Predicting neuronal firing from calcium imaging using a control theoretic approach

Nicholas A. Rondoni[1*], Fan Lu[1], Dan Turner-Evans[2], Marcella Gomez[1]

1 Department of Applied Mathematics, University of California Santa Cruz, Santa Cruz, CA, USA

2 Department of Molecular, Cell and Developmental Biology, University of California Santa Cruz, Santa Cruz, CA, USA

* nrondoni@ucsc.edu

Abstract

Two-photon calcium imaging has become a powerful tool to explore the functions of neurons and the connectivity of their circuitry. Frequently, fluorescent calcium indicators are taken as a direct measure of neuronal activity. These indicators, however, are slow relative to behavior, obscuring functional relationships between an animal's movements and the true neuronal activity. As a consequence, the firing rate of a neuron is a more meaningful metric. Converting calcium imaging data to the firing of a neuron is nontrivial. State of the art methods depend largely on neural networks or non-mechanistic processes, which may yield acceptable correlations between calcium dynamics and the frequency at which a neuron fires, but do not illuminate the underlying chemical exchanges within the neuron or require significant data to be trained on. Leveraging modeling frameworks from chemical reaction networks (CRN) coupled with modern control architectures, a new algorithm is presented based off of fully deterministic ordinary differential equations (ODEs). This framework utilizes model predictive control to challenge state of the art correlation scores while retaining the interpretability that comes inherent with a model. Moreover, these computations can be done in real time, thus enabling online experimentation informed by neuronal

firing rates. To demonstrate the use cases of this architecture, it is tested on a ground truth dataset courtesy of the *spikefinder* challenge. Finally, the predictive power of the methodology put forward here is underscored by demonstrating its ability to aid in the creation of novel indicators utilized in two-photon imaging.

Author summary

We put forward a novel approach to infer when neurons fire as a function of calcium concentration. These calcium recordings are useful for imaging whole populations of neurons, such as those found in the brain, but act only as a proxy for the true underlying spiking occurrences. Moreover, these calcium traces are unfortunately noisy. To uncover the actual firing times we apply modern control architectures to a framework built from chemical reaction equations. The result is accurate, fast enough to provide information in real time, and highly interpretable. More broadly this analysis process aids understanding in contexts where methods of measurement obfuscate the desired ground truth information. Not just capable of inference, prediction is possible thanks to the fully mechanistic formulation provided here. Specifically, we may quantify how indicators, which allow the tracking of calcium, behave in the presence of differing conditions. In more general terms, this framework also enables understanding of the tools used to measure the desired underlying signals.

Introduction

Information processing within populations of neurons is frequently uncovered via two-photon imaging [1]. This method indirectly tracks concentrations of intracellular calcium ions by their fluorescence, which is a noisy realization of the underlying calcium signal [2]. An action potential within a neuron results in an uptake in intracellular calcium, in turn yielding an increase of fluorescence as calcium ions bind to an indicator. Downstream analyses are frequently concerned with true spiking times, not fluorescent traces. Converting between a fluorescent trace and true spiking times is nontrivial, coming with many computational difficulties. Amongst them are disparities in the temporal resolution of imaging compared to neuronal dynamics, noise in fluorescence,

nonlinear relationships between calcium indicators and fluorescence, and biological parameters unknown a priori [3].

This work is concerned with inferring a true firing signal from noisy time series calcium imaging data. State of the art methods approach this in a handful of different ways. Ranging from supervised learning [4], generative methods [5], and nonnegative deconvolutions manifesting as convex optimization problems [6, 7], a litany of theory has arisen as a means to uncover the true spiking occurrences of a neuron. The works of [4, 8] and later [9] present thoughtful surveys of leading methods, while suggesting their preferred approach.

Of particular interest as of late is the ability to analyze systems of neurons in real time. Closed loop investigation of neural circuitry requires quick data processing and carefully crafted experimental conditions. With these conditions met, experiments can deliver sensory stimuli informed by system-wide neural dynamics. For example, a closed loop strategy was employed to tease out how activity in visual processing regions are connected to specific brain states [10]. Recent work from [11] has shown how real time inference of neural activity can inform what stimulation should be supplied to the neurons undergoing imaging, though these authors use calcium traces directly as their measures of activity. Model prediction control (MPC), tracing back to 1960s [12], has caught attention in biological control system, resulting in high levels of control accuracy [13, 14]. In this type of controller, a model, often based on ordinary differential equations, is used to predict optimal control strategy via optimization over a finite time window of data. With only a handful of methods fast enough to be considered real time [7], we aim to add to this body of existing methods by employing tools such as MPC from the controls community.

Methods and Models

This section utilizes a chemical reaction network to build a model, then discusses how model predictive control (MPC) can be tailored to infer underlying firing rates.

Chemical Reaction Formulation

Starting with first principles, we examine the chemical reaction network theorized to dictate the interplay between a calcium ion, a calcium indicator, and a bound (or fluoresced) calcium indicator. These quantities will be denoted $[\mathrm{Ca}^{2+}] = x(t)$, $[\mathrm{CI}] = y(t)$, and $[\mathrm{CI}^*] = z(t)$ respectively, where $[\,\cdot\,]$ denotes concentration. Intracellular calcium binds to an indicator at a rate of k_f. Simultaneously these compounds may unbind at a rate of k_r, leaving the indicator and calcium ion free within the cell once again. This gives

$$\mathrm{Ca}^{2+} + \mathrm{CI} \underset{k_r}{\overset{k_f}{\rightleftharpoons}} \mathrm{CI}^* \tag{1}$$

With this chemical reaction network, we may utilize the law of mass action, which posits the rate of a reaction is proportional to the product of reactant concentrations [15, 16].

