Improvement in model flexibility reveals a minimal signalling pathway that explains T cell responses to pulsatile stimuli

Jun Allard[a]

[a]Department of Mathematics, Department of Physics & Astronomy, Center for Complex Biological Systems, University of California Irvine

Abstract

Efforts to develop quantitative models must face a trade-off between interpretability and quantitative accuracy, which often disfavores interpretability. Here we adopt an operational definition of interpretability, specifically that a model is described by an arrow diagram wherein each arrow corresponds to a positive effect or negative effect of one component upon a process, and fewer arrows is more interpretable than more arrows. We then develop a method to add flexibility — and thus accuracy in fitting data — to mathematical models by relaxing functional form assumptions, while constrained by the same arrow diagram and thus the same interpretability. We apply this method to the T cell, where quantitative models are particularly needed, in part because of ongoing efforts to engineer T cells as therapeutics. Recent experiments exposed T cells to pulsatile inputs and measured their frequency response, and found several nonlinear frequency responses: high-pass, low-pass, band-pass, and band-stop. Using our modeling approach with enhanced flexibility, we show that a simple signaling model quantitatively captures the frequency response of CD69 surface expression, one of the correlates of T cells activation, with accuracy within the experimental inter-replicate standard error of the mean. Specific qualitative behaviors map to specific parts of the arrow diagram: Band-pass behavior can be explained by refractory de-sensitizing circuit (we refer to this as "first-aid icing a wound"). Band-stop behavior can be explained by removal-inhibition (we refer to this as "roommate interrupts my studying"). Apparent low-pass emerges naturally when total stimulation time is constant. Taken together, our results demonstrate the ability to achieve both quantitative prediction and interpretability in understanding cellular dynamics. Simple models may at first appear incapable of explaining complex data, but might indeed be able to by adding this modest flexibility.

Introduction

T cells accept input in the form of ligands, which activate receptors, to produce cellular outputs. For the T cell, a major input is antigenic peptide presented by MHC, which activates the T Cell Receptor, and the cellular outputs include an array of immunological function, like cytokine secretion and surface molecule expression [1, 2]. Due to their role in disease and therapeutics, there is an acute need for quantitative, predictive models of whole-cell responses from input to output [3–6].

Much work has been done on understanding the T cell's response to inputs of different molecular identities, e.g., ligands with different affinities or accessory ligands, or different density [1, 7]. One alternative way to experimentally characterize a complex system – with rich history of success across engineering and biology [8, 9] – is by presenting the system with pulsatile inputs, and varying the frequency or on-time (T_{on}) and off-time (T_{off}), as shown in Figure 1A. Indeed, it is hypothesised that T cells may be able to integrate pMHC signals across different cells, which would imply T cell behaviour may be regulated by pulses of antigen encounters on the timescales of minutes to hours [11]. Pulsatile inputs have been presented to T cells using optogenetics [10, 11]. Fascinatingly, researchers found all four basic categories depending on cell type and specific output observed, roughly speaking:

- High-pass, in which gaps between pulses reduces the output;
- Low-pass, in which gaps between pulses enhance the output, to an even higher response than no-gap stimulation;
- Band-stop, in which intermediate-duration gaps reduce the output [11];
- Band-pass, in which intermediate-duration gaps enhance the output.

Several categorically different responses are evident in a dataset from Harris et al. [10] reproduced in Figure 1B, where the output is surface expression of CD69, a correlate of T cell activation.

One way to mathematically characterize a complex system is with phenotypic modeling [2], in which we seek a minimalist pathway, often expressed as system of coupled nonlinear differential equations that represent an arrow diagram, as in Figure 2A. Components in the minimalist pathway represent a coarse-grain view of molecular interactions — where each component contains complex biophysics and nonlinearities [12–18] and may not correspond to a specific molecule.

Here, we ask, what is the phenotypic model that can explain the response of T cells to pulsatile stimulation observed in CD69 measurements by Harris et al. [10]?

Quantitative predictive models battle between interpretability and predictive accuracy [19]. Models with many degrees of freedom often represent data more accurately, but may be more difficult to interpret. There are several transformative approaches to defeat this trade-off [20–22]: approaches that approximate both a simple dynamical law and the data [23]; approaches that begin with a large list of symbolic forms for mathematical expressions and perform sparse regression to select from these [24]; and approaches that allow flexibility in specific terms in the differential equation using universal functions [25–28].

Here we adopt the goal of capturing the data to within its experimental uncertainty, with an arrow diagram where each arrow is a monotonic-increasing or monotonic-decreasing function (represented by arrow-head or bar-head, respectively, in Figure 2A). In other words, herein, our operational definition of interpretability is an arrow diagram where each arrow describes a postive effect of negative effect on its target process, and fewer arrows is more interpretable than more arrows. The method we develop to achieve this is thus a special case of the above-mentioned approaches to balance interpretability and accuracy. One crucial difference in

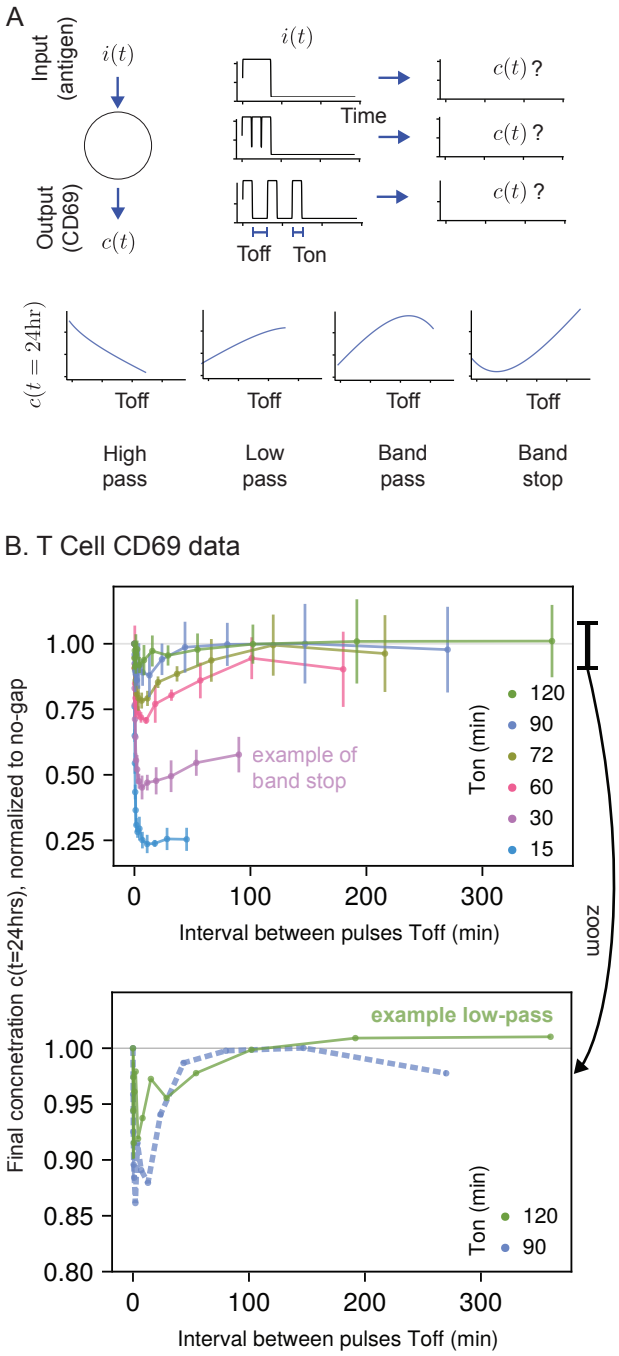

Figure 1: **Experimentally-observed response of T cells to pulsatile stimuli reveals several categorically different features.** A. Exposing a system to inputs $i(t)$ with cycles of on-times T_{on} and off-times (intervals between stimulus, T_{off}). In this case, the system is a T cell, and the input is optogenetically-activated antigen stimulation, and the output is surface expression of CD69. There are several responses, termed "frequency responses" and typically named after their response to frequencies (so, for example, "high pass" is the term for a strong response at high frequency, or short T_{off}). B. The response of T cells to pulsatile exposure to antigen stimulation from Harris et al. [10]. Surface expression of T cells was measured after 24 hours. Data is normalized to the response of a continuous 120 minutes of stimulation ($T_{\text{off}} = 0$). Error bars indicate the standard deviation between replicates (note that in figures with model fitting, the standard error of the mean is shown instead). "No-gap" refers to $T_{\text{off}} = 0$.

our goal, and therefore method, is that the arrow diagram is viewed as a hypothesis to test and potentially reject. And so, the additional model flexibility must be within this hypothesis.

Using this, we show that a simple signaling model quantitatively captures the frequency response of CD69 surface expression, with accuracy within the experimental inter-replicate standard error of the mean. Simple models may at first appear incapable of explaining complex data, but might indeed be able to by relaxing functional form assumptions.

Results

Subsets of CD69 frequency response data can be explained by simple models, when flexibility is added to model assumptions

We began by listing four pathway arrow diagrams of increasing complexity. The simplest model we study is a linear pathway with 3 components $a(t), m(t), c(t)$. In this model, each component is produced in response to the previous component, with the first component responding to the input $i(t)$ and the last component $c(t)$ identified as surface expression of CD69. Since the quantitative features of the CD69 data exhibit significant nonlinearity, we explored three more models with increasing complexity starting from the linear pathway model, shown in Figure 2A. (In the future, it may be possible to combine our method with methods that systematically explore the set of arrow-diagrams, see Discussion and [2, 29, 30]). So, the next model we study includes $a(t)$ inhibiting the removal of $m(t)$ (Removal inhibition, RI). Then, we add an incoherent feed-forward loop (IFF), a motif that appears across biology [31, 32] including T cell models [4, 33]. In the incoherent feed-forward model, a new component w acts as a refractory variable, accumulating in response to a and reducing the accumulation of m. The final model has both removal inhibition and incoherent feed-forward (RI+IFF).

