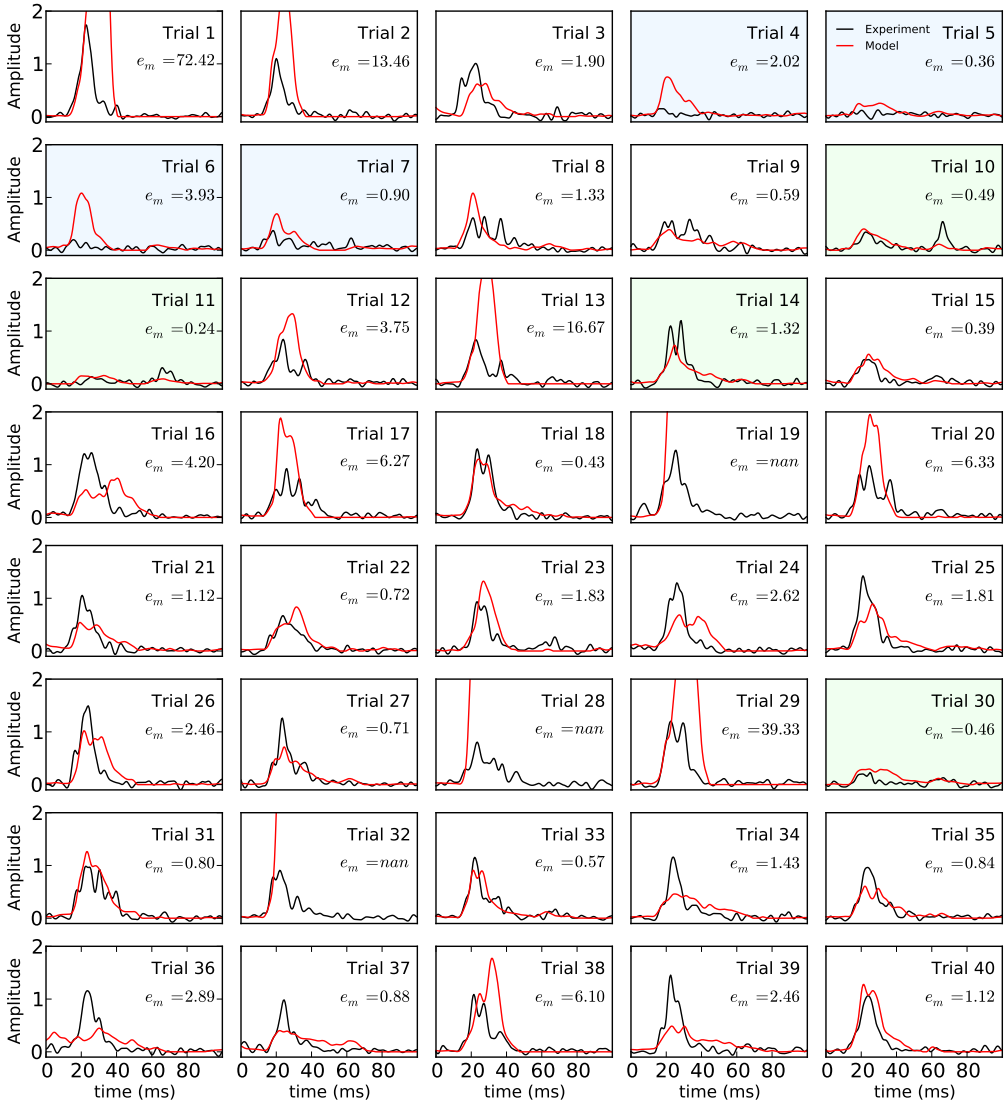

Figure A.42: Modeled layer 4 population firing rates.

A.22 Stimulus Condition 22

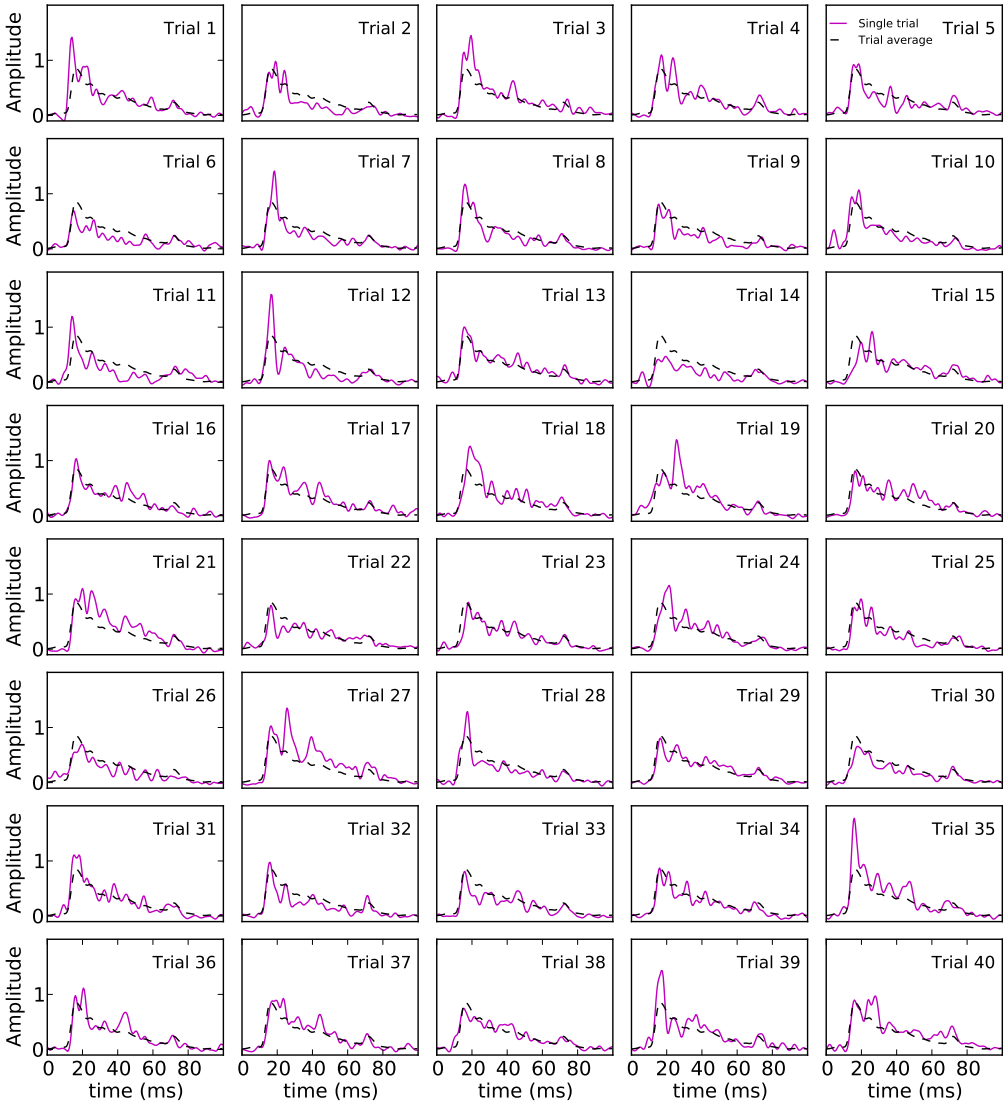

Figure A.43: Thalamic population firing rates.