Supposing calcium indicators do not passively diffuse out of the cell, the total concentration of calcium indicator is time invariant. That is the sum of concentrations of the indicator, both bound and unbound, should remain constant. Call this value L. Then for all time t

$$y(t) + z(t) = L \qquad\qquad L \in \mathbb{R}^+$$

and the equations that follow from (1) may be reduced to

$$\begin{cases} \dot{x} &= k_r z - k_f x(L - z) \\ \dot{z} &= k_f x(L - z) - k_r z \end{cases} \tag{2}$$

Note these ODEs are nonlinear due to the product of x and z appearing in both equations. Motivated by the observation that a neuron's firing results in an uptake of calcium [17], and ultimately is the dominating force controlling the balance of $x = [\mathrm{Ca}^{2+}]$, we add in a continuous function of time $s(t)$ to represent firing rate at a given time t, scaled by some constant α. We finally note calcium ion x could passively diffuse out of the cell, and subtract a γx term to account for this, arriving at the final governing ODE system

$$\begin{cases} \dot{x} &= \alpha s - \gamma x + k_r z - k_f x(L - z) \\ \dot{z} &= k_f x(L - z) - k_r z \end{cases} \tag{3}$$

where

$$x(t), z(t) : \mathbb{R}^+ \cup \{0\} \to \mathbb{R}, \qquad \alpha, \gamma, k_r, k_f, L \in \mathbb{R}^+$$

In this phrasing, the question of inferring spiking rates as a function of calcium imaging has been cast as an inverse problem, in which we learn the control signal responsible for driving the dynamics witnessed. A table of parameter values and their biophysical interpretation can be found in the supporting information section, S1 Table. Justification that this system is stable for all achievable values of s can be found in the supporting information section, S1 Appendix.

Model Predictive Control and the Neuron

MPC is an optimal control technique. This well tested method computes necessary control actions that minimize a cost function and adhere to an underlying model [18], in our case a constrained set of ODEs defined in equation (3). This control action is computed over a finite horizon, enabling changes to actuation that update with real time information. For all simulations in the results section, the open-source MPC package do-mpc [19], implemented in python, was utilized.

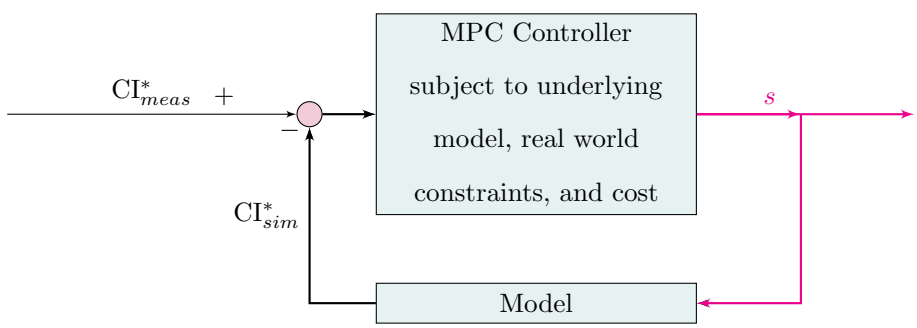

In the above block diagram, CI^*_{meas} denotes the measured amount of fluoresced indicator, while CI^*_{sim} represents the simulated amount of this quantity as dictated by the model (3). Per the definition of MPC, we require a cost defined over our horizon of

n timesteps. For simulations to follow, we consider the following minimization problem

$$\min_{s\in\mathbb{R}^+} \sum_{k=0}^{n-1} \left(\left(\mathrm{CI}^*_{k,sim} - \mathrm{CI}^*_{k,meas}\right)^2 + r(s_k - s_{k-1})^2 \right) + (\mathrm{CI}^*_{sim} - \mathrm{CI}^*_{meas})^2 \tag{4}$$

for parameter $r \in \mathbb{R}^+$. The first squared term enforces our simulated and measured calcium indicators track one another on a finite horizon, and $s_k - s_{k-1}$ penalizes changes in s, essentially limiting the derivative and encouraging smoothness. All simulations to follow take $r = 0.01$ as regularization term. Other authors have had success with the similar cost functions, sometimes using a l_0 or l_1 penalty on s instead of the quadratic cost above [6, 7]. Observe that the final term in (4) acts as terminal penalty function, which aids the effectiveness of MPC. An improperly chosen terminal penalty function may degrade performance and potentially destabilize the closed-loop system. Although increasing the prediction horizon can improve performance, this significantly raises computational costs. A well-established solution to these issues involves selecting the terminal penalty function as the infinite-horizon value function that satisfies the dynamic programming equations [20, 21], as is done above. The horizon length $n = 6$ was utilized for all simulations.

In order to partially address the noisy fluorescence signal, a shifted sigmoidal filter

$$\sigma(x) = \frac{1}{1 + e^{-(x+1)}} - h$$

is applied to CI^*_{meas}. To avoid saturation of this filter, e.g., CI^*_{sim} values becoming stuck at 1, a vertical shift downward of $h = 0.15$ is additionally used. If saturation still occurs, a more aggressive vertical shift of $h = 0.25$ is then applied and computations are restarted. The nature of a sigmoidal curve minimizes contributions from extreme values, mapping them closer to 0 or 1, while retaining a linear regime for middle of the pack measurements. More advanced filtering techniques to address stochasticity inherent in the system could undoubtedly improve this methodology, this is discussed briefly in the following section.

Results and Discussion

With a model formulated and the tools of inference defined, this section benchmarks performance using the *spikefinder* [22] challenge dataset.

Spikefinder Validation

This dataset contains not only time series calcium imaging data, but recorded "ground-truth" spikes via electrophysiology. It is one of the few public datasets to contain simultaneous recordings of spiking times and fluorescent traces, allowing for validation of many of the models discussed in this document. An application of our methodology, henceforth called the MPC approach, to a particular neuron produces a time series of simulated bound calcium indicator CI^*_{sim} and spiking rates s as a function of measured indicator CI^*_{meas}, visualized below in Fig 1.