One strategy to represent regulation, which has proven valuable over decades of mathematical biology research [6, 29, 31, 34–36], is with specific mathematical forms for the regulators containing 1-2 parameters (see equation in Figure 2B and Equations 5-12), and using these as terms in a coupled system of nonlinear ordinary differential equations. These functional forms can be derived from molecular interactions assuming mass-action kinetics plus, for example, multi-step irreversible processes [37] or intermediate-complex saturation [21, 38]. Graphs of these regulatory functions are shown in Figure 2B as dashed lines. We translate the arrow diagrams into models using these functional forms to represent the regulatory terms, and fit the parameters in these mathematical models to experimental data. Full equations and fitting algorithm are in Methods.

Before fitting to the full CD69 dataset, in this section we fit each model to a subset of the data corresponding to a specific on-time, out of the 6 on-times in the full dataset. (Thus, there are 4 models $\times$ 6 data subsets $= 24$ fits.) For each model and on-time data series, we measure the root-mean-square error (rms) and report these in Supplemental Figure S1 and Supplemental Figure S2. Examples are shown in the 3 columns of Figure 2C (dashed lines). Two of the individual on time series can be fit by these models to within the standard error of the mean (a representative is shown in Figure 2C right, dashed line), and 9 can be fit within the standard deviation of the data. Most (15 out of 24) of the models and on-time series cannot be fit (representatives shown in Figure 2C left and middle).

Yet, *in vivo*, the individual arrows in the phenotypic models of Figure 2 may not obey simple functional forms. Many biophysical phenomena (surface-confinement, mechanics, phase separation) imbue nonlinearity into individual molecular steps [12–18], and anyway the nodes in the arrow diagrams are coarse-grain

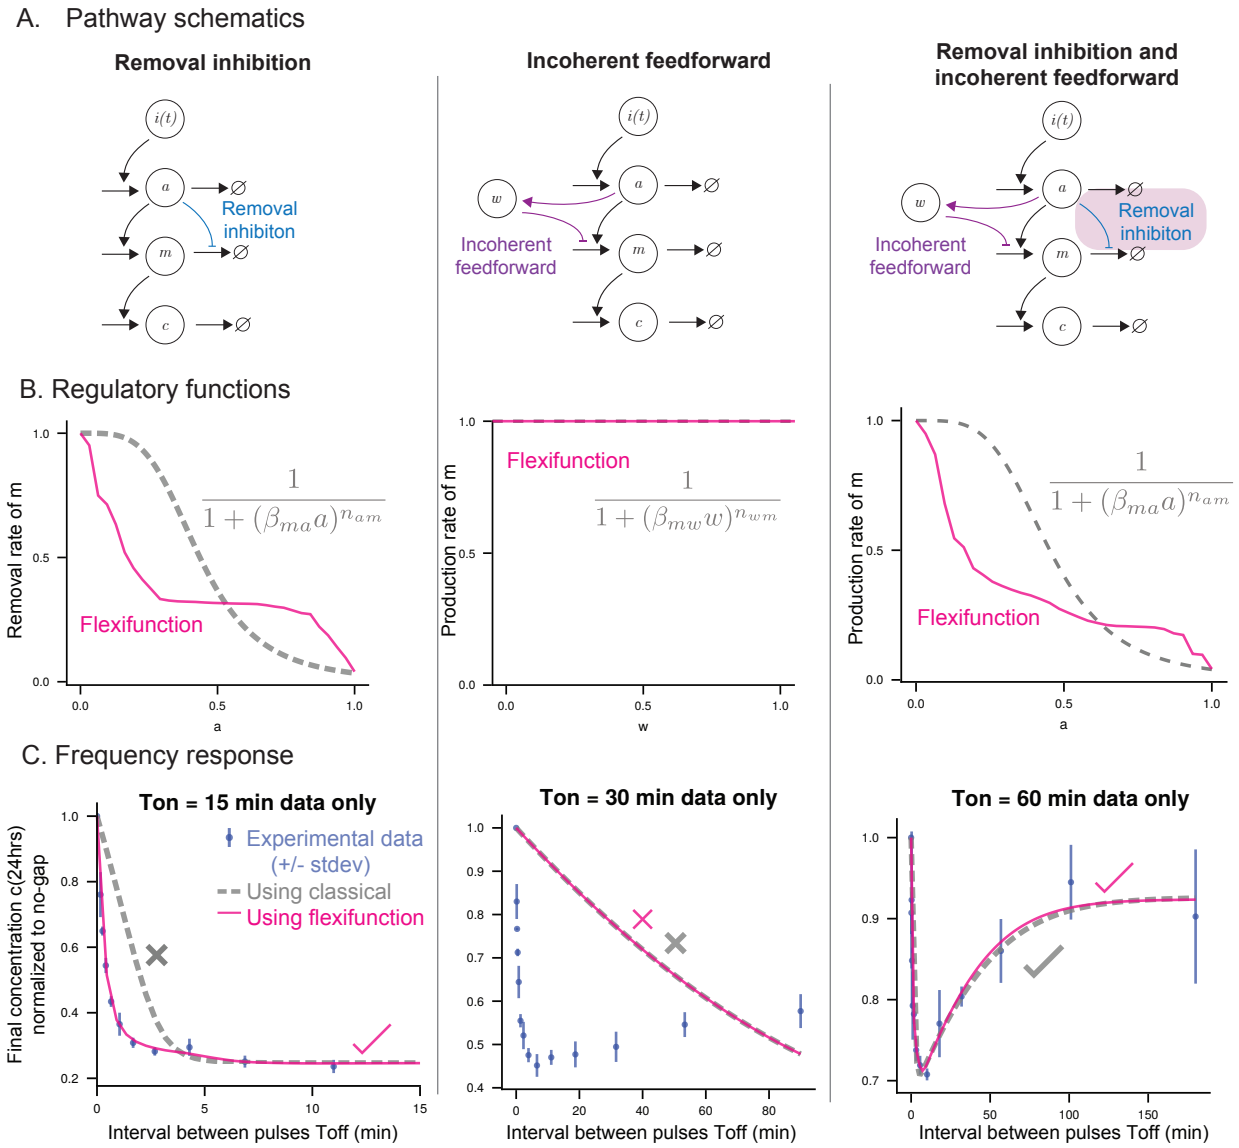

Figure 2: **Model fitting to subsets of CD69 frequency response data with and without model flexibility** A Model pathways described as arrow diagrams, where each arrow is restricted to be either a positive relationship or negative relationship. The highlighted arrows represent one component regulating another, and either obey a mathematical expression or are allowed to explore a universal continuous positive function restricted to either increase or decrease. B The regulatory functions corresponding to highlighted arrows. Dashed line is classical model obeying equation. Solid line is the universal function, which we term a "flexifunction" for convenience. C Frequency response with single on-time series for both classical and flexible regulator functions. There are 6 data subsets, 4 classical models and 5 fleximodels, so there are 24 classical models+dataset pairs and 30 fleximodel+dataset pairs. Representative examples are shown of for 3 on-times, and a corresponding classical and fleximodel fit (dashed and solid line, respectively). Error bars indicate between-replicate standard deviation of data from Harris et al. [10]. Checkmark or X indicate whether the root-mean-square discrepancy between model and single on-time series is above or below the experimental standard error of the mean. To see the root-mean-square fit of every model with every on-time subset, see Supplemental Figure S1.

representations of many molecular interactions. So, we developed a modeling approach in which the regulatory functions are represented as arbitrary functions with determined monotonicity, in this case, always decreasing as a function of their input, which we refer to as flexifunctions for convenience. For mathematical description and learning algorithm, see Methods. Note crucially that these models with flexifunctions correspond to the same arrow diagram as the classical models described above.

The rms performance of these fleximodels is shown in Supplemental Figure S1. We did not fit the linear pathway, since it was not even able to qualitatively fit the high-pass behavior. Since RI+IFF has two regulatory functions, either one or both could be represented as a flexifunction. Thus, there are 5 models $\times$ 6 data subsets $= 30$ fits. With this modest increase in modeling flexibility, 20 out of 30 agreed within the standard error of the mean, and 24 out of 30 were within the standard deviation. Representative examples are shown in the columns of Figure 2C (solid lines). There the left and right examples agree with the data, and the middle example fails to fit the middle data.

One potential concern is that adding flexibility to the model could lead to over-fitting, in which the model learns to capture experimental error in the specific training set, and would not generalize well to repeats or new data. The monotonic flexifunctions we employ are approximators of universal monotonic function (subject to the constraints described in Methods), and level of approximation is described by a parameter that controls the degrees of freedom in each flexifunction, n_{dofs}, defined in Methods. As the degrees of freedom increase, the flexifunction approaches the universal function. This allows us to study the possibility of over-fitting. We perform cross-validation with increasing flexibility in the flexifunctions, shown in Supplemental Figure S3 and discussed in Methods. As expected, increasing flexifunction flexibility, via degrees of freedom, at first improves performance, but after around 32 or 64 degrees of freedom per flexifunction, the performance is maintained on the training data but lost on the testing data, thus suggesting an optimal flexibility to avoid overfitting. We note two further lines of evidence suggesting that the flexibility is not leading to overfitting: First, many fleximodels failed to fit on-time series data (Figure 2C). Second, the performance in Supplemental Figure S3 plateaus at large degrees of freedom, even on the training data, demonstrating the structural constraints in the model. Both of those suggest the models are still sufficiently constrained to be used to reject hypotheses about the biological system.

RI+IFF arrow diagram model can quantitatively explain full CD69 frequency response data, by relaxing assumption of functional forms

Having developed the flexible modeling method, we next sought to fit the model to the full dataset, learning the regulatory functions directly from the data if necessary. This process can be visualized in the table in Figure 3A: From left to right, model complexity increases (more arrows, more components). From top to bottom, model complexity is constant (same arrow diagram), but flexibility of individual terms is increased. The first model to agree with the full on-time, off-time dataset is the RI+IFF model with both regulatory functions represented as flexifunctions. The regulatory functions and frequency response are shown in Figure 3B and C (solid lines), along with the best-performing equivalent classical model (dashed lines). The model has 4 biophysical parameters representing timescales of the 4 components, in addition to the flexifunctions. The parameters of the best-fit model are: $\tau_a = 2.45, \tau_m = 1.0, \tau_w = 129.4, \tau_c = 10^6$, in minutes. The timescale of τ_c attained its upper-bound and is much larger than the experimental timescale, so effectively $\tau_c \to \infty$.

Taken together, this demonstrates that a simple model may at first appear incapable of explaining complex data, but might indeed be able to do so, by relaxing functional form assumptions.

