An enhanced transcription factor repressilator that buffers stochasticity and entrains to an erratic external circadian signal

Steven A. Frank*

Department of Ecology and Evolutionary Biology, University of California, Irvine, Irvine, CA, United States

How do cellular regulatory networks solve the challenges of life? This article presents computer software to study that question, focusing on how transcription factor networks transform internal and external inputs into cellular response outputs. The example challenge concerns maintaining a circadian rhythm of molecular concentrations. The system must buffer intrinsic stochastic fluctuations in molecular concentrations and entrain to an external circadian signal that appears and disappears randomly. The software optimizes a stochastic differential equation of transcription factor protein dynamics and the associated mRNAs that produce those transcription factors. The cellular network takes as inputs the concentrations of the transcription factors and produces as outputs the transcription rates of the mRNAs that make the transcription factors. An artificial neural network encodes the cellular input-output function, allowing efficient search for solutions to the complex stochastic challenge. Several good solutions are discovered, measured by the probability distribution for the tracking deviation between the stochastic cellular circadian trajectory and the deterministic external circadian pattern. The solutions differ significantly from each other, showing that overparameterized cellular networks may solve a given challenge in a variety of ways. The computation method provides a major advance in its ability to find transcription factor network dynamics that can solve environmental challenges. The article concludes by drawing an analogy between overparameterized cellular networks and the dense and deeply connected overparameterized artificial neural networks that have succeeded so well in deep learning. Understanding how overparameterized networks solve challenges may provide insight into the evolutionary design of cellular regulation.

1 Introduction

Transcription factor (TF) networks control cellular response. The concentrations of the TF proteins feed into the TF network, which then stimulates or represses the expression of various genes. The TF network computes an input-output function: TF concentrations in, gene expression out.

A key puzzle of cellular design concerns how TF input-output functions solve the challenges of life. The puzzle poses two questions. First, what functional relation between

inputs and outputs solves the biological challenge? Second, how is the functional relation encoded molecularly?

This article focuses on the first problem, the functional solution to the biological challenge. Following a prior article, the challenge concerns maintaining a circadian rhythm of molecular concentrations (Frank, 2022c). The rhythm must buffer stochastic fluctuations in molecular concentrations.

An internal rhythm perturbed by stochasticity cannot retain perfect periodicity without an external signal. In this case, an external circadian signal appears and disappears randomly. That external signal provides an occasional opportunity for entrainment. However, the erratic external signal does not provide sufficient information for the cell to achieve good periodicity by simply mirroring the signal. Instead, the TF network must entrain to the external signal when available and otherwise buffer internal molecular stochasticity to maintain its internal rhythm in the absence of external information.

I use an artificial neural network to search for a TF input-output function that solves this circadian challenge. I embed the TF network in a stochastic differential equation that tracks the concentrations of mRNAs transcribed by gene expression and the associated TF proteins translated from the mRNAs. The TF network takes the TF protein concentrations as inputs and produces as outputs the mRNA transcription rates for the various TF genes.

I studied the same circadian challenge in the prior article (Frank, 2022c). That prior article tried to solve the functional and mechanistic parts of the puzzle simultaneously. Mechanistically, I encoded the TF input-output function using the currently favored thermodynamic model for TF binding to DNA gene promoter regions and the consequences of that binding for mRNA transcription rates (Bintu et al., 2005a; Bintu et al., 2005b).

Among many computer runs in that prior study, I found only one solution that provided reasonably good circadian tracking and entrainment based on explicit thermodynamic parameters for the TF network. That solution was very difficult to find computationally and was sensitive to changes in the thermodynamic parameters. Additionally, although the thermodynamic model I used is the most widely favored description for the molecular mechanism, it is very unlikely for that model to be an accurate and complete description of the actual molecular mechanism.

This article gains by focusing solely on the functional challenge of mapping inputs to outputs without concern for the mechanism. In particular, I use an artificial neural network to encode the functional mapping instead of the TF thermodynamic model. Once we have an understanding of the kinds of functional relations that solve the challenge, we can then search for candidate molecular mechanisms that could potentially encode those functional relations.