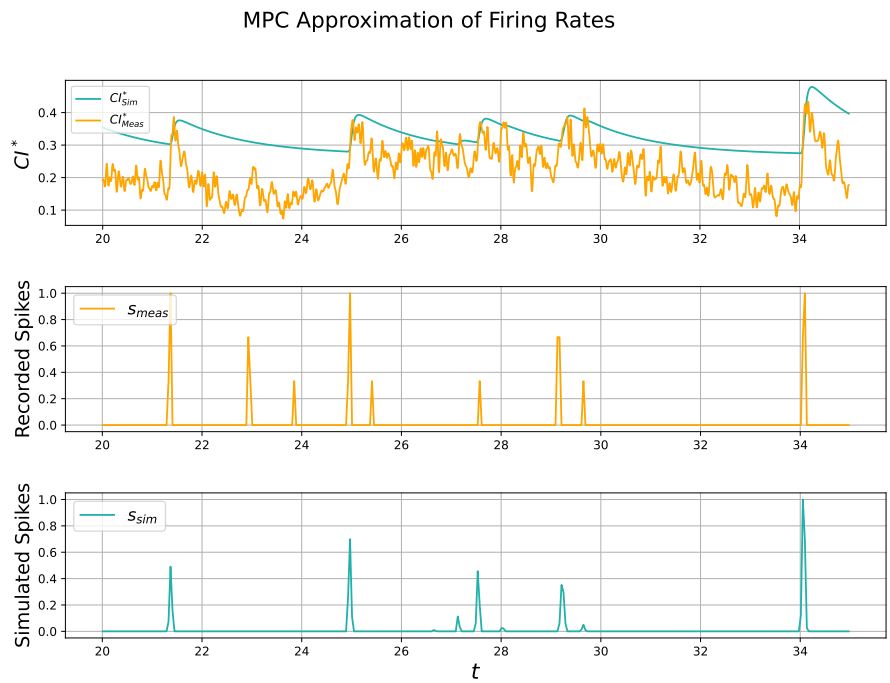

Fig 1. Visualization of Approximations Pearson correlation coefficient is 0.694 for this particular 15 s subset of test data set 1, neuron 0. The whole 12 min recording scores a 0.491. Spikes normalized to be within [0, 1] for ease of viewing. A few false spikes are present, while others are missed.

The primary metric valued during the *spikefinder* challenge was Pearson correlation

coefficient. While emphasis is still placed on correlation, alternative methods of measuring spike synchrony have been put forward such as Victor-Purpura and van Rossum metrics [23–25]. Since spikes are inherently discrete events, for comparison of two spike trains to be computationally tractable and informative we downsample in accordance with the *spikefinder* challenge methodology.

Beginning with correlation coefficient, we compare against the state of the art algorithms STM [4] and Oasis [7] in Fig 2. Here STM serves as a representative of supervised methods, while Oasis is a NND algorithm. Note the work of [7] is another method that is certainly fast enough for real time applications, alongside the MPC approach put forward here. Oasis was eventually benchmarked on the spikefinder dataset by [9], this NND method has a slight advantage in correlation score, but does not have the same amount of interpretability as the MPC approach.

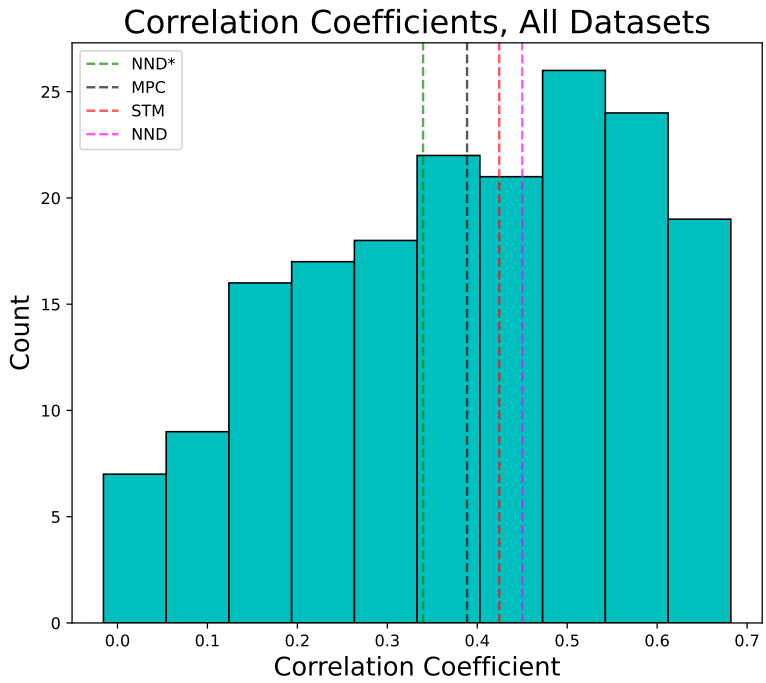

Fig 2. Comparison to state of the art. Histogram of correlation coefficients across all data sets, both train and test. The red dashed line denotes the mean of the STM algorithm, green denotes Oasis supplied with information about the indicator, pink denotes Oasis with no information about the indicator, and the blacked dashed line denotes the mean of the plotted dataset. Means presented are across all data sets, train and test, as well. Resultant firing signals downsampled by a factor of 4 in accordance with the *spikefinder* challenge, yielding a bin width of 40 ms.

As can be seen in Fig 2, our Δ mean relative to STM is only 0.03. While STM does perform better, it requires ample training datasets. The authors of STM note their algorithm does have the capability to perform well when applied to datasets not seen during training, but this is after a wide variety of datasets have already been processed. For the most accurate predictions, such as the results that produced the mean in figure 3, STM is trained on datasets similar to its eventual testing datasets. Put more concisely, STM's distribution of training data resembles the distribution of testing, at least for the spikefinder challenge. When considering specific datasets the difference in performance is less clear. There exist individual datasets on which the STM algorithm performs notably better than this MPC approach and vis versa. For example, consider Fig 3.