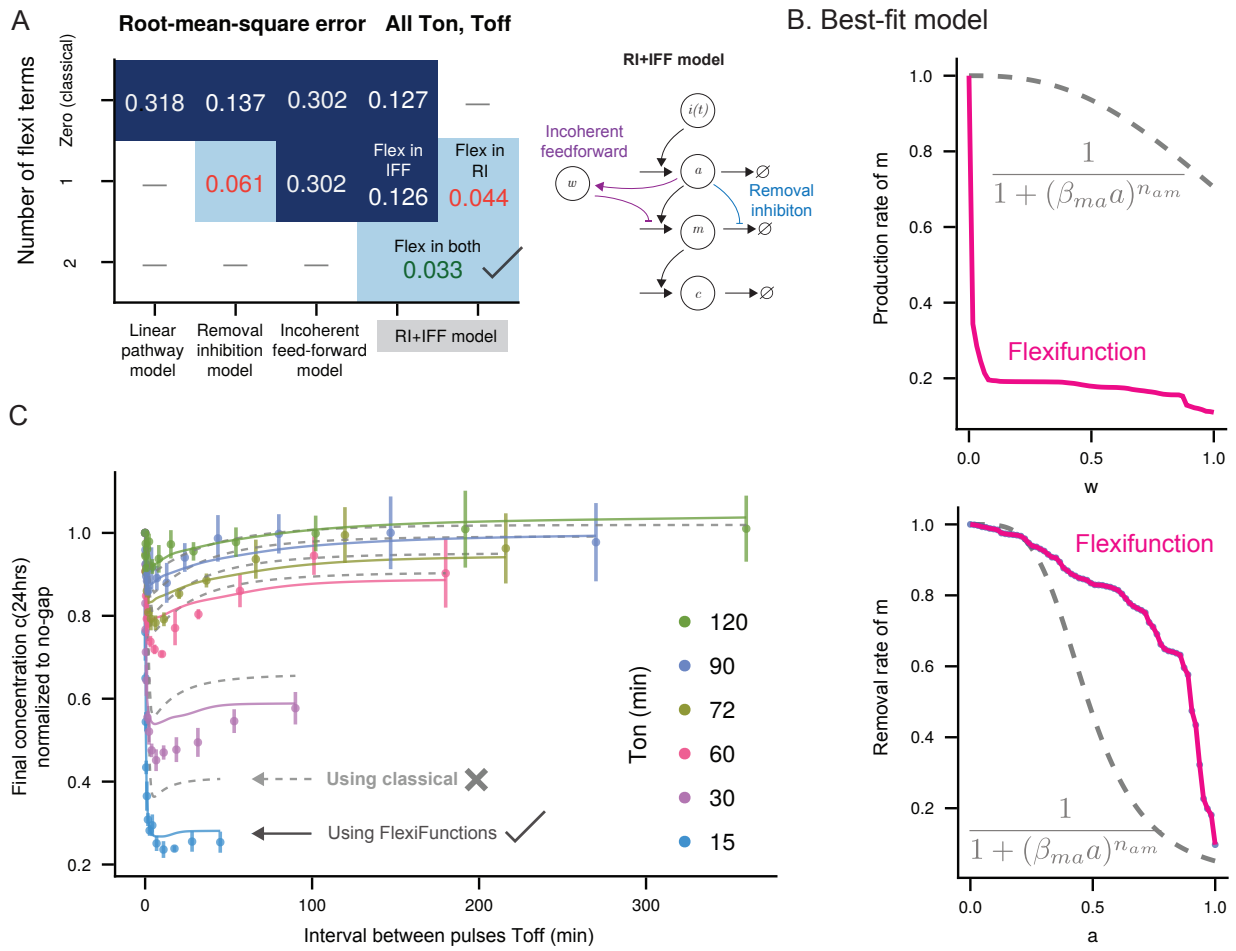

Figure 3: **A simple model with removal-inhibition and incoherent feedforward can fit full CD69 dataset, when flexibility is added to regulatory functions.** A Best-fit root-mean-square error when fitting to all on-times and off-times simultaneously. Four arrow diagram models, with either zero (classical), 1 or 2 flexi-function terms. Models with best-fit rms below the standard error of the mean shown in green with a checkmark. Models with best-fit rms below the standard deviation shown in orange. B The best-fit model and the only model we explored that has root-mean-square error below the experimental standard error of the mean is the model with both removal-inhibition and incoherent feed-forward, and both of those associated regulatory functions were allowed to be flexible. Regulator functions of best-fit RI-IFF model (solid line) deviate significantly from the best-fit classical regulatory functions. (Note that the other parameters have also changed between models.) C Best-fit frequency response. Error bars indicate between-replicate standard deviation from [10]. Solid lines are best-fit flexi model. Dashed lines are best-fit classical model.

Inspection of the regulatory functions suggests an explanation for why classical model failed to fit. The data exhibits a dramatic decrease in response for even very short off-times, when on-times are short, and yet only moderate decrease for longer off-times or on-times. To capture this, the refractory inhibitor gives dramatic inhibition after a small build-up in w, and only moderately more inhibition as more w accumulates. We include the full time-series for all on-times and off-times in **Supplemental Movie 1**. A notable feature is that the first pulse leads to much larger m than subsequent pulses, when on-time and off-time are small. This is consistent with the regulatory function and the data.

While the best-fit line has a root-mean-square error below the average standard error of the mean of all data, it does not run through every data point's error bars. Indeed, there are systematic mismatches, for example the steep dips in the data at low T_{off}. It is possible that these could be included in the same arrow diagram by replacing the $m \to c$ term with a flexifunction, but we did not attempt this, to maintain our strict commitment to the original desired threshold of goodness of fit.

Band-stop and low-pass arise in RI+IFF due to "roommate interrupts studying" and "icing a wound" phenomena

A widely-valued power of classical mathematical biology modeling is that their qualitative behaviors are often interpretable. We take a single on-time series that exhibits both band-stop (a local minimum in response versus T_{off}) and low-pass (an enhanced response for larger T_{off}, even above the no-gap stimulus), shown in Figure 4A.

First, band-stop behavior, or the decrease in response for moderate gaps between pulses. The m variable accumulates during stimulation, but decays rapidly when stimulus from a is removed due to its role in inhibiting removal. As a result, even short gaps in the stimulus prevent m from reaching its no-gap value (green arrow in Figure 4Bii). This harms the rate of increase in c. We refer to this as "roommate interrupts my studying" or "P.I. interrupts my coding" [39]: in this analogy, the intermediate quantity m is a latent ability to focus and study or write code.

Next, low-pass behavior, or, the increase in response for large gaps between pulses, even higher response than no-gap simulation. During no-gap stimulation, the refractory variable w accumulates gradually (blue curve in Figure 4Bi), limiting the steady rate of c accumulation. If there are gaps longer than the timescale of w, and the on-times are below the timescale of w accumulation, then w never has a chance to impact c to its full effect (blue arrow in Figure 4Biii), leading to a stronger response. This is a known feature of biological systems that exhibit adaptation [40] and negative self-regulation [41–43]. One example is offered by blood flow to a wound after an injury: common first-aid advice is to apply ice to a wound, but to do so intermitently. In this analogy, the refractory variable w is your body's homeostatis signal (compensatory vasodilation), which de-sensitizes you to the ice.

Note that both effects (band-stop arising from roommate-interrupts-my-studying and low-pass arising from adaptation) are features of the arrow diagram models, and arise in the classical mathematically-expressed models too.