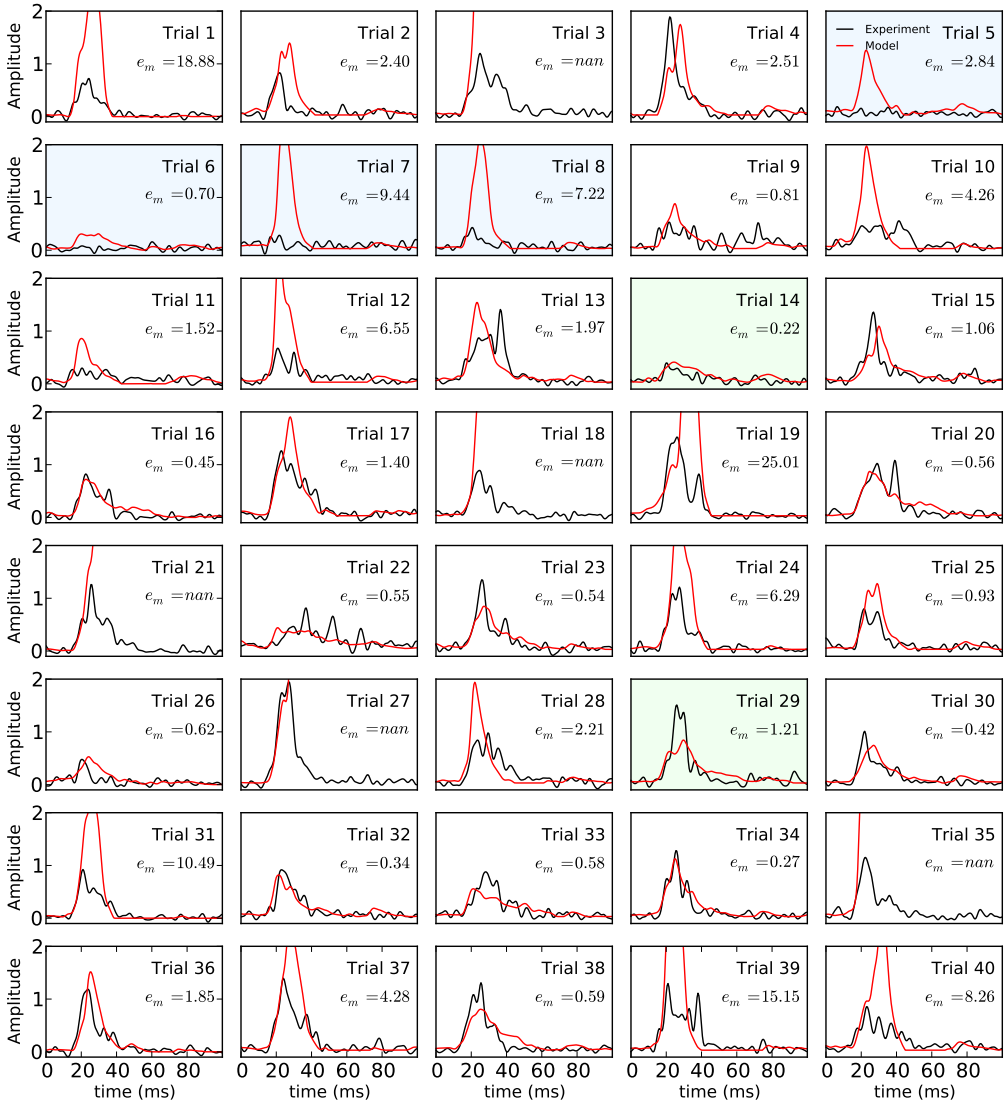

Figure A.44: Modeled layer 4 population firing rates.