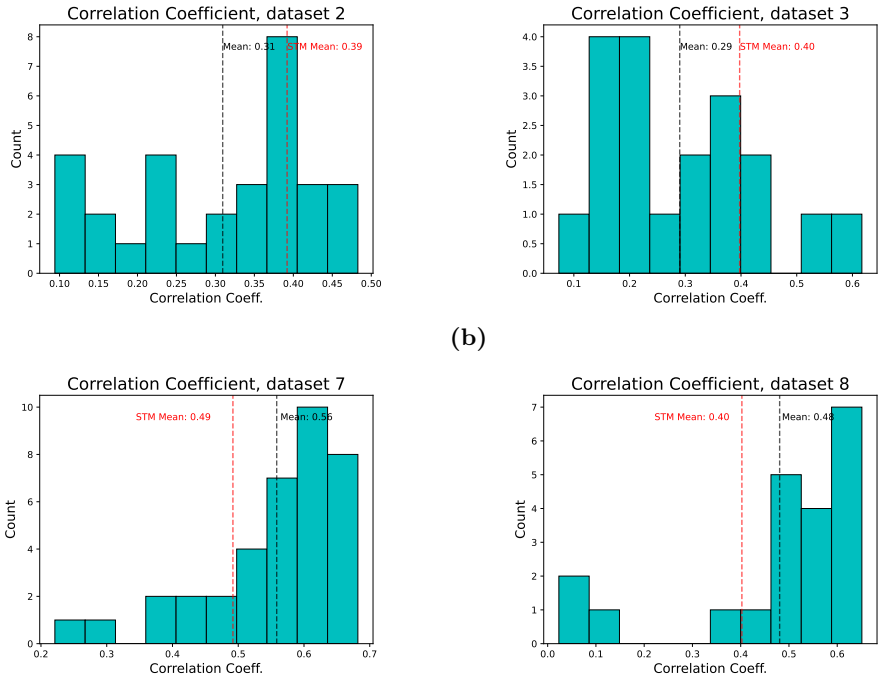

(a) **(b)**

(c) **(d)**

Fig 3. Comparison to STM on specific datasets. Histogram of correlation coefficients computed for four different datasets. (a) and (b) both showcase STM outperforming the MPC approach, while (c) and (d) show the MPC approach outperforming. Datasets 2 and 3 had both train and test data available - both are included. For datasets 7 and 8 only training data was made available.

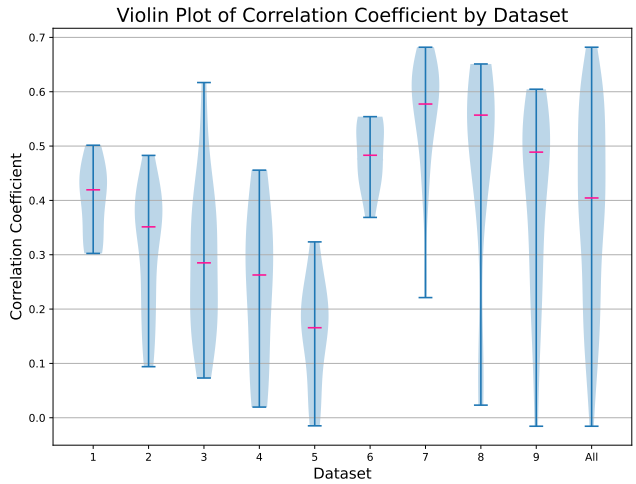

Fig 4. Violin plot of correlation score Pearson correlation coefficients from MPC approach organized by data set, 40 ms bin sizes. The violin plot uses kernel density estimation to compute an empirical distribution of the samples by dataset. Medians of MPC approach are included in pink.

The MPC algorithm does not require training in the same way a supervised approach such as STM would. Rather, parameters may be selected a priori from known chemical reaction constants, however this undoubtedly results in a sub-optimal parameter regime. For the simulations shown in this section parameters were selected for their biophysical relevance then manually adjusted on training sets in order to maximize correlation coefficients. Thus there is a small training phase, though much shorter and less efficient than that of a neural net's. With enough biophysical insights, an auto-calibration of parameters should be possible similar to that of MLspike [5], completely alleviating the need to hand train. STM has an auto-calibration functionality, but it requires introducing stimulus and spike history dependence [4].

Prior to examining the aforementioned metric Victor-Purpura distance [23], recall that lower scores are better. In Fig 5 we compare the naive Oasis, meaning Oasis supplied with no information, to the MPC approach. The cost of moving a spike a single timestep is 1 for this simulation.

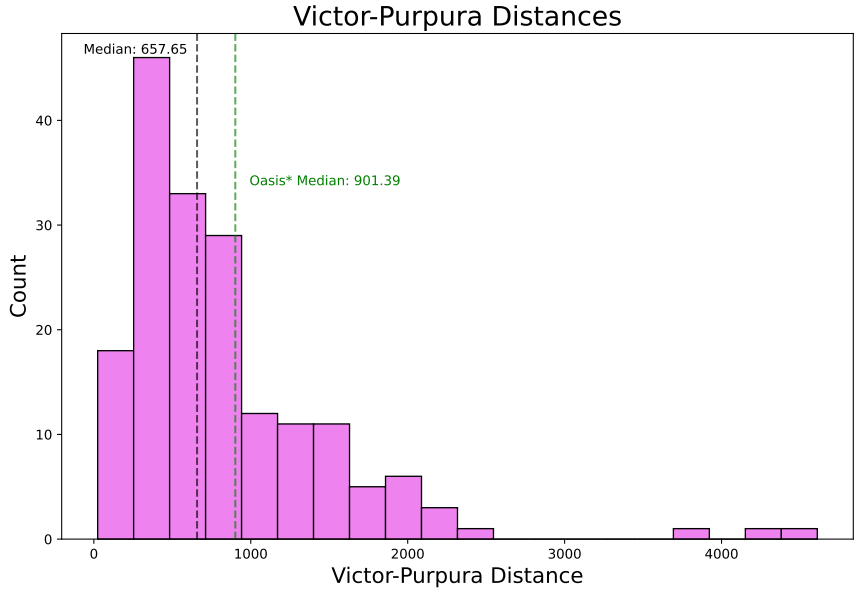

Fig 5. Comparison of Victor-Purpura distance Victor-Purpura distances across all datasets. Median value of MPC approach shown with a black dashed line, while the median of Oasis (supplied with no information) is shown in green. Since lower is better, we outperform the naive oasis method. A single outlier at 8000 has been excluded from the plot for ease of viewing, leaving 178 different samples visualized in this histogram. Medians were computed with all samples included.

Reflections on Noise

Though this MPC approach does not require training in a conventional neural net sense, careful selection of parameters is still imperative. This can be seen by considering Figs 6 and 7 in tandem.

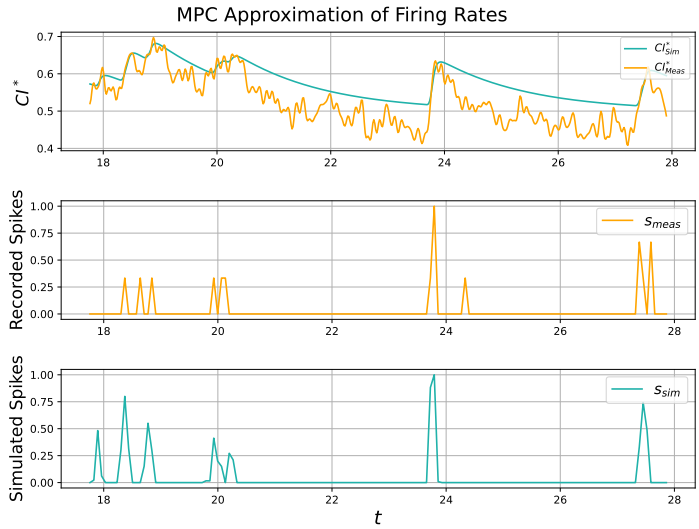

Fig 6. Pearson correlation coefficient for this specific 10 s subset is 0.61, though the whole 12 min recording scores a 0.50. Spikes normalized to be within [0, 1] for ease of viewing. Though most are detected, one is missed at around 24s, with timings slightly off for other spikes.