For example, given a functional solution encoded by a neural network, one could then fit the classic thermodynamic model to that functional solution. That fitting of the molecular mechanism to the functional solution would be useful when attempting to engineer the mechanism in actual cells and measure their molecular dynamics.

Prior modeling approaches searched for TF networks to achieve particular dynamics (Bolouri and Davidson, 2002; Kauffman et al., 2003; De Jong et al., 2004; Mishra et al., 2018; Patel and Bush, 2021). The approach in this article has several technical advantages. Neural networks are often ideal function approximators (Goodfellow et al.,

2016). Computationally, they scale very well to higher dimensions, can easily use automatic differentiation for optimization, and can be embedded within stochastic differential equations while maintaining the great computational advantages of automatic differentiation (Baydin et al., 2018; Rackauckas et al., 2020; Frank, 2022a).

Without these technical advantages, computational optimization of stochastic TF networks is difficult and has not previously been studied in a widely applicable way. The computer code provided with this article can easily be adapted to study other biological challenges. Additional studies will eventually give a sense of the kinds of input-output mappings required of TF networks to solve the demands of life.

As future studies accumulate, we may find that the evolutionary success of TF networks and the computational success of neural networks arise from a common foundation. Both networks may induce essentially the same geometric manifold of evolutionary or learning dynamics on which improving performance plays out (Frank, 2017).

2 Methods

This article extends the approach in Frank (2022c). The most important change replaces the prior thermodynamic model for the TF input-output response function with a computational neural network, which leads to greatly improved optimization outcomes.

I wrote the computer code in the Julia programming language (Bezanson et al., 2017). I used the SciML packages to optimize differential equations (Rackauckas et al., 2020). Efficient optimization depends on automatic differentiation (Baydin et al., 2018; Margossian, 2019), which is built into the SciML system. The source code for this article provides the details for all calculations and plotting of results (Frank, 2022b).

The goal is to optimize a TF regulatory control system in order to track an environmental target. I first describe the TF system and then describe the environmental target.

The deterministic component of the TF system dynamics is given by the temporal derivatives for numbers of mRNA molecules, x, and the TFs produced by those mRNAs, y, as

$$\begin{aligned}\dot{x}_i &= m_i f_i(\mathbf{y}) - \delta_i x_i \\ \dot{y}_i &= s_i x_i - \gamma_i y_i,\end{aligned} \tag{1}$$

for m_i as the maximum mRNA production rate, δ_i as the mRNA decay rate, s_i as the TF production rate per mRNA, and γ_i as the decay rate of the $i = 1, \ldots, n$ TFs (Marbach et al., 2010).

The function f_i transforms the numbers of TFs in the vector $\mathbf{y}$ into the production level of the ith mRNA, varying between 0 for complete repression and 1 for maximum production. In this article, each f_i is a separate neural network that takes n inputs as $\log(\mathbf{y})$.