A.23 Stimulus Condition 23

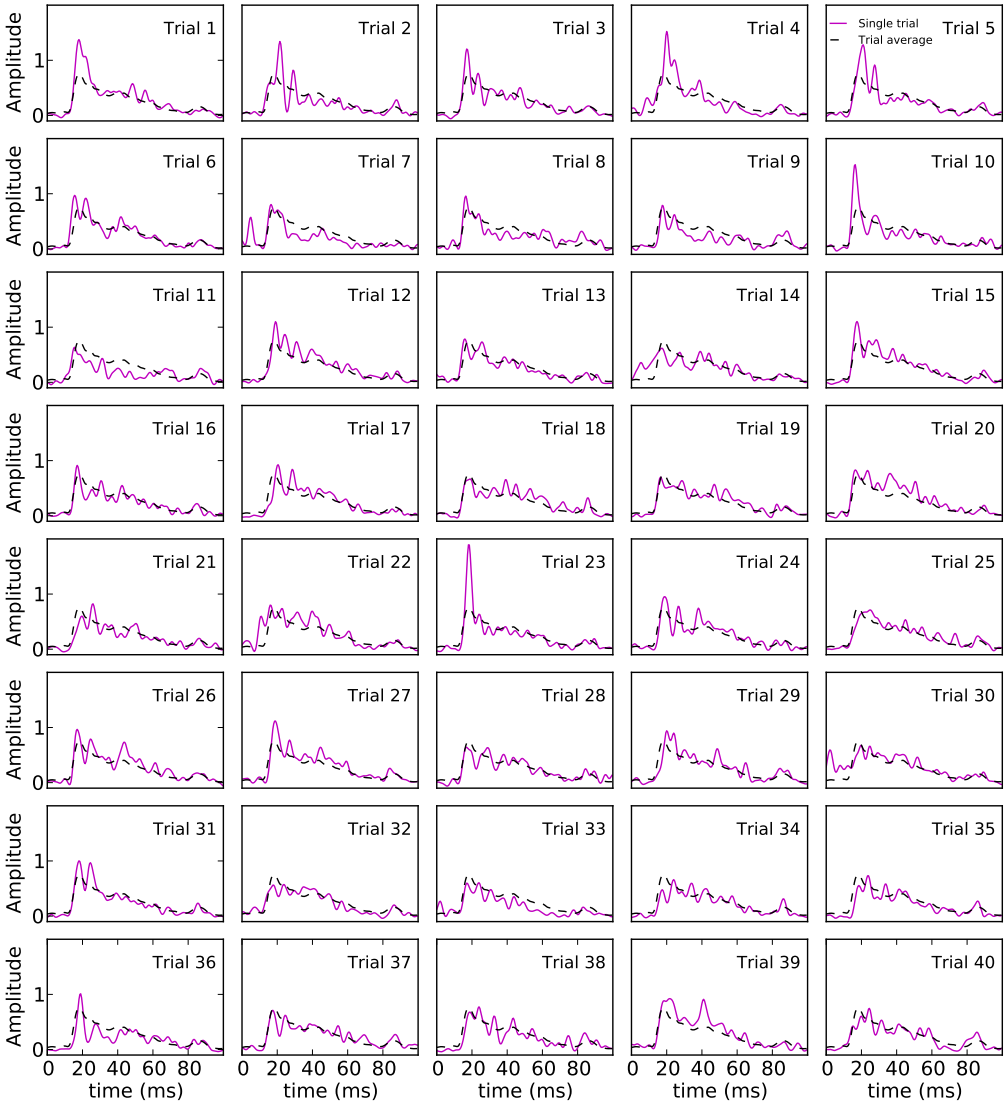

Figure A.45: Thalamic population firing rates.

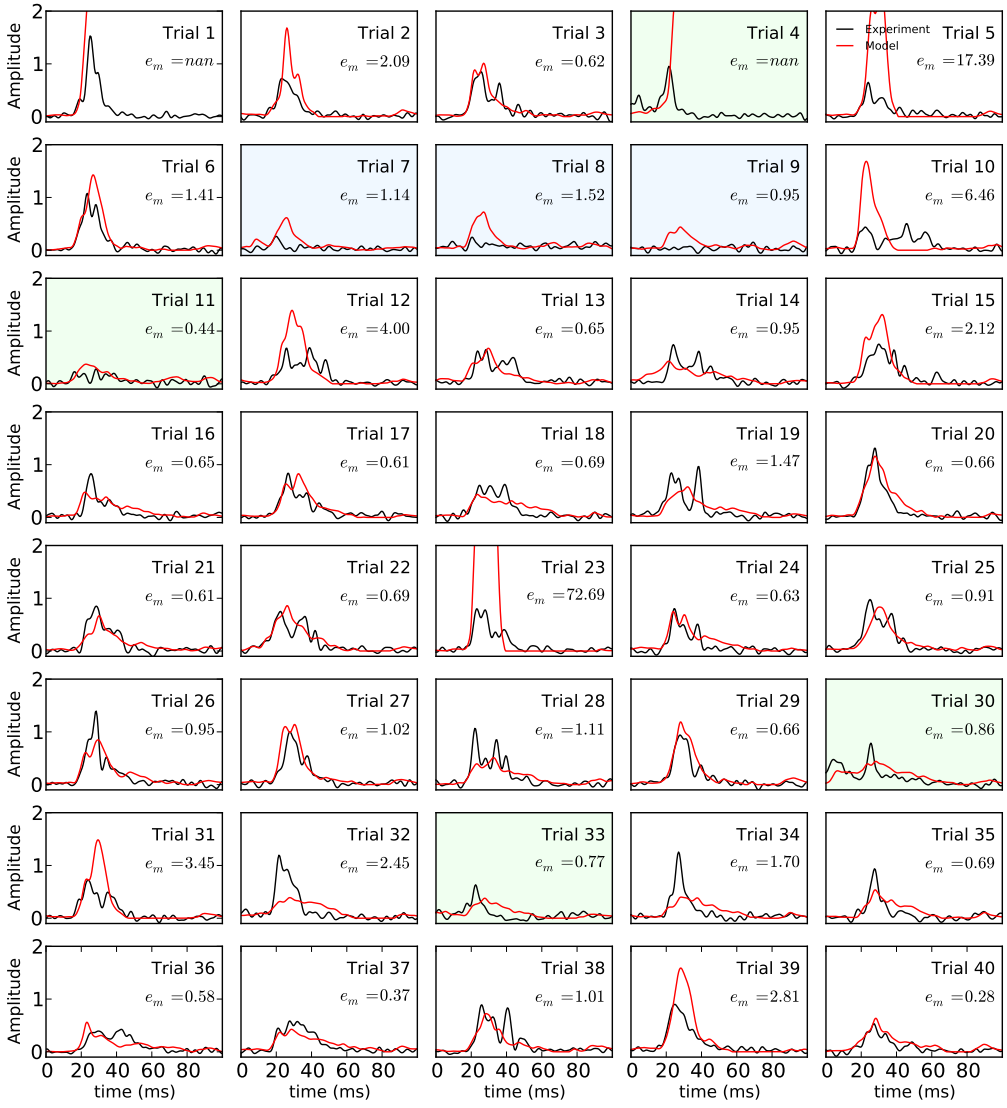

Figure A.46: Modeled layer 4 population firing rates.