Fig 7. Noise in CI^*_{meas} causes the MPC approach to record many false spikes. Correlation coefficient is 0.19

In Fig 6 near optimal parameters are utilized, which incorporates a slow decay rate. This slow decay rate abstracts away noise in the fluorescence signal by avoiding small increases not worthy of a spike. In contrast, Fig 7 has parameters selected to encourage almost perfect tracking of CI_{meas}. While this is easily accomplished via the MPC algorithm, the resultant signal is riddled with false spikes, and thus has a low correlation coefficient. This is because under these circumstances the algorithm tracks the *noise in imaging* instead of the true underlying calcium signal. As mentioned in the introduction, fluorescence is a noisy realization of the true calcium trace, and as such the MPC approach would certainly be improved by incorporating a nonlinear conversion between normalized fluorescence and CI^*_{meas}. Authors of MLspike [5] incorporate such ideas to great success. In contrast our current formulation equates these two quantities.

Remarks on Real Time Control

Regarding the usability of methods presented in this document for online control, when training is accounted for only our MPC approach and the NND algorithms are fast enough or structurally able to complete in real time. That is, for all solutions presented the MPC algorithm was able to predict spiking signals in equivalent or less time than the duration of the recording. Though these signals were computed offline, the methodology needs no adaption to run out of the box in real time.

While Jewell's updated methodology (along with other efficient NND approaches) can impressively compute spiking times for data sets with as many as 100,000 time instants in mere seconds on a laptop, this first requires knowledge of a parameter or training to be accurate. In Jewell's approach the relevant parameter was discovered by training on the first half of a data set then testing on the second half, suggesting training is still important if high accuracy is desirable. It is worth noting the authors have a way to infer a reasonable parameter regime, though this was not used for testing. In the case of the suite2p implementation of Oasis [9], training is not required but knowledge of the system must be supplied in some way to bring scores in line with means presented in plot 2. Otherwise, the means will be closer to the naive Oasis* instead of the more accurate Oasis. NND models complete on the order of milliseconds however, which is notable considering their stand out accuracy.

The MPC algorithm takes longer than these NND methods; with run times almost always 10-30% less than the duration of their recording. However there is a rich legacy of MPC used for real time control, with many packages optimized for speed and ease of implementation, and as such is relevant in this context for its accuracy and tractable nature. Moreover, this MPC approach allows for other quantities of interest to be inferred, such as the concentration of unbound calcium indicator within the cell. This is a result of our model's rooting in biophysical mechanisms.

Biophysical Insights

With the MPC approach and the model underlying it justified, analysis of the system may be considered. Thus far we have been told the status of state variables and inferred s from this signal. Reaping the benefits of a fully mechanistic model, we may now suppose s and analyze the impacts on state variables.

To accomplish this, we perform frequency response analysis on the system of ODEs (3),

$$\begin{cases} \dot{x} &= \alpha s - \gamma x + k_r z - k_f x(L - z) \\ \dot{z} &= k_f x(L - z) - k_r z \end{cases} \tag{3}$$

then move to show how this process can be reverse engineered such that an indicator may be designed to work within a specific regime of frequencies.

Since our ODEs are nonlinear many of the elegant results pertaining to frequency response of linear systems is unavailable. Instead, we must suppose a range of frequencies at which s oscillates then note the impacts on our state variables of interest. Specifically, take

$$s(t) = A\sin(\omega t) + c$$

for parameters $A, \omega, c \in \mathbb{R}$. This process is showcased for a single ω in Fig 8. The response in calcium ion $[\mathrm{Ca}^{2+}] = x$ and indicator $[\mathrm{CI}^*] = z$ follow from the supposed s in the leftmost plot.

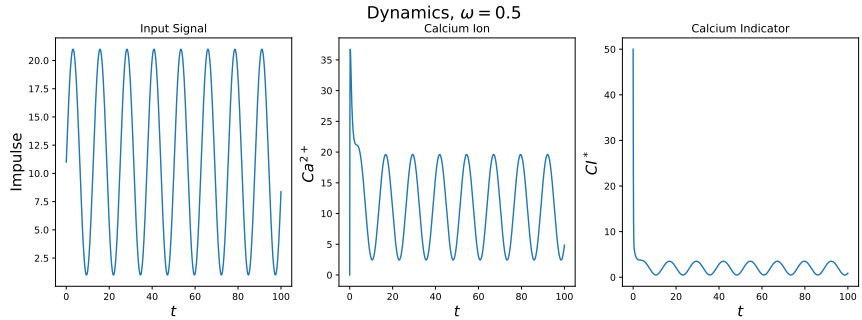

Fig 8. Parameters: $A = 20$, $c = 1$, $\omega = 0.5$. Notice after transient dynamics the resultant amplitudes for the calcium ion, indicator are not as large as the amplitude of the s signal. This suggests these state variables do not keep up with changes in s at this frequency.

To avoid confusion, note $s(t)$ is the firing rate at given time t and is measured in Hz. Here we examine how frequently s must change before state variables begin to lag. This is determined by the parameter ω, which determines the frequency of the s signal, also measured in Hz. For the previous example, $\sin(0.5t)$ has period 4π and thus oscillates at a frequency of $1/(4\pi) \approx 0.080$ Hz between the values of 1 and 21 Hz.

To evaluate the ODE system's performance for a range of ω, we perform simulations akin to those done in Fig 8. However after transient dynamics conclude, we note the ratio of amplitudes between s and our state variables. This process is showcased in Fig 9. In this way we are able to tease out at which frequencies our state variables are able to reasonably respond to.