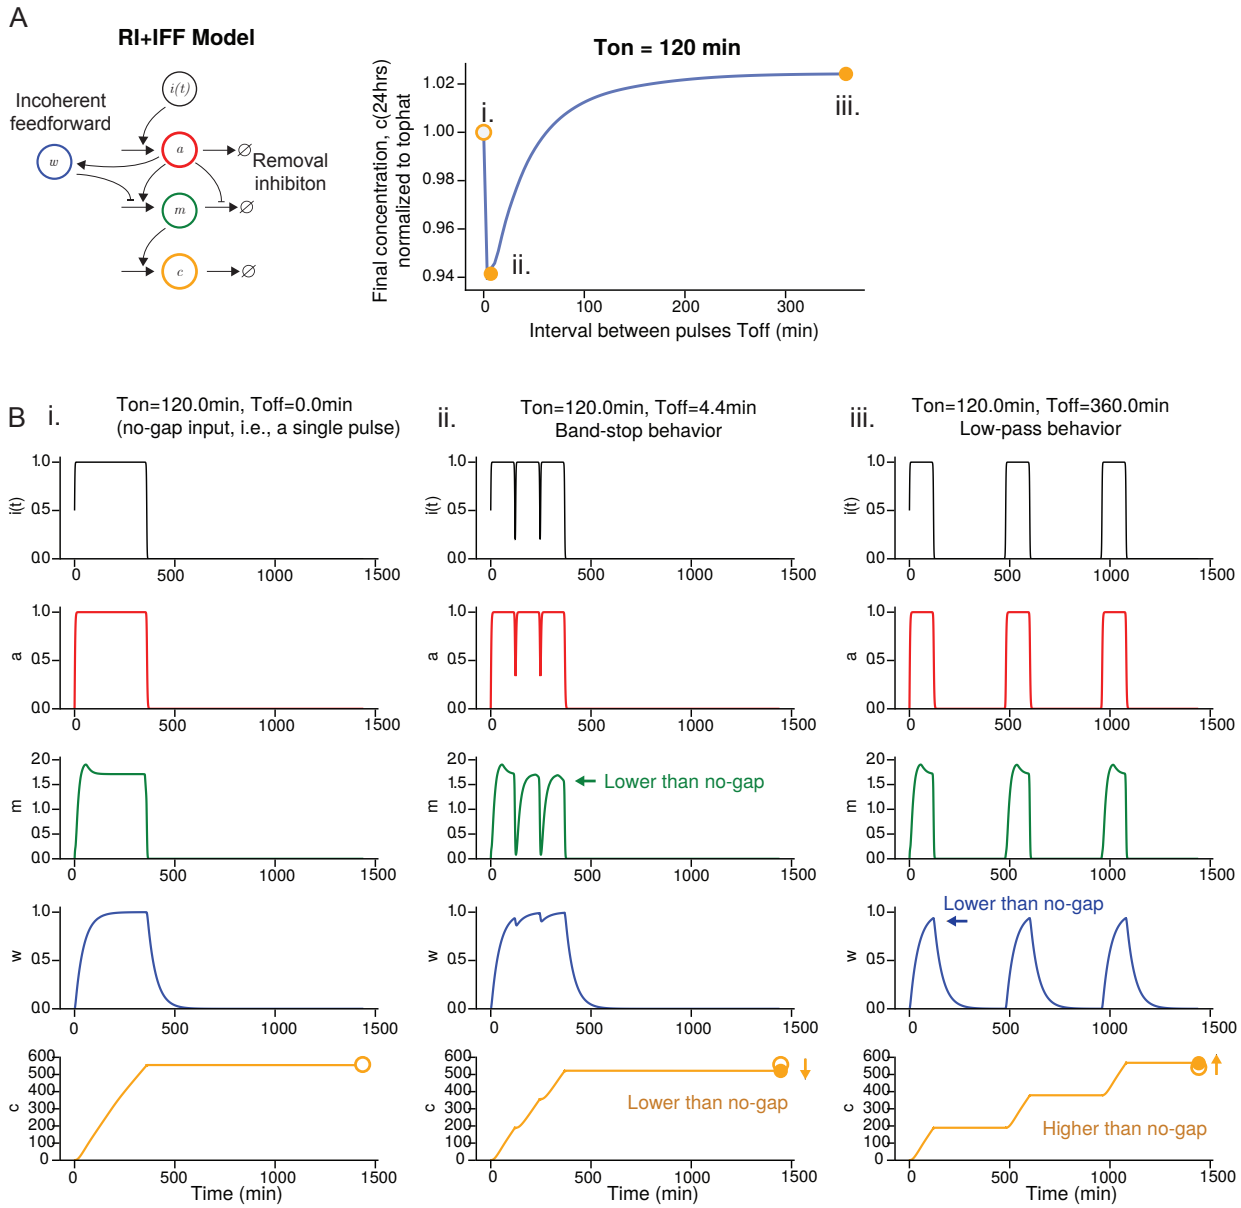

Figure 4: **Qualitative features in frequency response have simple heuristic explanations.** A Frequency response with single on-time series showing band-stop (the response is lower when there are intermediate gaps between pulses) and low-pass (the response is higher when there are large gaps between pulses). B Three selected time series, at off-times indicated by letters in (A). (i) tophat, i.e., one continuous pulse, (ii) intermediate gaps between pulses leading to lower response at 24 hours, (iii) long gaps between pulses leading to higher response at 24 hours.

Apparent band-pass arises with constant total stimulation time, in all models

The experimental protocol of Harris et al. [10] used a constant total stimulation time in all measurements, for technical reasons (to control for consistent damage accumulation). And, as mentioned above, the response is measured at a fixed time (24 hours after the first pulse begins). In this protocol, it is possible that larger gaps between pulses (T_{off}) give rise to stronger response — which we refer to as apparent band-pass — simply because there is less time between the last pulse and the measurement.

Indeed, even the linear pathway model can exhibit apparent band-pass under this stimulus protocol. This is shown in Figure 5A. Note that even though the linear pathway model can give apparent band-pass, it is still incapable of explaining the CD69 data, since the data exhibits both band-stop and band-pass in the same T_{on} data series, whereas the linear pathway model can only give apparent band-pass. To confirm that this is a feature of the protocol and not the model, we simulate the linear pathway model with a modified protocol where pulses continue until the measurement at 24 hours, in Figure 5C. Here, the model was incapable of increasing response strength.

This raises the question of whether the real system would exhibit a true band-pass response, if pulses continue until measurement. While an experiment with continuing pulses is challenging, we can use the trained cross-validated RI+IFF model to simulate the T Cell's response to continuing pulses. Interestingly, in Figure 5D, we find that there is predicted to be a local maximum at intermediate off-times, but this maximum does not attain a higher response than a continuous stimulus. Note that the exploration of hypothetical continuous pulses in this section does not invalidate the results of previous sections: linear pathway model still fails to explain the experimentally-observed data.

Discussion

The cross-validation we performed demonstrated that the model presented here — removal-inhibition and incoherent feedforward, with nonlinear regulatory function — can predict datapoints withheld from training. This suggests it may be more generally predictive. If so, the model may have the ability to aid in the design of stimulus schedules for optimal T cell responses. For example, it could be used to design a stimulus schedule that maximes activation, or minimizes activation, or maximizes it over a particular period before minimizing it. Similar predictive mathematical models have contributed to the design of schedules for administering cancer treatement drugs [43] and have been suggested for T-cell-based therapies [33].

The phenotypic model revealed by pulsatile stimulus contains two common motifs: removal-inhibition of an intermediate component, and an incoherent feed-forward loop. Removal-inhibition provides the ability to respond rapidly to changes [31]. This motif gives rise to band-stop behavior, which has been suggested to allow T cells to filter out interactions with healthy cells, to allow more effort to focus on its targets [11]. Incoherent feed-forward motifs are found throughout biology [31, 32] and endow systems with adaptation: they become desensitized following prolonged exposure, effectively redefining their baseline.

When a model fails to fit every feature of a dataset, including small quantitative features that do not change the categorical behavior, it is tempting to add model complexity. The modeling approach presented here allows the same, interpretable arrow diagram model to be given additional quantitative flexibility to be learned from data. The method is similar in spirit to a class of approaches in which the right-hand-side of a dynamical model is replaced with a universal function [25–27], with a few key distinctions. First is our viewpoint that monotonic-increasing or monotonic-decreasing connections between nodes are desirable for interpretability. Second is that all the universal functions we used were single-input, single-output. Indeed,

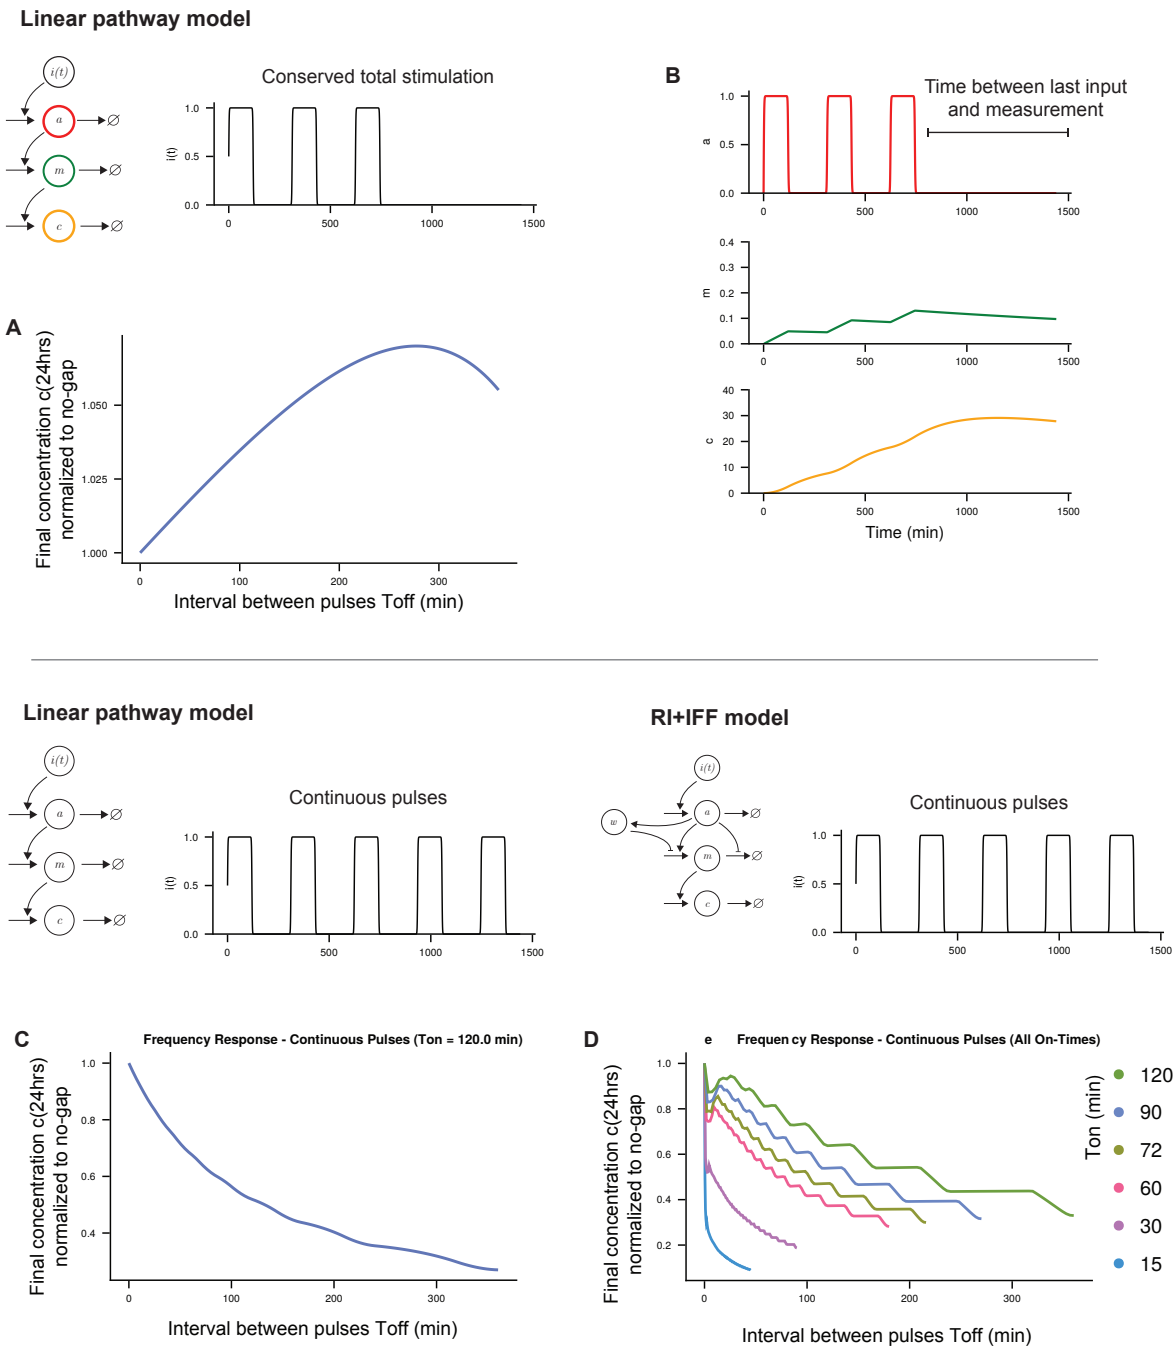

Figure 5: All models including the minimal model with a linear pathway can exhibit an increase in response when $T_{\text{off}} > 0$, compared to a single no-gap pulse. A Linear arrow archicture, with constant total stimulation (as in experimental data in Harris et al. [10]). B Given constant total stimulation time, longer gaps imply shorter time until the measurement, leading to increase in measured response. C Linear arrow architecture, but with continuous pulsatile input. D The RI+IFF best-fit model, but with continuous pulsatile input.