A.24 Stimulus Condition 24

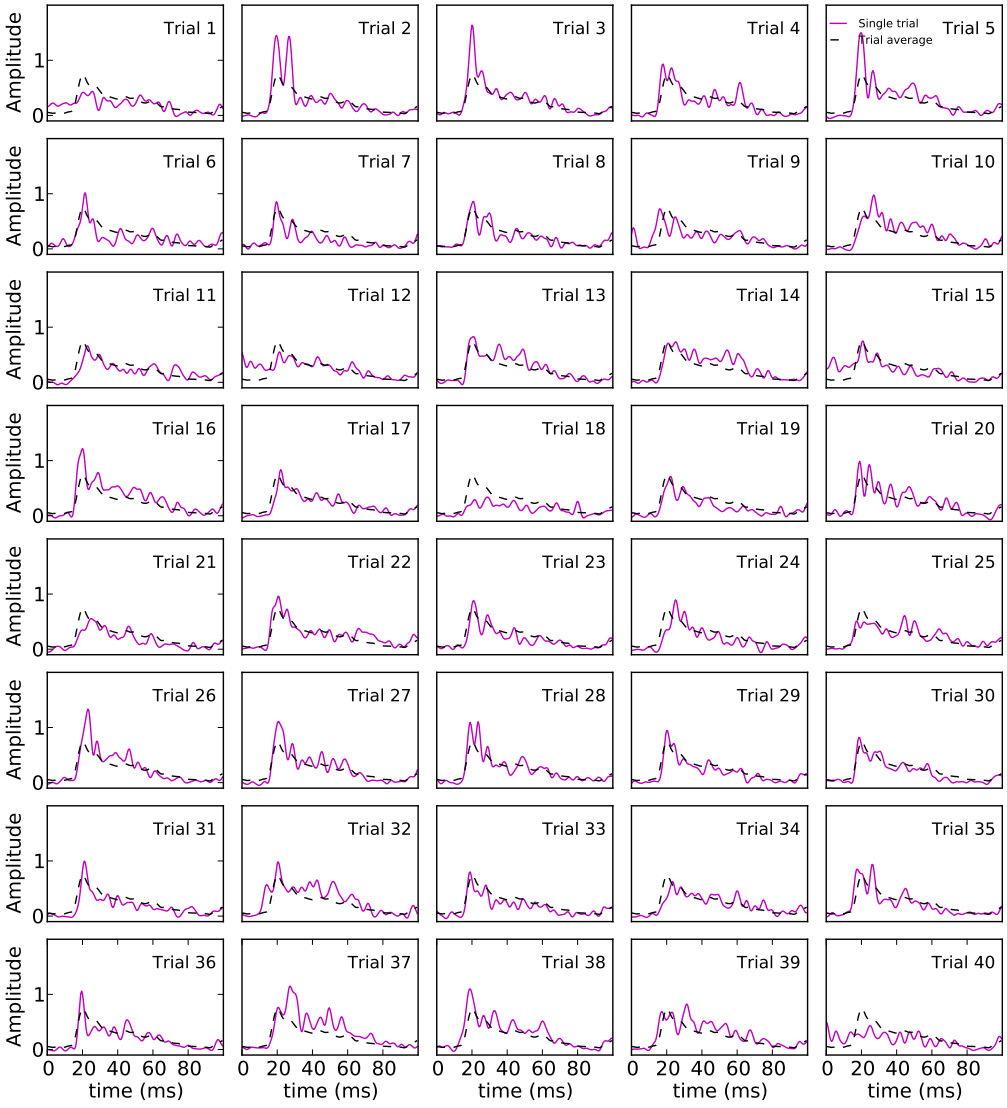

Figure A.47: Thalamic population firing rates.

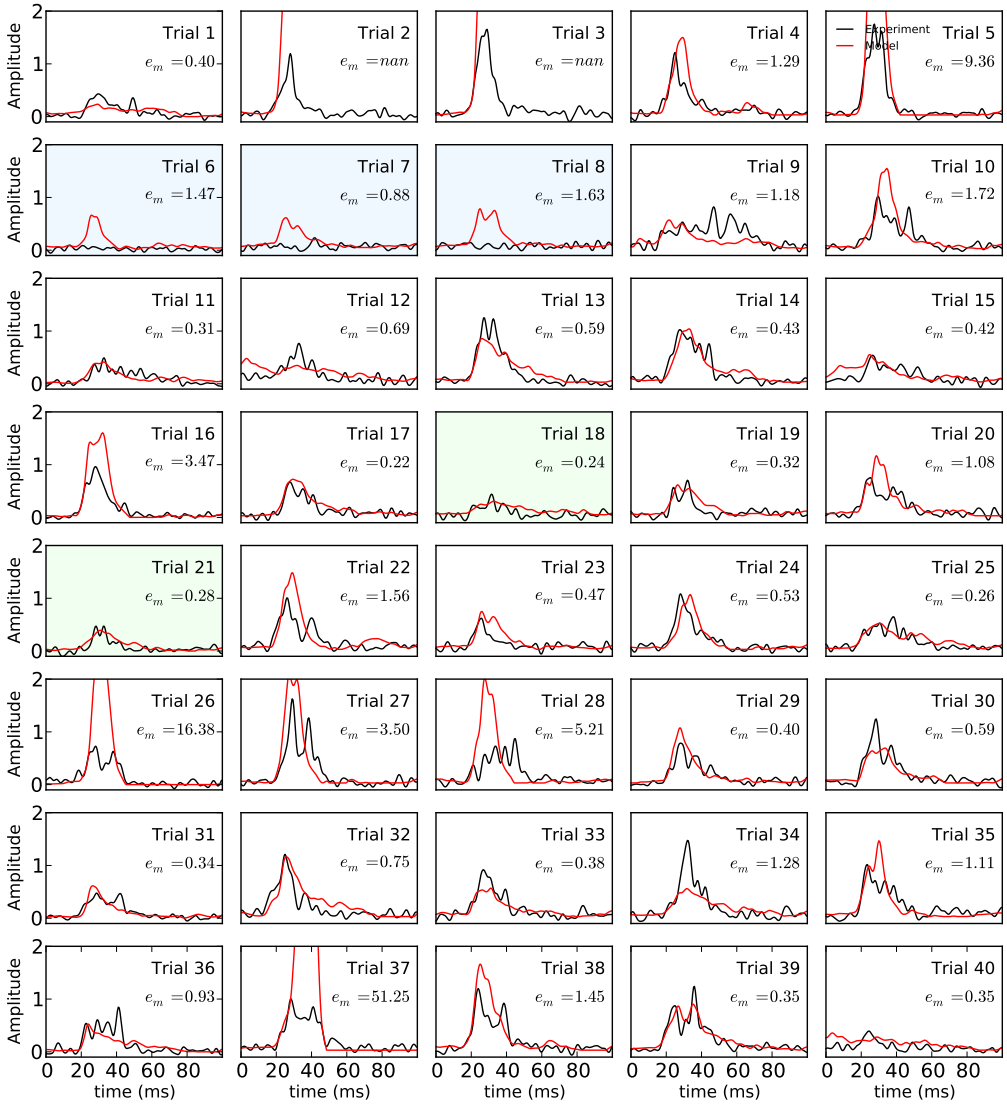

Figure A.48: Modeled layer 4 population firing rates.