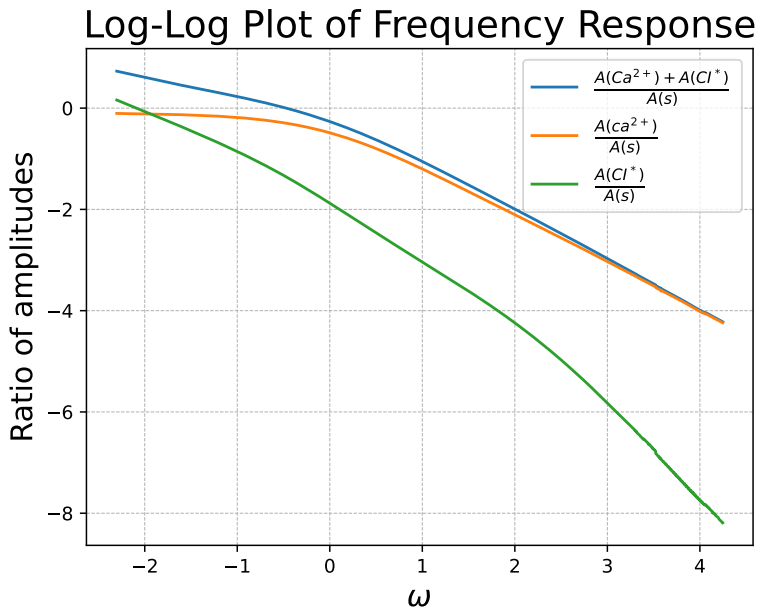

Fig 9. Frequency response analysis of commonly used indicator GCaMP6s, $k_f = 0.0514$, $k_r = 7.6$, $L = 30$. Shows for higher frequency oscillations in firing rate calcium dynamics cannot keep up, evident in the amplitude of s dominating the amplitude of calcium. Blue shows the total amount of calcium in the system, both bound and unbound. Orange depicts the unbound calcium ions themselves, while green shows the calcium ions bound to an indicator. Justifies utilizing firing rate as the metric for neural activity.

The above calcium indicator keeps up with changes in firing rate poorly for higher frequency oscillations, highlighting the need to convert recorded calcium indicator quantities to firing rate for certain downstream analyses in which a neuron may change the rate at which it fires frequently.

This not only has the power to analyze frequency response for existing indicators, but the potential to aid in the development of novel indicators. For example, if we desired an indicator which operated better for smaller values of ω than GCaMP6s, we might try setting new k_f and k_r then examine the resultant frequency response plot. To this end, consider Fig 10, which achieves precisely this by utilizing $k_f = 0.01$ and $k_r = 10$ as reaction rates for this toy example.

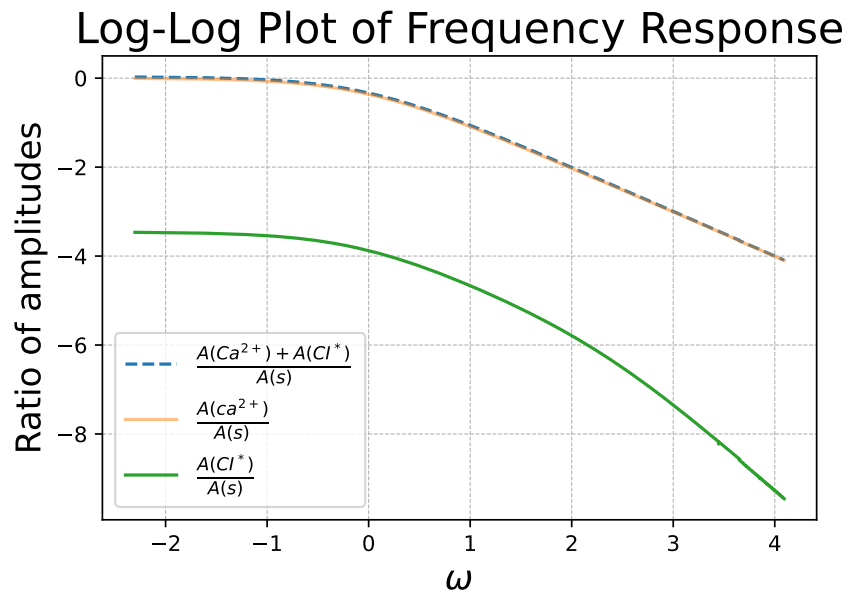

Fig 10. Frequency response of a hypothetical indicator. All parameters are held constant aside from $k_f = 0.01$ and $k_r = 10$. Notice an improvement in ratio of amplitudes at small ω when compared to figure 9. A drop off still begins at $\omega = 0$ or about $e^0 = 1$ Hz.

The task at hand then becomes creating an indicator with binding affinity $k_d = k_f/k_r = 0.01/10$. If additional control over indicator dynamics is desirable, each parameter in (3) may be tuned such that the frequency response is satisfactory. Constructing an indicator which agrees with all parameters in our governing ODE equations is likely difficult, but nonetheless this has the potential to serve as an approximate guide for the creation of future indicators.

As discussed in [26], general performance criteria to consider in the construction of GECIs include

(i) large dynamic range

(ii) high calcium sensitivity

(iii) faster response kinetics

(iv) linear response properties.

The frequency response analysis pipeline outlined in this section could prove invaluable to evaluating the above performance metrics.

Future Work

An auto-calibration of parameters α, γ in equation (3) would increase the usability of this software. To accomplish this, some connection between these parameters and the dynamic qualities of the fluorescent trace must be established. In the style of suite2p, k_f and k_r could be chosen by simply supplying knowledge of the indicator's speed. Finally, any of these terms could be learned through supervised learning if labeled data is available.

As another avenue for improvement, consider the terminal penalty function of (4). This function can also be approximated via the Q function in reinforcement learning (RL) [27], presenting an opportunity for future directions of this work in applying RL-based approaches for efficient MPC design.

Finally, a more sophisticated treatment of noise would increase accuracy further. This could be achieved by coupling another ODE connecting fluorescence to calcium, as is done by MLspike, or by employing particle filtering methods. Sequential Monte Carlo methods appear uniquely suited to address the stochasticity present in the system [28].

Conclusion

In this paper, we propose a novel modeling framework and strategy to infer underlying firing rates of neurons provided calcium imaging traces. This first principles approach underscores the potential of adapting modern control architectures to this space of problems. Accurate enough to challenge state of the art methods, this formulation joins the select few methodologies capable of providing real time spiking information. This fully mechanistic model is highly interpretable and capable of not just inference, but can uncover and predict biophyiscal nuances as they relate to hyperparameters such as reaction rates of indicators.