The neural network architecture follows a standard general form, with the following details. The first layer of the network has $5n$ output nodes. Each of those nodes sums an affine transformation of each input, r, of the form $\alpha + \beta r$, in which each of those $5n^2$ transformations has its own α and β parameters. The value of each output node for this first layer is transformed by the mish activation function (Misra, 2019)
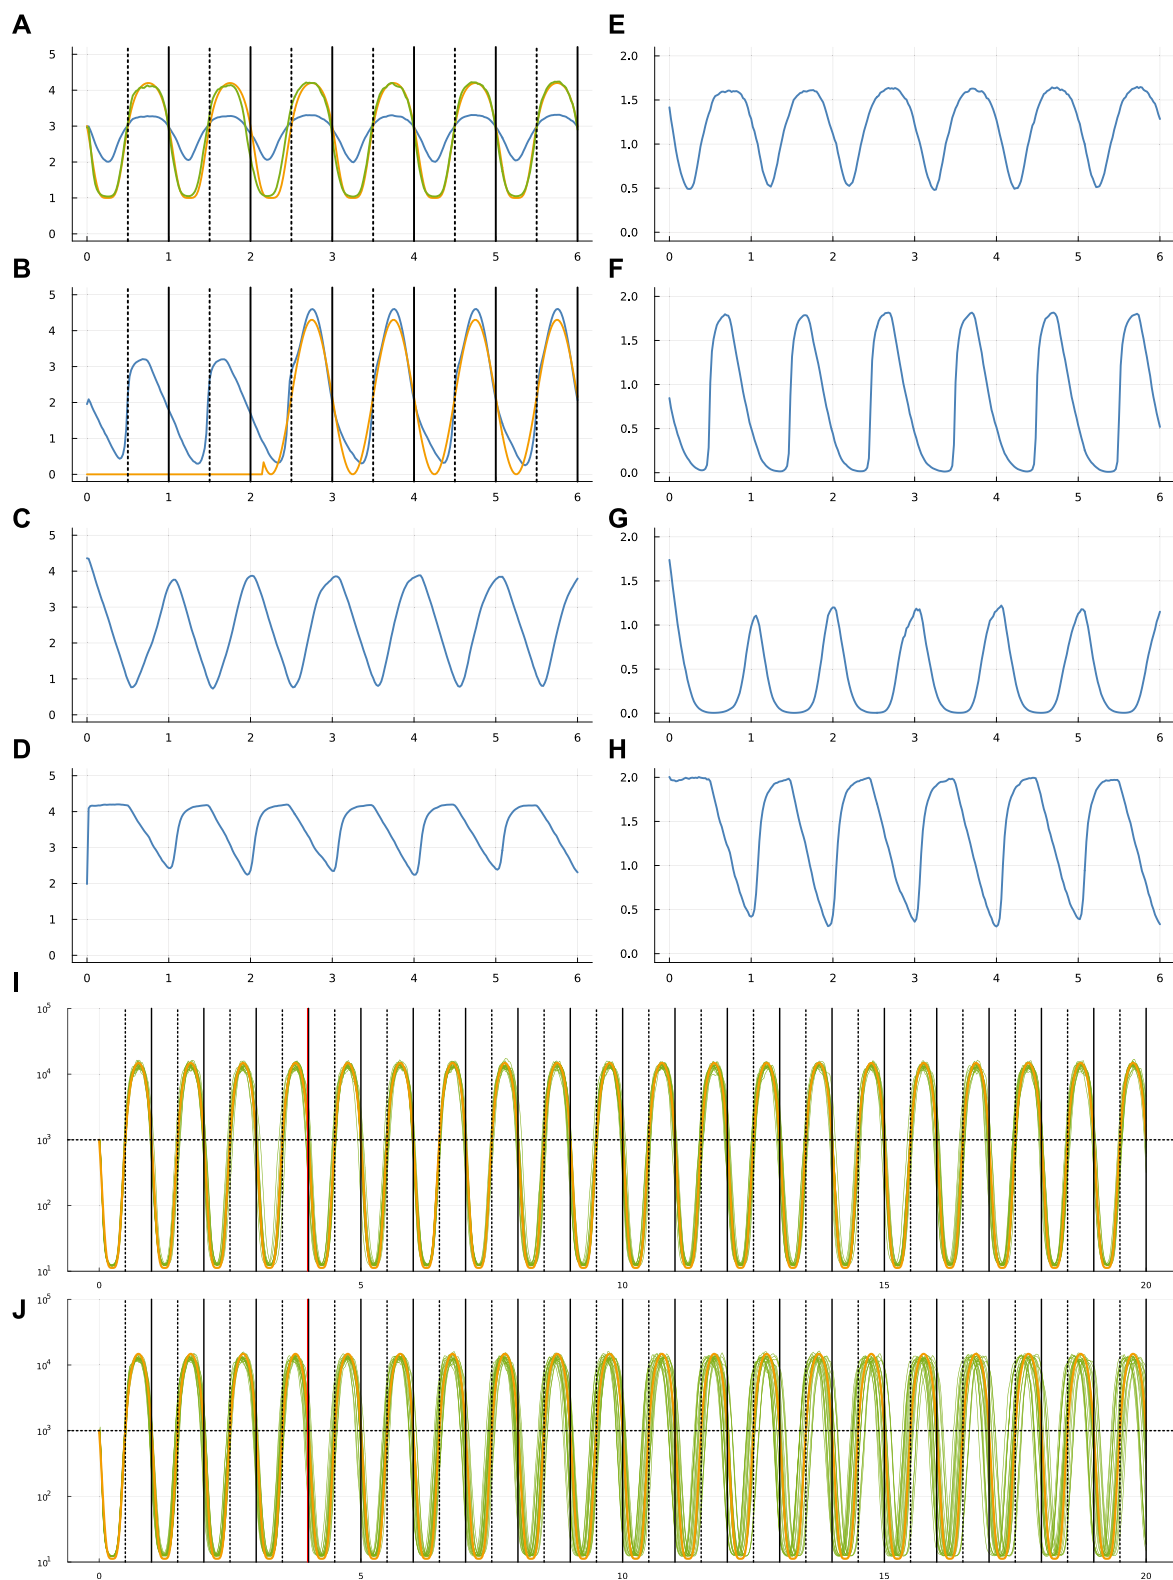
FIGURE 1
Circadian dynamics of the TF network from run sde−4_1_t4. Panels **(A–D)** show the stochastic dynamics of the TF proteins. The dynamics of the matching mRNAs that produce the proteins are shown in panels **(E–H)**. The time period is 6 days, along the x-axis of each of these panels. The vertical lines in **(A,B)** show entry into daylight (dotted) and nighttime (solid). The y-axis is $\log_{10}(1 + y)$ for number of molecules per cell, y. In **(A)**, the optimization procedure attempts to match the number of TF 1 molecules (blue curve) to a circadian rhythm by minimizing a loss value. To calculate the loss value, begin with the number of TF 1 molecules, y, transformed by a Hill function, $\bar{y} = y^2/(10^6 + y^2)$, to yield the green curve, which traces $1 + 4\bar{y}$. The gold curve traces the target circadian pattern. The loss value to be minimized is the sum of the squared deviations between the gold and green curves at 50 equally spaced time points per day. The number of TF 2 proteins in **(B)** is influenced by the internal cellular dynamics and is also increased in response to an