A.25 Stimulus Condition 25

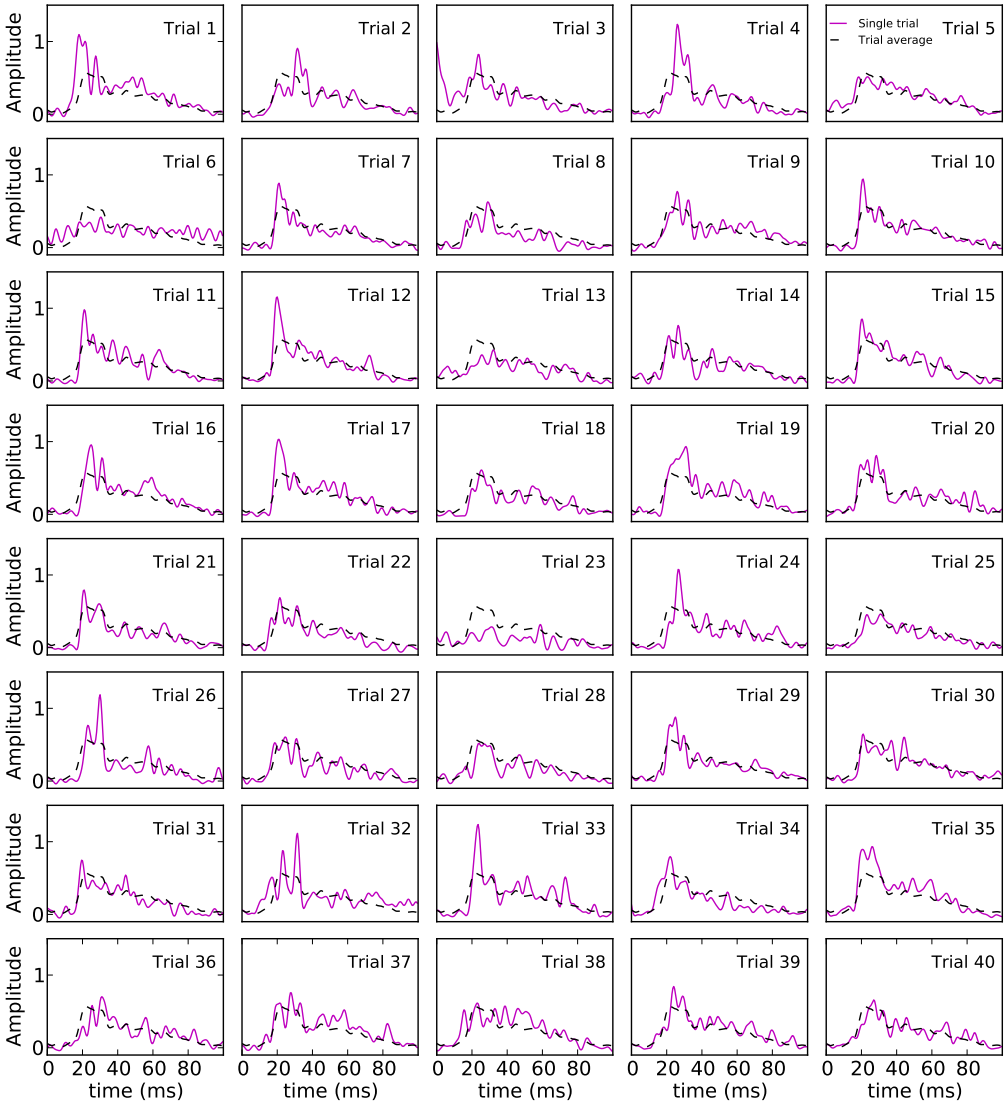

Figure A.49: Thalamic population firing rates.

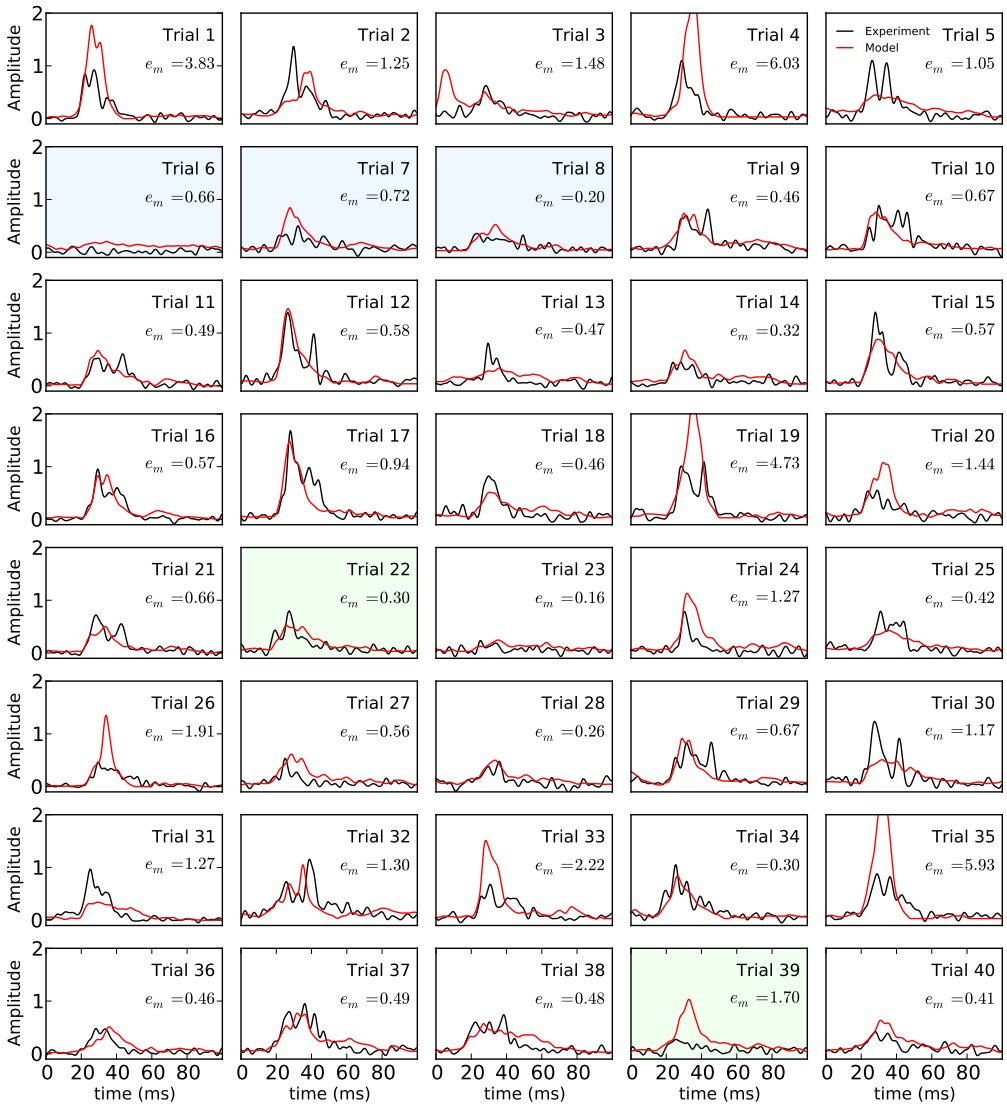

Figure A.50: Modeled layer 4 population firing rates.