Supporting information

S1 Appendix. Stability analysis. In this appendix the governing ODE system (3) is shown to be stable for all realistic values of s. Begin the stability analysis by considering the intersection of nullclines. Let $\dot{x} = 0$, $\dot{z} = 0$. After a bit of algebra the

coordinates of fixed point (x^*, z^*) is found to be

$$(x^*, z^*) = \left(\frac{k_r z}{k_f(L-z)}, \frac{\alpha s k_f L}{\alpha s k_f + \gamma k_r} \right) = \left(\frac{\alpha s}{\gamma}, \frac{\alpha s k_f L}{\alpha s k_f + \gamma k_r} \right)$$

The Jacobian of the system is given by

$$J_f(x,z) = \begin{bmatrix} -\gamma - k_f \left(L - \frac{\alpha s k_f L}{\alpha s k_f + \gamma k_r} \right) & k_r + k_f x \\ k_f L - k_f \left(\frac{\alpha s k_f L}{\alpha s k_f + \gamma k_r} \right) & -k_f \left(\frac{\alpha s}{\gamma} \right) - k_r \end{bmatrix}$$

Examining the eigenvalues of this matrix upon evaluation at (x^*, z^*), the condition put on s in order yield two negative eigenvalues, and hence stability, is simply that $s \in \mathbb{R}$. Note these bounds are found from numerical approximations, as an exact solution becomes intractable.

Since the range of s is restricted to the positive reals, it follows that we have stability for all attainable values of s. To visualize that these ODEs produce a stable fixed point as a function of s consider the following phase diagram, in which s is fixed to be 30. Since our fixed point is a function of s, it follows that taking s as a time varying function $s(t)$ would simply shift the location of our attractor. Taking parameters $\alpha = \gamma = 1, k_r = 7.6, k_f = 0.05135, L = 100, s = 30$, we would expect our fixed point to be at

$$(x^*, z^*) = \left(\frac{\alpha s}{\gamma}, \frac{\alpha s k_f L}{\alpha s k_f + \gamma k_r} \right) = (30, 16.854)$$

This is visualized in Fig 11.

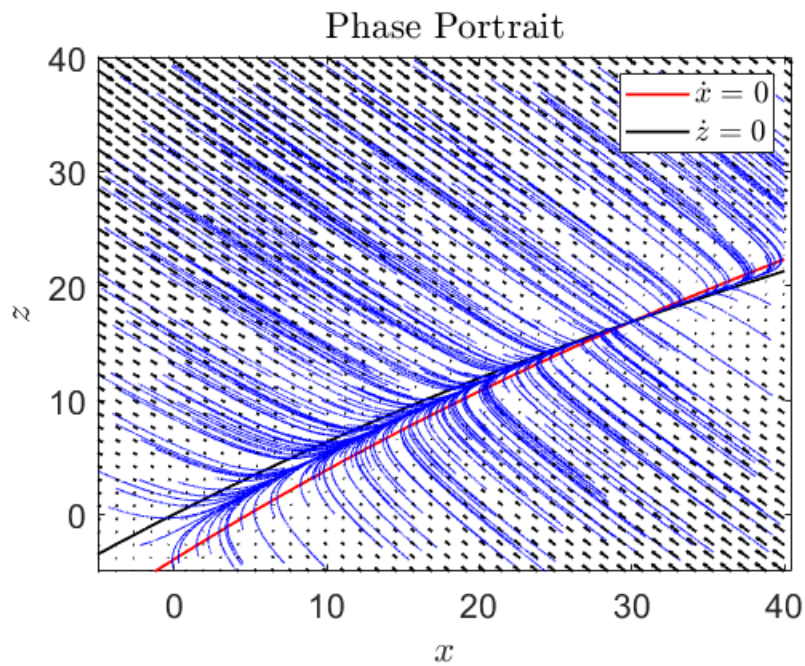

Fig 11. Flows of equation (3). For this simulation, $\alpha = \gamma = 1$, $k_r = 7.6$, $k_f = 0.05135$, $L = 100$, $s = 30$. Note our fixed point is, as was expected, at $(30, 16.854)$.

S1 Table Parameter values and physical interpretation. The following tables provide physical interpretation and units for governing equation (3). Values provided were utilized the in *spikefinder* simulations.

Table 1. Datasets 1-3, and 6-9

Parameter	Value	Units	Description
k_r	0.1	$\mu M^{-1}s^{-1}$	Forward reaction rate
k_f	10	s^{-1}	Reverse reaction rate
α	16.66	μM/spike	Weighting of firing
γ	1.366	s^{-1}	Passive diffusion scalar
L	100	μM	Total indicator available
Δt	0.01	s	Simulation step size

Table 2. Dataset 4

Parameter	Value	Units	Description
k_r	0.1	$\mu M^{-1}s^{-1}$	Forward reaction rate
k_f	10	s^{-1}	Reverse reaction rate
α	10	μM/spike	Weighting of firing
γ	0.73	s^{-1}	Passive diffusion scalar
L	100	μM	Total indicator available
Δt	0.01	s	Simulation step size

Table 3. Dataset 5

Parameter	Value	Units	Description
k_r	0.2	$\mu M^{-1}s^{-1}$	Forward reaction rate
k_f	10	s^{-1}	Reverse reaction rate
α	10	μM/spike	Weighting of firing
γ	0.73	s^{-1}	Passive diffusion scalar
L	100	μM	Total indicator available
Δt	0.01	s	Simulation step size

Acknowledgments

Thank you to the Gomez lab group, in particular Ksenia Zlobina, for insightful discussions pertaining to modeling and analysis.

References

1. Jason N D Kerr and Winfried Denk. Imaging in vivo: watching the brain in action. *Nature Reviews Neuroscience*, 9(3):195–205, March 2008.

2. Carsen Stringer and Marius Pachitariu. Computational processing of neural recordings from calcium imaging data. *Current Opinion in Neurobiology*, 55:22–31, 2019. Machine Learning, Big Data, and Neuroscience.

3. Joshua T Vogelstein, Brendon O Watson, Adam M Packer, Rafael Yuste, Bruno Jedynak, and Liam Paninski. Spike inference from calcium imaging using sequential monte carlo methods. *Biophysical Journal*, 97(2):636–655, July 2009.