(Continued)

$$\mathrm{mish}\,(r) = r \tanh\left(\log\left(1 + e^r\right)\right).$$

The $5n$ outputs from the first layer form the inputs for another neural network layer, which leads to the final single-valued output that controls the mRNA transcription rate for the associated TF gene. That output arises by first summing affine transformations from all $5n$ inputs and then applying the sigmoid function to that sum, in which the sigmoid function outputs values between 0 and 1, as

$$\mathrm{sigmoid}\,(r) = \frac{1}{1 + e^{-r}}.$$

The same architecture is repeated for each of the n TFs, with the single-valued output of each network influencing the transcription rate of the associated TF gene.

I added stochastic perturbations to the molecular abundances in Eq. 1, making the system a set of stochastic differential equations. Stochastic fluctuations for a molecule with abundance z follow $\sqrt{z}\mathrm{d}W$ in which W is a standard normal variate with mean 0 and standard deviation 1, thus $\sqrt{z}W$ has a standard deviation of $\sqrt{z}$.

The normalized magnitude of the fluctuations is the ratio of the standard deviation relative to the current abundance, $\sqrt{z}/z = 1/\sqrt{z}$. As z drops, the normalized fluctuations increase. To prevent fluctuations from becoming too large relative to the abundance, which can cause negative abundance values in the numerical analysis, I replaced $\sqrt{z}$ with $z/4$ for $z \leq 16$. Thus, we may write the stochastic component as $g(z) = \sqrt{z}$ for $z > 16$ and $z/4$ for $z \leq 16$, leading to the stochastic differential equations

$$\begin{aligned}\mathrm{d}x_i &= \left[m_i f_i(\mathbf{y}) - \delta_i x_i\right]\mathrm{d}t + g(x_i)\mathrm{d}W \\ \mathrm{d}y_i &= \left[s_i x_i - \gamma_i y_i\right]\mathrm{d}t + g(y_i)\mathrm{d}W. \end{aligned} \tag{2}$$

I calculated the trajectories of these stochastic differential equations with the Julia DifferentialEquations.jl package (Rackauckas and Nie, 2017), using the Ito implicit method ISSEM. The Julia SciML libraries optimize the performance of the TF system by passing the automatic differentiation procedure through the stochastic differential equations, including the embedded neural networks.

The design goal is for the stochastic TF system to maintain a circadian rhythm given only a sporadically present external circadian signal. I first describe the design goal. I then follow with details about the external circadian signal.

The design goal is for TF 1 (y_1) to follow a 24 h period. TF abundance above $S = 10^3$ molecules per cell corresponds to an "on" state for daytime. Below that threshold, the cell is in an "off" nighttime state.

The optimization procedure seeks a parameter combination that minimizes the loss measured as the distance between the target circadian pattern shown in the gold curve and the transformed

abundance of TF 1 in the green curve in Figure 1A. In particular, the daily target rhythm follows

$$v_t = \frac{\sin\left(2\pi t + \pi\right) + 1}{2}, \tag{3}$$

with t in days and passage across 0.5 corresponding to transitions between day and night. The loss is

$$L = \sum_t \left(\frac{y_t^2}{S^2 + y_t^2} - \frac{v_t^2}{0.5 + v_t^2}\right)^2,$$

in which y_t is the abundance of TF 1, and the transformations of y and v are Hill functions that normalize values to make the different scales comparable, and $S = 10^3$ is the cellular state switch rate given above. For the summation, the time values start at $t = 0$ and increment by 0.02 until the end of time for a particular analysis. The values for y and v in Figure 1A are normalized to the scaling of the plots, as described in the figure legend.

A stochastic system inevitably diverges from the target circadian trajectory. In this model, the system may use an external daylight entrainment signal to correct deviations. To pass information about the external signal into the system in Eq. 2, I added a strong boost to the production rate of TF 2 (y_2) in proportion to the intensity of external light. In particular, the rate of change in TF 2 abundance in the presence of the external light signal is augmented by $u_t = 10^{6v_t - 4}$, in which v_t is the daily rhythm in Eq. 3.

Thus, when the external light signal is present, I added $u_t\mathrm{d}t$ to the value of $\mathrm{d}y_2$ in Eq. 2, noting that all y and x values are functions of t. The augmented system becomes

$$\begin{aligned}\mathrm{d}x_i &= \left[m_i f_i(\mathbf{y}) - \delta_i x_i\right]\mathrm{d}t + g(x_i)\mathrm{d}W \\ \mathrm{d}y_i &= \left[s_i x_i - \gamma_i y_i\right]\mathrm{d}t + g(y_i)\mathrm{d}W \quad i \neq 2 \\ \mathrm{d}y_2 &= \left[s_2 x_2 - \gamma_2 y_2 + u_t\right]\mathrm{d}t + g(y_2)\mathrm{d}W. \end{aligned} \tag{4}$$

Initially, the external light signal is absent. The signal switches on and off randomly. In Figure 1B, the gold curve shows the external light signal, which switches on in the third day and stays on for the remaining days shown. The blue curve traces the abundance of TF 2. Thus, the overall challenge is for TF 1 to track the circadian pattern, given internal stochastic molecular dynamics and a randomly occurring external entrainment signal that influences the abundance of TF 2.