A.26 Stimulus Condition 26

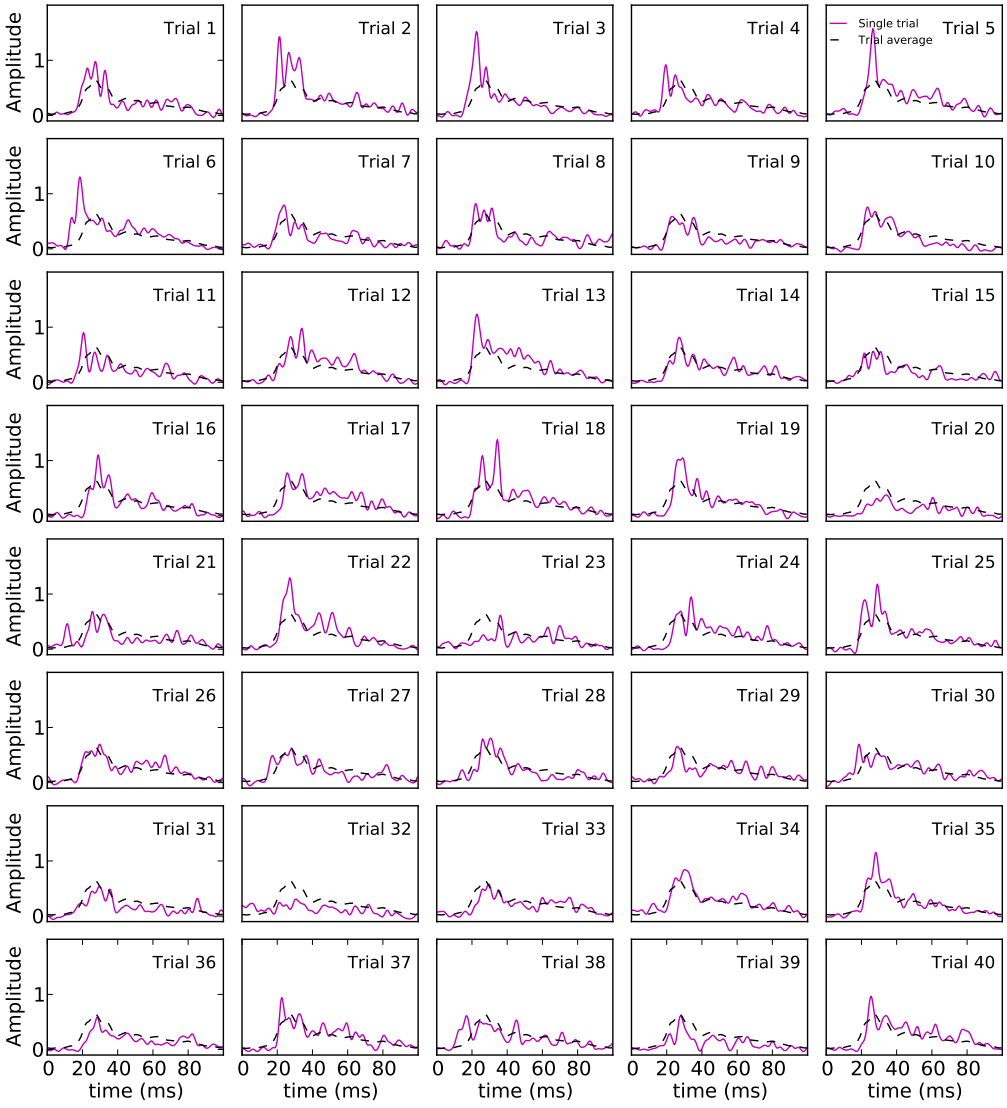

Figure A.51: Thalamic population firing rates.