the approach can adapt to the specific project's definition of interpretability. The approach could be repeated with non-monotone regulatory functions — if your objective of interpretability were simply the connections between components. Or, it could be repeated with more constrained functions, for example functions that have log-derivatives no larger than 1.0, ensuring no arrow exhibits ultrasensitivity, and any ultrasensitive response arises from the arrow diagram itself, not individual arrows. This is connected to another distinction between our approach and similar approaches, which is our use of piece-wise linear approximation to represent the universal function. Imposing the monotonicity constraint proved most computationally efficient for the piecewise-linear universal function approximator (Equation 17) compared to Chebyshev expansions, Fourier series, Taylor series, and multi-layer perceptrons. Indeed, early versions of our code used multi-layer perceptrons but the piece-wise linear version provides a computationally cheaper and, in our view, more elegant implementation for this specific application.

In this work, we manually explored the space of possible arrow-diagrams, represented visually as moving left-to-right in the table in Figure 3A. More sophisticated methods exist to explore the set of arrow-diagrams [4, 29, 30, 44, 45]. For this dataset, a suprisingly simple model proved able to fit the data with the desired root-mean-square residual, making manual, ad-hoc model exploration feasible. However, future data sets may necessitate combining the flexifunction approach with automated exploration of the space of arrow-diagrams.

Methods

All code to define, simulate, and learn models is provided at:

`https://github.com/allardjun/PulsatileModelLearning`

Models

The mathematical models were scaled (non-dimensionalized) in quantity value, while time was kept in physical units of minutes. In all cases, $i(t)$ represents the input signal of optically-activated TCR, either 0 or 1 at any time t. We identify $c(t)$ as the surface expression of CD69, in nondimensional units so that only comparisons of $c(t)$ are meaningful. Model equations were constrained so that if $i(t) = 0$ then the steady-state $c(t) = 0$ in agreement with Harris et al. [10]. When rescaling, there are two mathematically equal choices of whether the τ_i multiplies the full right-hand side or just the linear decay term. We found that when fitting data, τ_c was larger than the other τ_i's, so for numerical integration stability and accuracy, we rescaled c with τ_c as shown below.

Linear architecture

$$\frac{da}{dt} = \frac{1}{\tau_a}\left(i(t) - a\right) \tag{1}$$

$$\frac{dm}{dt} = \frac{1}{\tau_m}\left(a - m\right) \tag{2}$$

$$\frac{dc}{dt} = m - \frac{c}{\tau_c} \tag{3}$$

Parameters: τ_a, τ_m, τ_c.

Removal-Inhibition (RI)

$$\frac{da}{dt} = \frac{1}{\tau_a}\left(i(t) - a\right) \tag{4}$$

$$\frac{dm}{dt} = \frac{1}{\tau_m}\left(a - \frac{1}{1 + (\beta_{ma} a)^{n_{am}}} m\right) \tag{5}$$

$$\frac{dc}{dt} = m - \frac{c}{\tau_c} \tag{6}$$

Parameters: τ_a, τ_m, τ_c, β_{ma}, n_{am}.

Incoherent Feed-forward (IFF)

$$\frac{da}{dt} = \frac{1}{\tau_a}\left(i(t) - a\right) \tag{7}$$

$$\frac{dm}{dt} = \frac{1}{\tau_m}\left(\frac{a}{1 + (\beta_{mw}(w))^{n_{wm}}} - m\right) \tag{8}$$

$$\frac{dw}{dt} = \frac{1}{\tau_w}\left(a - w\right) \tag{9}$$

$$\frac{dc}{dt} = m - \frac{c}{\tau_c} \tag{10}$$

Parameters: τ_a, τ_m, τ_c, τ_w, β_{mw}, n_{wm}.

Removal-Inhibition and Incoherent Feedforward Loop (RI+IFFL)

$$\frac{da}{dt} = \frac{1}{\tau_a}\left(i(t) - a\right) \tag{11}$$

$$\frac{dm}{dt} = \frac{1}{\tau_m}\left(\frac{a}{1 + (\beta_{mw}w)^{n_{wm}}} - \frac{1}{1 + (\beta_{ma}a)^{n_{am}}}m\right) \tag{12}$$

$$\frac{dw}{dt} = \frac{1}{\tau_w}\left(a - w\right) \tag{13}$$

$$\frac{dc}{dt} = m - \frac{c}{\tau_c} \tag{14}$$

Parameters: τ_a, τ_m, τ_c, τ_w, β_{mw}, n_{wm}, β_{ma}, n_{am}.

Generation of frequency response plot for comparison with data

The experimental data from Harris et al. [10] consists of CD69 protein expression measurements in engineered T cells subjected to pulsatile optical stimulation with varying pulse durations ($T_{\text{on}} = 15$, 30, 60, 72, 90, 120 minutes) and interpulse intervals ($T_{\text{off}} = 0$–480 minutes).

To match these conditions, we implement a pulsatile input function that generates time-dependent stimulation patterns with constant time-integral across different pulsing protocols. This constraint ensures that different pulse/interval combinations deliver equivalent total stimulation, matching the experimental design where pulse amplitude was adjusted to maintain constant integrated dose. Each experimental condition ($6 \times 14 = 84$ total conditions) is simulated by integrating the ODE system with the corresponding pulsatile input, producing model predictions for the 24-hour CD69 readout that can be directly compared to experimental measurements.

Frequency response analysis provides quantitative comparison between model predictions and experimental data across the full parameter space of pulsing conditions. For each model and parameter set, we systematically simulate all 84 experimental conditions by solving the ODE system with appropriate pulsatile inputs using high-precision numerical integration (Vern9 algorithm with adaptive tolerance).

The final output concentration $c(24\text{h})$ serves as the model prediction for CD69 expression at the experimental endpoint. To match experimental normalization, all model outputs are normalized to the first off-time condition within each on-time series:

$$\text{c24_normalized}[i,j] = \frac{\text{c24}[i,j]}{\text{c24}[i,1]}. \tag{15}$$

This produces a 6×14 frequency response matrix that directly corresponds to the experimental data structure. Model-data comparison is visualized through heatmaps and line plots showing normalized CD69 expression as a function of interpulse interval for each pulse duration. The root-mean-square (rms) error between model and data serves as the primary metric for model performance:

$$\text{rms} = \sqrt{\frac{\sum(\text{model} - \text{data})^2}{6 \times 14}}, \tag{16}$$

providing normalized error estimates across all experimental conditions.

Flexifunctions implementation

Flexifunctions provide a data-driven approach to extend classical mechanistic models by replacing classical functions with learnable functions. Each flexifunction is parameterized by a n_{dofs} parameters, $\theta_1..\theta_{n\text{dofs}}$, that define breakpoints for monotonic interpolation over the unit interval $[0,1]$.

The function is built using cumulative sums of squared parameters,

$$\Theta_j = \sum_{k=1}^{j} \theta_k^2 \quad \text{for } j = 1, 2, \ldots, N \tag{17}$$

and $\Theta_0 = 0$. This guarantees the Θ_j values are increasing and have zero-intercept. The piecewise linear interpolation occurs between evenly-spaced adjacent cumulative values,

$$f_{\text{increasing}}(x) = \Theta_{i-1} + (N \cdot x - (i-1)) \cdot (\Theta_i - \Theta_{i-1}) \tag{18}$$

where $i = \min(\lfloor N \cdot x \rfloor + 1, N)$ is the interval index. Note the optimization algorithm can freely adjust $\theta_k \in \mathbb{R}$ without violating the constraints (increasing and zero-intercept). To obtain a continuous, strictly-positive, decreasing function with intercept at unity, we take

$$\text{flex}(x) = \frac{1}{1 + f_{\text{increasing}}(x)}. \tag{19}$$

Thus, these functions satisfy

$$\text{flex}\,(0) = 1, \tag{20}$$

$$\text{flex}\,(x) > 0 \tag{21}$$

and

$$\frac{\partial \text{flex}}{\partial x} < 0. \tag{22}$$

As $n_{\text{dofs}} \to \infty$, the piece-wise linear functions captured in this basis approach the set of arbitrary decreasing functions subject to Equations 20-22. These monotonic-increasing, zero-intercept functions are then inserted into the right-hand side of differential equations, for example for the RI+IFF model,

$$\begin{aligned}
\frac{da}{dt} &= \frac{1}{\tau_a}\left(i(t) - a\right) \\
\frac{dm}{dt} &= \frac{1}{\tau_m}\left(a \cdot \text{flex1}\left(w\right) - \text{flex2}\left(a\right) \cdot m\right) \\
\frac{dw}{dt} &= \frac{1}{\tau_w}\left(a - w\right) \\
\frac{dc}{dt} &= m - \frac{c}{\tau_c}
\end{aligned}$$

Unless otherwise noted, flex1 and flex2 each contain n_{dofs} parameters, where $n_{\text{dofs}} = 64$ as determined by cross-validation, see below.

Learning Algorithm

The framework computes rms errors compared to experimental between-replicate variability (standard deviation for 3 replicates = 0.0638). The codebase uses a dual parameter representation system to ensure all optimization parameters are well-behaved across the entire real number line $(-\infty, +\infty)$. Before parameters are visible to the optimization algorithms, they are transformed into optimization-friendly "represented" parameters using log and square-root transformations. After optimization, they are "derepresented" back to physically meaningful values that appear in the ODEs above.

Learning Algorithm for Classical Models

Classical models are optimized using global optimization methods. The primary algorithm is the adaptive differential evolution variant `BBO_adaptive_de_rand_1_bin()` from the `Optimization.jl` Julia package, which performs global parameter space exploration for the 3–8 classical parameters.