In future studies, one may wish to consider how organisms use additional environmental signals to synchronize internal rhythmicity. For example, there may be additional TFs that capture those signals, as TF 2 does for light in the model given here.

This model sets the influence of the external light signal on the internal system by the level of u_t. In future studies, one may consider how changing the strength of this influence affects the dynamics. Control theory provides some clues (Frank, 2018).

Think of the magnitude of u_t as the gain of a control signal. Simplifying slightly, increasing the gain of a signal typically has two potential consequences. First, higher gain means that the system moves toward its target more quickly. Second, higher gain reduces system stability, with the potential for disastrous instability. Much of control theory is devoted to analyzing this design tradeoff between performance and stability.

3 Results

3.1 Overview

Computational optimization finds parameters for Eq. 1 that influence the stochastic dynamics of mRNAs and TF proteins. Different runs of the optimization procedure converged to different parameters. Figure 1 illustrates the results for one run.

In that figure, panel (a) shows the target circadian pattern in gold and the system's circadian pattern in green. The optimization goal is to minimize the total squared distance between those two curves (Frank, 2022c). The blue curve in that panel shows the molecular abundance of the TF 1 protein. The green curve arises by transforming the blue curve by a Hill function, as explained in the figure caption. The x-axis shows time measured in days.

The optimization begins by attempting to fit the parameters to the circadian pattern for part of the first day. Then, the algorithm slowly increases the temporal range in a sequence of steps, for each step allowing the fit to adjust to the additional time range. The final time range for fitting the parameters in these runs was 4 days. Once the fit is obtained, I studied the match between the target circadian pattern and the system dynamics over a time period greater than the 4 days used for fitting.

I wrote the code to emphasize the chance of finding a good fit rather than to maximize the speed of the fitting process. A single optimization run takes about 10 days when using 12 computational threads on a 2022 Apple Studio Ultra M1 computer. I analyzed eight full runs for this study, which was sufficient to find several good fits and to illustrate a range of outcomes.

3.2 Dynamics

In Figure 1B shows the dynamics of the TF 2 protein. The blue curve is the abundance of that TF. The gold curve shows the external circadian entrainment signal, which may be interpreted as the external light intensity experienced by the cell. The entrainment signal comes and goes randomly. For these runs, the signal begins in the off state and then switches on and off randomly with a waiting time drawn from an exponential distribution with a mean of 2 days. The TF 2 protein increases in abundance in response to the light intensity (Frank, 2022c).

Overall, panels (a-d) show the dynamics of the four TF proteins, which are the y_i values in Eq. 4 for a system with $n = 4$ TFs. For each TF protein, the matching panel to the right shows the dynamics of the mRNAs that encode the TFs.

The gold curve in panel (i) shows the target circadian pattern over 20 days. Each of the 20 green curves shows the cellular circadian state for a run of the stochastic dynamics. As in panel (a), cellular state is calculated as a Hill transformation of the abundance of TF 1. For each run in panel (i), the external entrainment signal comes on and off randomly with a mean waiting time of 2 days, as in panel (b). The parameters are optimized for runs with a time period of 4 days, indicated by the vertical red line at day 4. The dynamics over the subsequent 16 days shows how the system performs beyond the training optimization period.

Panel (j) repeats the setup in panel (i), except that the external signal comes on and off randomly with a mean waiting time of 1,000 days. Because the signal starts in the off state, almost always the signal never comes on. The green curves therefore show the ability of the system to maintain its circadian rhythm in the absence of any external signal. Inevitably, the stochastic fluctuations of molecular concentrations cause the cellular rhythm to diverge from the target pattern.

Figures 2, 3 show the same data for two other optimization runs. The following subsections compare the different optimization runs. For all runs, the settings, optimized parameters, output values, and code for making the figures are included in the online repositories listed in the Data Availability statement at the end of the article.