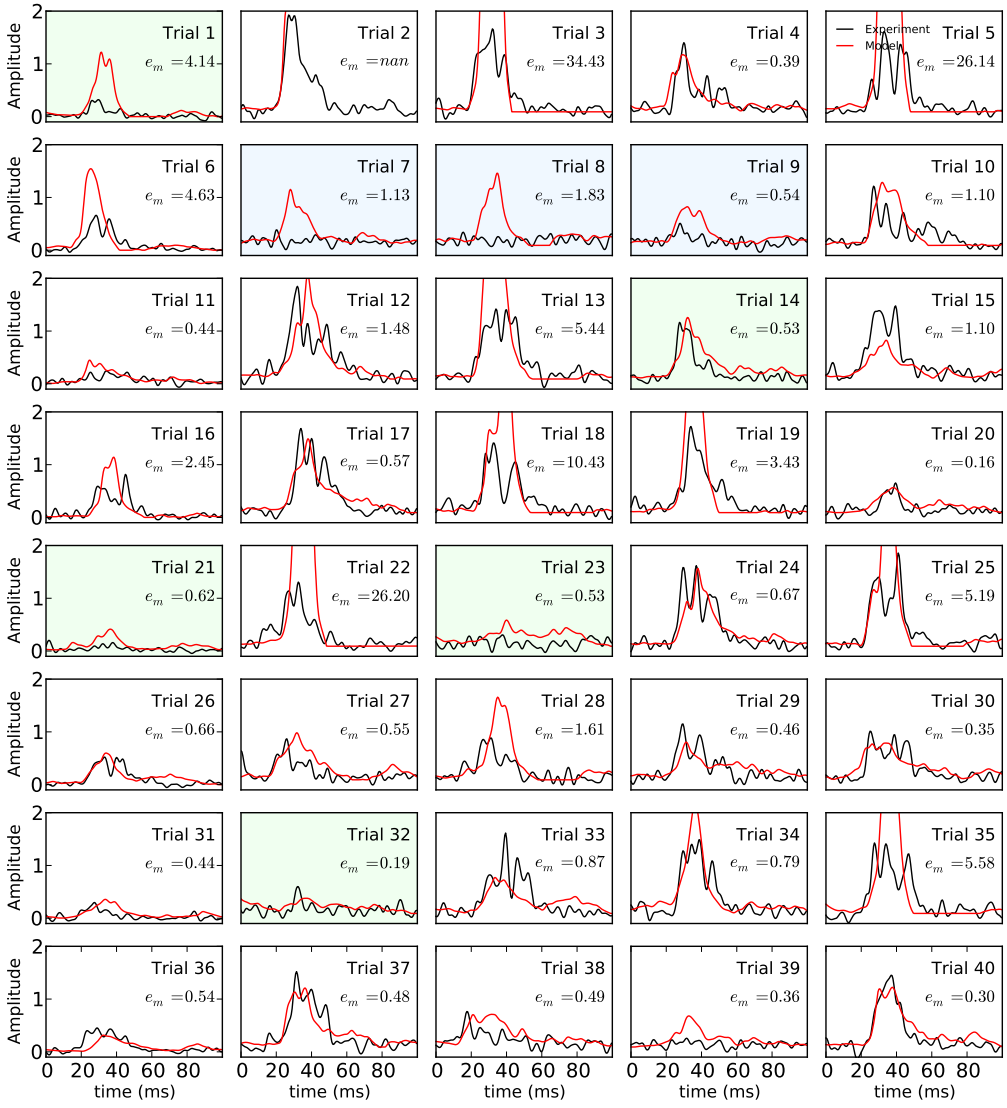

Figure A.52: Modeled layer 4 population firing rates.

A.27 Stimulus Condition 27

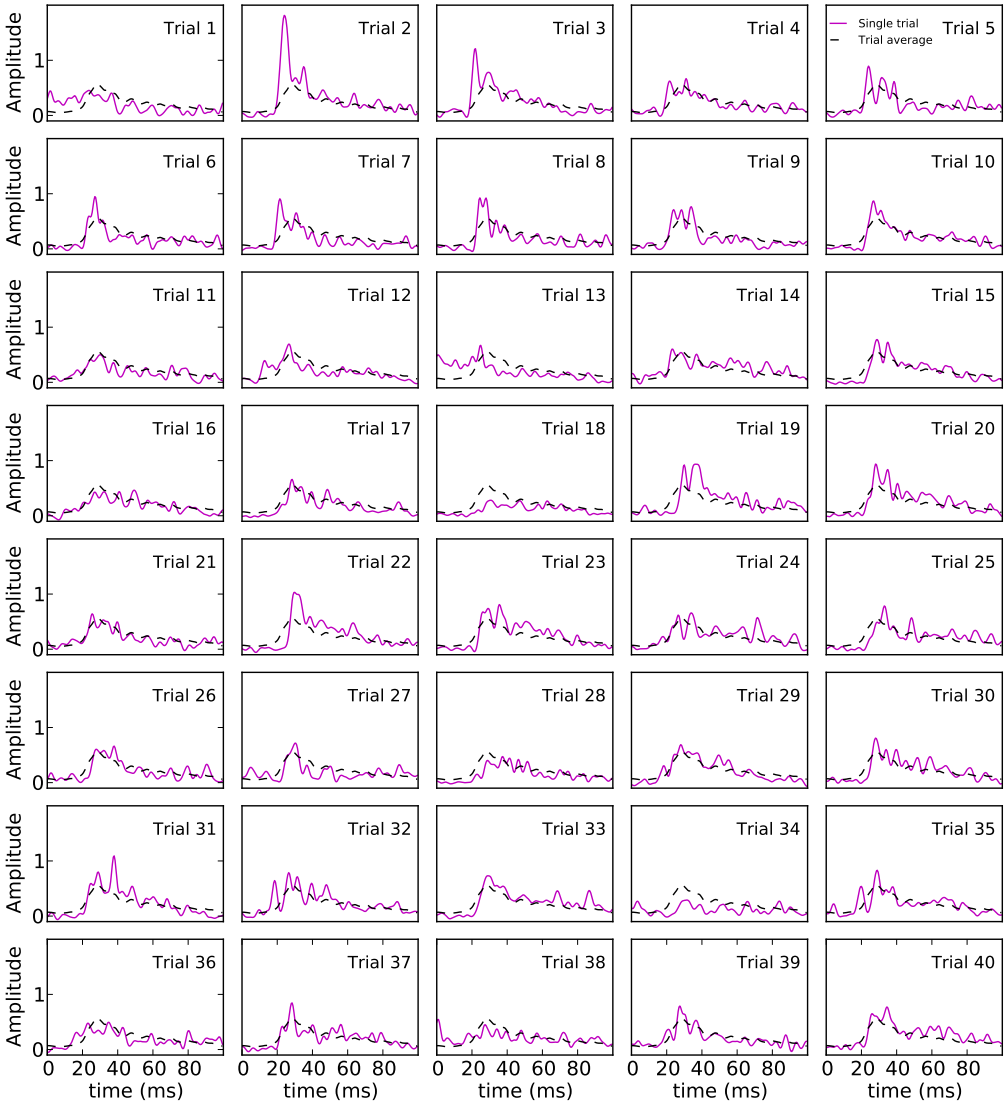

Figure A.53: Thalamic population firing rates.