The algorithm converges by 2.5×10^4 iterations (Supplemental Figure S4).

Learning Algorithm for Models Containing flexifunctions

Flexi models present a significantly more challenging optimization landscape due to the high-dimensional parameter space (20–256 flexifunction parameters plus classical parameters). The learning protocol implements a three-stage "corduroy" optimization strategy that alternates between different parameter subsets to manage complexity and avoid local optima. An example learning trajectory is shown in Supplemental Figure S4.

First, we run the corresponding classical model as above to generate an initial guess. The corduroy learning protocol then alternates between CMA-ES optimization of flexifunction parameters (holding classical parameters fixed) and a simplex refinement (Rowan's Subplex algorithm [46]) of classical parameters (holding flexifunctions fixed). CMA-ES (Covariance Matrix Adaptation Evolution Strategy) is particularly well-suited for flexifunction optimization due to its adaptive covariance matrix that captures parameter correlations and its robust performance in high-dimensional spaces. Interestingly, the simplex iterations decrease

the loss only slightly (not visible in Supplemental Figure S4), and yet are sufficient to allow the following CMA-ES iterations to discover a significant improvement, even though the previous CMA-ES round of iterations had converged.

The algorithm maintains adaptive step-size control (σ_0) computed based on parameter bounds and implements custom hyperparameter schedules for population size and termination criteria. Multiple corduroy rounds ensure convergence, with each iteration building upon improved parameter estimates from the previous round. This approach effectively decomposes the complex joint optimization into manageable subproblems while maintaining global search capabilities for the high-dimensional flexifunction parameter space.

We find that 120 CMA-ES iterations, 200 simplex iterations, each repeated 5 times leads to unchanging solutions and appears converged (Supplemental Figure S4). This procedure takes around 72 hours on Intel Xeon Gold 6230 with 8GB of memory.

Note that the optimization function is differentiable, but not convex, so gradient descent methods fail to improve the fit.

Cross-validation and degrees-of-freedom selection

Flexifunction complexity is determined through systematic cross-validation to balance model expressiveness against overfitting. We implement k-fold cross-validation with $k = 7$ (since the number of data points is divisible by 7). This partitions the 84 experimental conditions into randomly-selected but non-overlapping training and validation sets, across all pulsing conditions. For each n_{dofs} value (2, 4, 8, 16, 32, 64, 128, 256), for each of the 7 random subsets, we train flexi models on 72 conditions and evaluate performance on the remaining 12 conditions, rotating through all fold combinations. The validation metric is RMS error computed on held-out data. This process generates learning curves showing training and validation error as functions of model complexity, with the optimal n_{dofs} selected as the point where validation error is minimized before overfitting occurs.

As shown in Supplemental Figure S3, the fit quality generally improved as more degrees of freedom were added to the flexifunctions. Then, the root-mean-square error was measured on the withheld subset. The error in the testing subset exhibits a pronounced minimum around $n_{\text{dofs}} = 32$ degrees of freedom. Note importantly that adding more degrees of freedom did not allow unlimited improvements to the fit even on the training set, but rather plateaus at large n_{dofs}. This suggests that these models have structural constraints that would allow a model to be rejected, even with infinite degrees of freedom.

Acknowledgements

We thank Omer Dushek, Brady Berg, Katie Bogue, Jack Corrette, Austin Marcus for useful discussion. This work was supported by NSF DMS 2052668.

References

[1] Jordan Kramer, P Anton van der Merwe, and Omer Dushek. Platforms for studying cell–cell recognition by immune cells. *Immunology and Cell Biology*, 103: 636–647, 2025.

[2] Melissa Lever, Philip K. Maini, P. Anton van der Merwe, and Omer Dushek. Phenotypic models of T cell activation. *Nature Reviews Immunology*, 14(9):619–629, 2014.

[3] Sahak Z Makaryan and Stacey D Finley. Enhancing network activation in natural killer cells: predictions from in silico modeling. *Integrative Biology*, 12(5):109–121, 2020.

[4] Melissa Lever, Hong-Sheng Lim, Philipp Kruger, John Nguyen, Nicola Trendel, Enas Abu-Shah, Philip Kumar Maini, Philip Anton van der Merwe, and Omer Dushek. Architecture of a minimal signaling pathway explains the T-cell response to a 1 million-fold variation in antigen affinity and dose. *Proceedings of the National Academy of Sciences*, 113(43):E6630–E6638, 2016.

[5] Kenneth Y. Wertheim, Bhanwar Lal Puniya, Alyssa La Fleur, Ab Rauf Shah, Matteo Barberis, and Tomáš Helikar. A multi-approach and multi-scale platform to model CD4+ T cells responding to infections. *PLoS Computational Biology*, 17(8):e1009209, 2021.

[6] Taisuke Kondo, François X.P. Bourassa, Sooraj Achar, Justyn DuSold, Pablo F. Céspedes, Makoto Ando, Alka Dwivedi, Josquin Moraly, Christopher Chien, Saliha Majdoul, Adam L. Kenet, Madison Wahlsten, Audun Kvalvaag, Edward Jenkins, Sanghyun P. Kim, Catherine M. Ade, Zhiya Yu, Guillaume Gaud, Marco Davila, Paul Love, James C. Yang, Michael L. Dustin, Grégoire Altan-Bonnet, Paul François, and Naomi Taylor. Engineering TCR-controlled fuzzy logic into CAR T cells enhances therapeutic specificity. *Cell*, 188(9):2372–2389.e35, 2025.

[7] Ashna Patel, Violaine Andre, Sofia Bustamante Eguiguren, Michael I Barton, Jake Burton, Eleanor M Denham, Johannes Pettmann, Alexander M Mørch, Mikhail A Kutuzov, Jesús A Siller-Farfán, Michael L Dustin, P Anton van der Merwe, and Omer Dushek. Using CombiCells, a platform for titration and combinatorial display of cell surface ligands, to study T-cell antigen sensitivity modulation by accessory receptors. *The EMBO Journal*, 43(1):132–150, 2024.

[8] John Bechhoefer. *Control Theory for Physicists*. Cambridge University Press, 2021.

[9] Candela L. Szischik, Juliana Reves Szemere, Rocío Balderrama, Constanza Sánchez de la Vega, and Alejandra C. Ventura. Transient frequency preference responses in cell signaling systems. *npj Systems Biology and Applications*, 10(1):86, 2024.

[10] Michael J Harris, Muna Fuyal, and John R James. Quantifying persistence in the T-cell signaling network using an optically controllable antigen receptor. *Molecular Systems Biology*, 17(5):e10091, 2021.

[11] Geoff P. O'Donoghue, Lukasz J. Bugaj, Warren Anderson, Kyle G. Daniels, David J. Rawlings, and Wendell A. Lim. T cells selectively filter oscillatory signals on the minutes timescale. *Proceedings of the National Academy of Sciences*, 118(9):e2019285118, 2021.

[12] Eva M. Schmid, Matthew H. Bakalar, Kaushik Choudhuri, Julian Weichsel, Hyoung Sook Ann, Phillip L. Geissler, Michael L. Dustin, and Daniel A. Fletcher. Size-dependent protein segregation at membrane interfaces. *Nature Physics*, 12(7):704–711, 2016.

[13] Robert Taylor, Jun Allard, and Elizabeth L. Read. Simulation of receptor triggering by kinetic segregation shows role of oligomers and close contacts. *Biophysical Journal*, 121(9):1660–1674, 2022.

[14] Ying Zhang, Lara Clemens, Jesse Goyette, Jun Allard, Omer Dushek, and Samuel A. Isaacson. The Influence of Molecular Reach and Diffusivity on the Efficacy of Membrane-Confined Reactions. *Biophysical Journal*, 117(7):1189–1201, 2019.

[15] Lara Clemens, Mikhail Kutuzov, Kristina Viktoria Bayer, Jesse Goyette, Jun Allard, and Omer Dushek. Determination of the molecular reach of the protein tyrosine phosphatase SHP-1. *Biophysical Journal*, 120(10):2054–2066, 2021.

[16] Jesse Goyette, David Depoil, Zhengmin Yang, Samuel A. Isaacson, Jun Allard, P. Anton van der Merwe, Katharina Gaus, Michael L. Dustin, and Omer Dushek. Dephosphorylation accelerates the dissociation of ZAP70 from the T cell receptor. *Proceedings of the National Academy of Sciences*, 119 (9):e2116815119, 2022.

[17] Kai Liu, Brian Chu, Jay Newby, Elizabeth L. Read, John Lowengrub, and Jun Allard. Hydrodynamics of transient cell-cell contact: The role of membrane permeability and active protrusion length. *PLoS Computational Biology*, 15(4):e1006352, 2019.

[18] Lara Clemens, Omer Dushek, and Jun Allard. Intrinsic Disorder in the T Cell Receptor Creates Cooperativity and Controls ZAP70 Binding. *Biophysical Journal*, 120(2):379–392, 2021.

[19] Alexis N Prybutok, Jason Y Cain, Joshua N Leonard, and Neda Bagheri. Fighting fire with fire: deploying complexity in computational modeling to effectively characterize complex biological systems. *Current Opinion in Biotechnology*, 75:102704, 2022.

[20] Chaopeng Shen, Alison P. Appling, Pierre Gentine, Toshiyuki Bandai, Hoshin Gupta, Alexandre Tartakovsky, Marco Baity-Jesi, Fabrizio Fenicia, Daniel Kifer, Li Li, Xiaofeng Liu, Wei Ren, Yi Zheng, Ciaran J. Harman, Martyn Clark, Matthew Farthing, Dapeng Feng, Praveen Kumar, Doaa Aboelyazeed, Farshid Rahmani, Yalan Song, Hylke E. Beck, Tadd Bindas, Dipankar Dwivedi, Kuai Fang, Marvin Höge, Chris Rackauckas, Binayak Mohanty, Tirthankar Roy, Chonggang Xu, and Kathryn Lawson. Differentiable modelling to unify machine learning and physical models for geosciences. *Nature Reviews Earth & Environment*, 4(8):552–567, 2023.

[21] Ben Noordijk, Monica L. Garcia Gomez, Kirsten H. W. J. ten Tusscher, Dick de Ridder, Aalt D. J. van Dijk, and Robert W. Smith. The rise of scientific machine learning: a perspective on combining mechanistic modelling with machine learning for systems biology. *Frontiers in Systems Biology*, 4: 1407994, 2024.