The Supplementary Material provides three Supplementary Figures S1–S3, that show dynamics for the same systems as in Figures 1–3 but with different patterns for the stochastic external light signal. Comparing two different runs of the same system helps to visualize how different stochastic inputs alter system dynamics.

3.3 Deviation from the target pattern

Figure 4A shows the deviations of each optimized system from the target circadian pattern. For a single realization, such as a green curve in panel (i) or (j) of the previous figures, I calculated the deviations as follows. The entry into external daytime occurs at the vertical dotted line in each daily interval, and the gold curve crosses the horizontal dotted line at 10^3. For a cell, the entry into its internal daytime state occurs as the green curve crosses above the horizontal dotted line at abundance 10^3.

For each day, the cellular deviation is the horizontal distance along the dotted line between the green curve and the gold curve. When the green curve is to the left of the gold curve, the deviation is negative. When it is to the right, the deviation is positive. I transformed the deviation into hours, the daily deviation amount divided by 24. For a realization of the dynamics over d days, there are d deviation measures for that realization.

For realizations that ran over d days, I repeated the measurements over 1,000 samples for a total of $1000d$ deviation measures. I then calculated percentiles of the deviations. Looking at the far left of Figure 4A, the two vertical lines and the circle show those percentiles. The lower line ranges from the 5th to the 25th percentile. The circle marks the 50th percentile. The upper line ranges from the 75th to the 95th percentile.