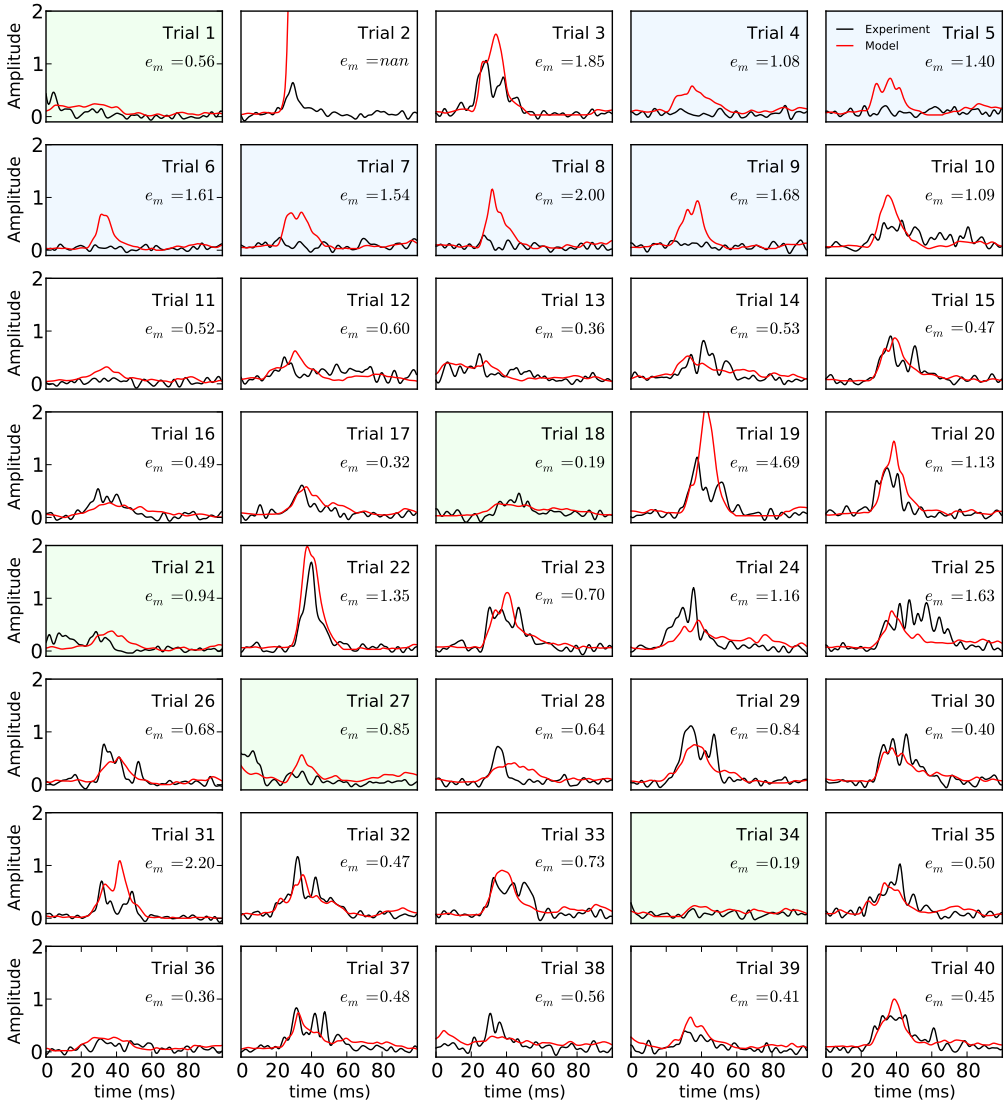

Figure A.54: Modeled layer 4 population firing rates.

Appendix B

Python scripts

B.1 Implementation of Thalamocortical Model

```
"""
Implements the Thalamocortical Model described in
Blomquist et. al []

Solves input of different time resolutions.
"""

__author__ = "Eivind Hennestad"
__email__ = "ehennestad@gmail.com"


import numpy as np
import matplotlib.pyplot as plt
import ODESolver

import os


param = { 'tau_ef': 3.7,
          'delta_ef': 2.5,
          'tau_er': 9.3,
          'tau_ir': 13.7,
          'beta_er': 4.27,
          'beta_ir': 4.81,
          'I_lin': -0.06,
          'I_sq': 0.41,
          'a_n': 0.55,
          'b_n': 1.48 }


def H(t):
    """
    Evaluate Heaviside unit step function

```

```
35      Accepts both integers and arrays
36      """
37      if isinstance(t, np.ndarray) or isinstance(t, list):
38          return np.array([ 0 if i < 0 else 1 for i in t ])
39      else:
40          return 0 if t < 0 else 1
41
42
43  def h(tau, delta, dt):
44      """Return temporal coupling kernel as np.array"""
45
46      t = np.arange(0, 5*tau, dt)
47      kernel = np.exp( -(t-delta) / tau ) / tau * H(t-delta)
48
49      return kernel
50
51
52  def F_n(I_n):
53      """Evaluate activation function for current I_n"""
54
55      I_lin = param['I_lin']
56      I_sq = param['I_sq']
57      a_n = param['a_n']
58      b_n = param['b_n']
59
60      if I_n < I_lin:
61          return 0
62
63      elif I_n < I_sq:
64          return a_n * (I_n - I_lin)
65
66      else:
67          return a_n * (I_n - I_lin) + b_n * (I_n - I_sq)**2
68
69
70  def f_T(r_T, dt):
71      """
72      Return thalamus contribution to soma current in
73      L4 population
74      """
75
76      h_ef = h(param['tau_ef'], param['delta_ef'], dt)
77      return np.convolve(h_ef, r_T) * dt
78
79
80  def solve(t, r_T):
81      """
82      Solve thalamocortical model and return solution.
83
84      Use RungeKutta Solver
85
86      Arguments:
```

```
87          -- t: time steps
88          -- r_T: thalamic firing rate
89      Return:
90          -- t: time steps
91          -- r_4: L4 population firing rate
92      """
93
94      #Adjust if input times are negative (prestimulus)
95      t_i = t[0]
96      if t_i < 0:
97          t += np.abs(t_i)
98
99      #Model parameters
100     tau_er =  param['tau_er']
101     beta_er = param['beta_er']
102     tau_ir =  param['tau_ir']
103     beta_ir = param['beta_ir']
104
105     dt = np.abs(t[1]-t[0])  # Time interval between steps
106     nt = int(1 / dt) # Number of time steps per unit.
107
108     #Convolve thalamic firing rate with kernel
109     f_T_arr = f_T(r_T, dt)
110
111     def f(u, t):
112         return [ - (u[0] / tau_er) + ( (beta_er/tau_er) \
113                     * F_n(u[0] - u[1] + f_T_arr[t*nt])),
114                  - (u[1] / tau_ir) + ( (beta_ir/tau_ir) \
115                     * F_n(u[0] - u[1] + f_T_arr[t*nt])) ]
116
117     method = ODESolver.RungeKutta4(f)
118     method.set_initial_condition( [0.0, 0.0] )
119
120     x, t = method.solve(t)
121
122     r_4 = np.zeros(len(r_T))
123
124     for i in range(len(r_T)):
125         r_4[i] = F_n(x[i, 0] - x[i, 1] + f_T_arr[i])
126
127     #Reset time steps
128     if t_i < 0:
129         t += t_i
130
131     return t, r_4
132
133
134 if __name__ == '__main__':
135
136     pass
```