[22] Rose Yu and Rui Wang. Learning dynamical systems from data: An introduction to physics-guided deep learning. *Proceedings of the National Academy of Sciences*, 121(27):e2311808121, 2024.

[23] John H. Lagergren, John T. Nardini, Ruth E. Baker, Matthew J. Simpson, and Kevin B. Flores. Biologically-informed neural networks guide mechanistic modeling from sparse experimental data. *PLoS Computational Biology*, 16(12):e1008462, 2020.

[24] Steven L. Brunton, Joshua L. Proctor, and J. Nathan Kutz. Discovering governing equations from data by sparse identification of nonlinear dynamical systems. *Proceedings of the National Academy of Sciences*, 113(15):3932–3937, 2016.

[25] Christopher Rackauckas, Yingbo Ma, Julius Martensen, Collin Warner, Kirill Zubov, Rohit Supekar, Dominic Skinner, Ali Ramadhan, and Alan Edelman. Universal Differential Equations for Scientific Machine Learning. *arXiv* 2001.04385, 2020.

[26] Maren Philipps, Nina Schmid, and Jan Hasenauer. Universal differential equations for systems biology: Current state and open problems. *bioRxiv* 2024.11.29.626122, 2024.

[27] Colby Fronk, Jaewoong Yun, Prashant Singh, and Linda Petzold. Bayesian polynomial neural networks and polynomial neural ordinary differential equations. *PLOS Computational Biology*, 20(10): e1012414, 2024.

[28] Nazanin Ahmadi Daryakenari, Mario De Florio, Khemraj Shukla, and George Em Karniadakis. AI-Aristotle: A physics-informed framework for systems biology gray-box identification. *PLOS Computational Biology*, 20(3):e1011916, 2024.

[29] Jonathan Rodriguez, Abdon Iniguez, Nilamani Jena, Prasanthi Tata, Zhong-Ying Liu, Arthur D Lander, John Lowengrub, and Richard A Van Etten. Predictive nonlinear modeling of malignant myelopoiesis and tyrosine kinase inhibitor therapy. *eLife*, 12:e84149, 2023.

[30] Justin D. Finkle, Jia J. Wu, and Neda Bagheri. Windowed Granger causal inference strategy improves discovery of gene regulatory networks. *Proceedings of the National Academy of Sciences*, 115(9): 2252–2257, 2018.

[31] U. Alon. *An Introduction to Systems Biology: Design Principles of Biological Circuits*. CRC Press, 2020.

[32] Claus Kadelka, Matthew Wheeler, Alan Veliz-Cuba, David Murrugarra, and Reinhard Laubenbacher. Modularity of biological systems: a link between structure and function. *Journal of the Royal Society Interface*, 20(207):20230505, 2023.

[33] Can Liu, Timothy Qi, J. Justin Milner, Yong Lu, and Yanguang Cao. Speed and Location Both Matter: Antigen Stimulus Dynamics Controls CAR-T Cell Response. *Frontiers in Immunology*, 12:748768, 2021.

[34] Leah Edelstein-Keshet. Mathematical Models in Biology. 2005.

[35] James Keener and James Sneyd. Mathematical Physiology, I: Cellular Physiology. *Interdisciplinary Applied Mathematics*, 2009.

[36] Yixuan Wang, Daniel R Bergman, Erica Trujillo, Alexander T. Pearson, Randy F. Sweis, and Trachette L. Jackson. Mathematical model predicts tumor control patterns induced by fast and slow cytotoxic T lymphocyte killing mechanisms. *Scientific Reports*, 13(1):22541, 2023.

[37] T W McKeithan. Kinetic proofreading in T-cell receptor signal transduction. *Proceedings of the National Academy of Sciences*, 92(11):5042–5046, 1995.

[38] Rob Phillips, Jane Kondev, Julie Theriot, Hernan G. Garcia, and Nigel Orme. Physical Biology of the Cell. 2012.

[39] Gloria Mark, Daniela Gudith, and Ulrich Klocke. The cost of interrupted work. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 107–110, 2008.

[40] Mustafa H. Khammash. Perfect adaptation in biology. *Cell Systems*, 12(6):509–521, 2021.

[41] Y. Li and A. Goldbeter. Frequency specificity in intercellular communication. Influence of patterns of periodic signaling on target cell responsiveness. *Biophysical Journal*, 55(1):125–145, 1989.

[42] Y. Li and A. Goldbeter. Pulsatile signaling in intercellular communication. Periodic stimuli are more efficient than random or chaotic signals in a model based on receptor desensitization. *Biophysical Journal*, 61(1):161–171, 1992.

[43] Philipp M. Altrock, Lin L. Liu, and Franziska Michor. The mathematics of cancer: integrating quantitative models. *Nature Reviews Cancer*, 15(12):730–745, 2015.

[44] Maria Secrier, Tina Toni, and Michael P. H. Stumpf. The ABC of reverse engineering biological signalling systems. *Molecular BioSystems*, 5(12):1925–1935, 2009.

[45] Tina Toni, David Welch, Natalja Strelkowa, Andreas Ipsen, and Michael P.H Stumpf. Approximate Bayesian computation scheme for parameter inference and model selection in dynamical systems. *Journal of The Royal Society Interface*, 6(31):187–202, 2009.

[46] T Rowan. The subplex method for unconstrained optimization. *PhD thesis, PhD thesis, Department of Computer Sciences, University of Texas*, 1990.

1 Supplemental material

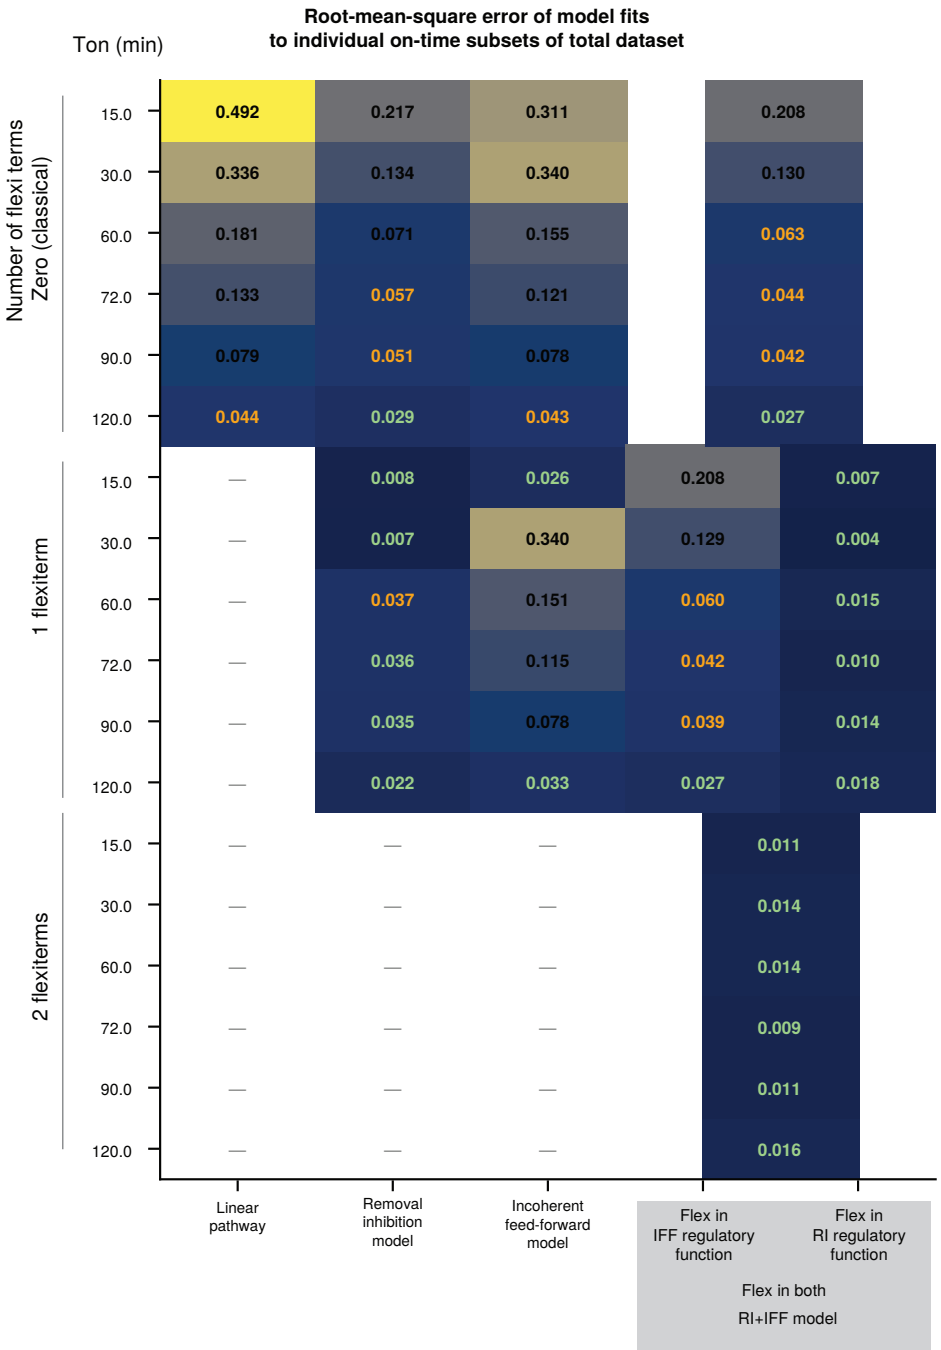

Figure S1: Best-fit rms error of models, fit to individual on-time series, for models with classical mathematical biology functional forms, and 1 or 2 flexifunctions. Text font indicates whether the fit root-mean-square error was smaller than the experimental data's standard deviation (orange) or smaller than the experimental data's standard error of the mean (green). Note that the RI+IFF model, which has both removal-inhibition and incoherent-feedforward regulatory functions, there are two options for where to put a single flexifunction.