4. Lucas Theis, Philipp Berens, Emmanouil Froudarakis, Jacob Reimer, Miroslav Román Rosón, Tom Baden, Thomas Euler, Andreas S Tolias, and Matthias Bethge. Benchmarking spike rate inference in population calcium imaging. *Neuron*, 90(3):471–482, May 2016.

5. Thomas Deneux, Attila Kaszas, Gergely Szalay, Gergely Katona, Tamás Lakner, Amiram Grinvald, Balázs Rózsa, and Ivo Vanzetta. Accurate spike estimation from noisy calcium signals for ultrafast three-dimensional imaging of large neuronal populations in vivo. *Nature Communications*, 7(1):1–17, November 2016.

6. Sean W Jewell, Toby Dylan Hocking, Paul Fearnhead, and Daniela M Witten. Fast nonconvex deconvolution of calcium imaging data. *Biostatistics*, 21(4):709–726, 02 2019.

7. Johannes Friedrich, Pengcheng Zhou, and Liam Paninski. Fast online deconvolution of calcium imaging data. *PLOS Computational Biology*, 13(3):1–26, 03 2017.

8. Lucas Theis, Philipp Berens, Emmanouil Froudarakis, Jacob Reimer, Miroslav Román Rosón, Tom Baden, Thomas Euler, Andreas S. Tolias, and Matthias Bethge. Benchmarking spike rate inference in population calcium imaging. *Neuron*, 90(3):471–482, 2016.

9. Marius Pachitariu, Carsen Stringer, and Kenneth D. Harris. Robustness of spike deconvolution for neuronal calcium imaging. *Journal of Neuroscience*, 38(37):7976–7985, 2018.

10. Chun-Feng Shang, Yu-Fan Wang, Mei-Ting Zhao, Qiu-Xiang Fan, Shan Zhao, Yu Qian, Sheng-Jin Xu, Yu Mu, Jie Hao, and Jiu-Lin Du. Real-time analysis of large-scale neuronal imaging enables closed-loop investigation of neural dynamics. *Nature Neuroscience*, 27(5):1014–1018, May 2024.

11. Zac Bowen, Dulara De Zoysa, Kelson Shilling-Scrivo, Samira Aghayee, Giorgio Di Salvo, Aleksandr Smirnov, Patrick O. Kanold, and Wolfgang Losert. Neuroart: Real-time analysis and targeting of neuronal population activity during calcium imaging for informed closed-loop experiments. *eNeuro*, 11(10), 2024.

12. Carlos E Garcia, David M Prett, and Manfred Morari. Model predictive control: Theory and practice—a survey. *Automatica*, 25(3):335–348, 1989.

13. Remy Chait, Jakob Ruess, Tobias Bergmiller, Gašper Tkačik, and Călin C Guet. Shaping bacterial population behavior through computer-interfaced control of individual cells. *Nature communications*, 8(1):1535, 2017.

14. Zachary R Fox, Steven Fletcher, Achille Fraisse, Chetan Aditya, Sebastián Sosa-Carrillo, Julienne Petit, Sébastien Gilles, François Bertaux, Jakob Ruess,

and Gregory Batt. Enabling reactive microscopy with micromator. *Nature Communications*, 13(1):2199, 2022.

15. Péter Érdi and János Tóth. *Mathematical models of chemical reactions : theory and applications of deterministic and stochastic models*. Princeton University Press, Princeton, N.J, 1989.

16. Martin Feinberg. *Correction to: Foundations of Chemical Reaction Network Theory*. Springer International Publishing, 2022.

17. Christof Koch. *Biophysics of Computation: Information Processing in Single Neurons* . Oxford University Press, 11 1998.

18. J.M. Maciejowski. *Predictive Control: With Constraints*. Prentice Hall, 2002.

19. Felix Fiedler, Benjamin Karg, Lukas Lüken, Dean Brandner, Moritz Heinlein, Felix Brabender, and Sergio Lucia. do-mpc: Towards fair nonlinear and robust model predictive control. *Control Engineering Practice*, 140:105676, 2023.

20. David Q Mayne, James B Rawlings, Christopher V Rao, and Pierre OM Scokaert. Constrained model predictive control: Stability and optimality. *Automatica*, 36(6):789–814, 2000.

21. Rong-Rong Chen and Sean Meyn. Value iteration and optimization of multiclass queueing networks. *Queueing Systems*, 32:65–97, 1999.

22. Philipp Berens, Jeremy Freeman, Thomas Deneux, Nikolay Chenkov, Thomas McColgan, Artur Speiser, Jakob H Macke, Srinivas C Turaga, Patrick Mineault, Peter Rupprecht, Stephan Gerhard, Rainer W Friedrich, Johannes Friedrich, Liam Paninski, Marius Pachitariu, Kenneth D Harris, Ben Bolte, Timothy A Machado, Dario Ringach, Jasmine Stone, Luke E Rogerson, Nicolas J Sofroniew, Jacob Reimer, Emmanouil Froudarakis, Thomas Euler, Miroslav Román Rosón, Lucas Theis, Andreas S Tolias, and Matthias Bethge. Community-based benchmarking improves spike rate inference from two-photon calcium imaging data. *PLoS Comput Biol*, 14(5):e1006157, May 2018.

23. Jonathan D Victor and Keith P Purpura. Metric-space analysis of spike trains: theory, algorithms and application. *Network: Computation in Neural Systems*, 8(2):127–164, 1997.

24. Conor Houghton and Kamal Sen. A New Multineuron Spike Train Metric. *Neural Computation*, 20(6):1495–1511, 06 2008.

25. Mark van Rossum. A novel spike distance. *Neural Computation*, 13:751–763, 04 2001.

26. Masatoshi Inoue. Genetically encoded calcium indicators to probe complex brain circuit dynamics in vivo. *Neurosci. Res.*, 169:2–8, August 2021.

27. Fan Lu, Joel Mathias, Sean Meyn, and Karanjit Kalsi. Convex q-learning in continuous time with application to dispatch of distributed energy resources. In *2023 62nd IEEE Conference on Decision and Control (CDC)*, pages 1529–1536. IEEE, 2023.

28. Nicolas Chopin and Omiros Papaspiliopoulos. *An introduction to sequential monte carlo*. 12 2020.