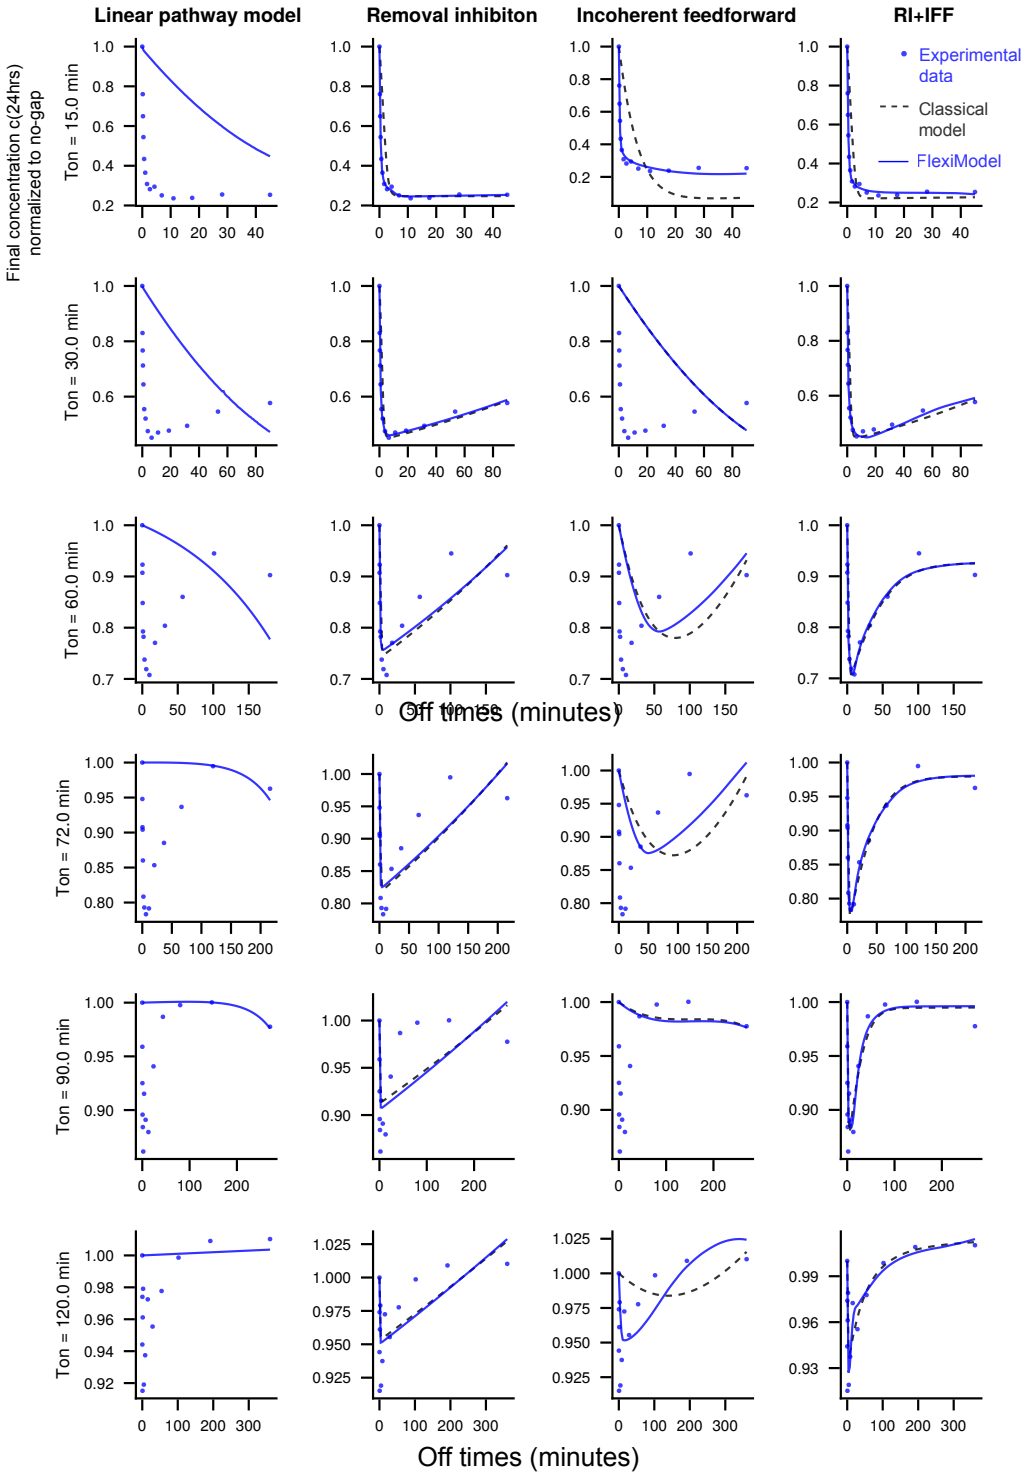

Figure S2: Best-fit rms error of models, fit to individual on-time series, for models with classical mathematical biology functional forms (for 4 models), and 1 flexifunction (for 3 models). These correspond to a subset of rms errors reported in Supplemental Figure S1.

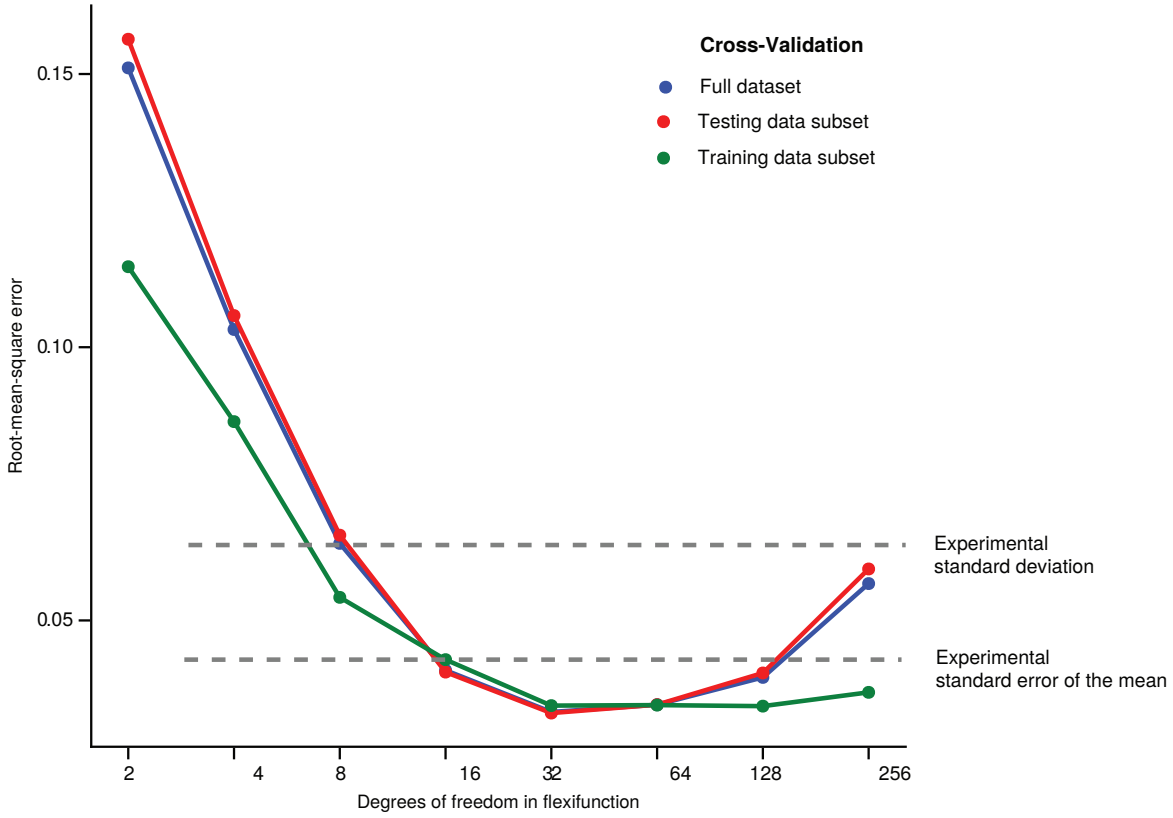

Figure S3: Cross-validation analysis allows choice of degrees of freedom in flexifunctions to avoid over-fitting. RI+IFF model (two flexifunctions) was fit to subsets of the full dataset, specifically 7 subsets were generated where 14 datapoints were removed without replacement (7 was chosen since the total number of datapoints was $7x14$). The fit quality, measured by root-mean-square error to the training set (green), generally improved as more degrees of freedom were added to the flexifunctions. Then, the root-mean-square error was measured on the removed subset (testing data subset, red). The error in the testing subset exhibits a pronounced minimum around 32 degrees of freedom. Note importantly that adding more degrees of freedom did not allow unlimited improvements to the fit to the training set — this suggests that these models have structural constraints that would allow a model to be rejected, even with infinite degrees of freedom. The error on the training data is slightly above the testing data for 2 points, which may be due to stochastic optimization. The slight increase in error on the training data at very large degrees of freedom may indicate the algorithm was not yet converged at these very large model sizes.

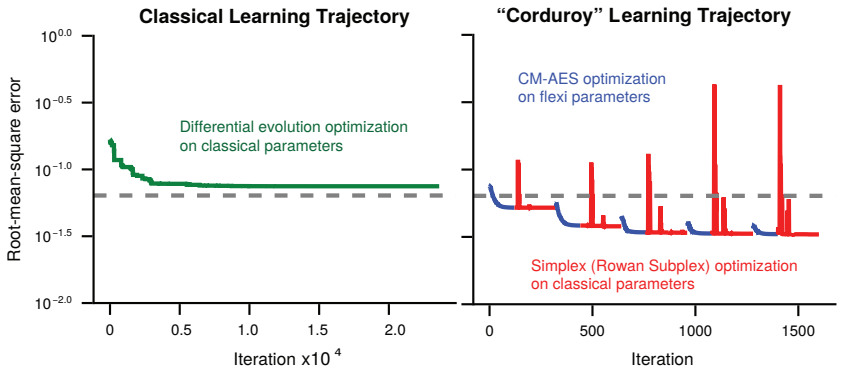

Figure S4: Iterative learning algorithm. To find the optimal parameter set, we first perform differential evolution optimization on just the classical model (left). Then, use these as initial guesses and repeat CMA-ES optimization on only the flexifunction parameters, keeping the classical parameters fixed, and simplex (Rowan's Subplex) on just the classical parameters, keeping the flexifunction parameters fixed (right)