Thus, each vertical set summarizes the distribution of deviations for a sample of 1,000 realizations over d days. In the figure, the vertical summary for each distribution occurs within a set of three such distributions. For each set, the left, middle, and right distributions correspond to $d = 10, 20, 30$, respectively, showing the distributions of deviations when realizations run over different numbers of days.
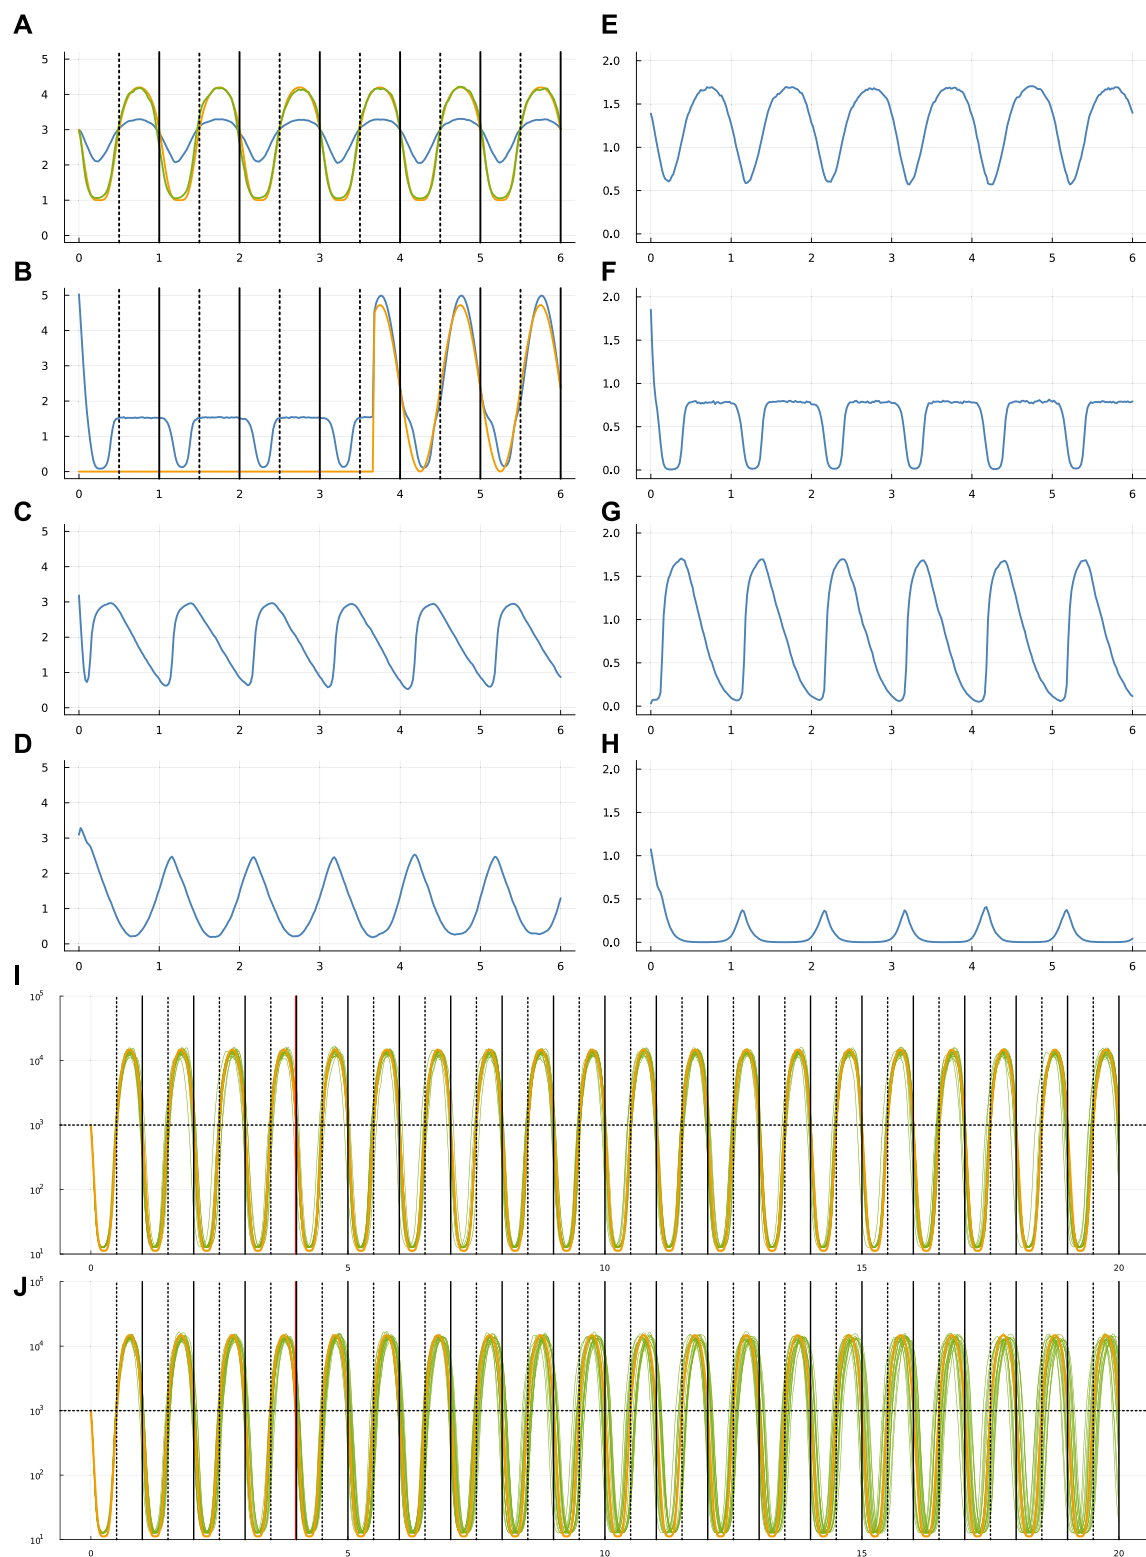
FIGURE 2
Circadian dynamics for run sde−4_2_t4.

The eight labels below the sets in Figure 4A provide the names for the optimization runs. I chose for further analysis the two sets on the left and the set on the right. Those three optimization runs had the narrowest distributions for the deviations, in which smaller deviations correspond to better tracking of the target circadian pattern.

The leftmost run, sde−4_1_t4, corresponds to the dynamics in Figure 1 and to the more detailed summary of deviations in
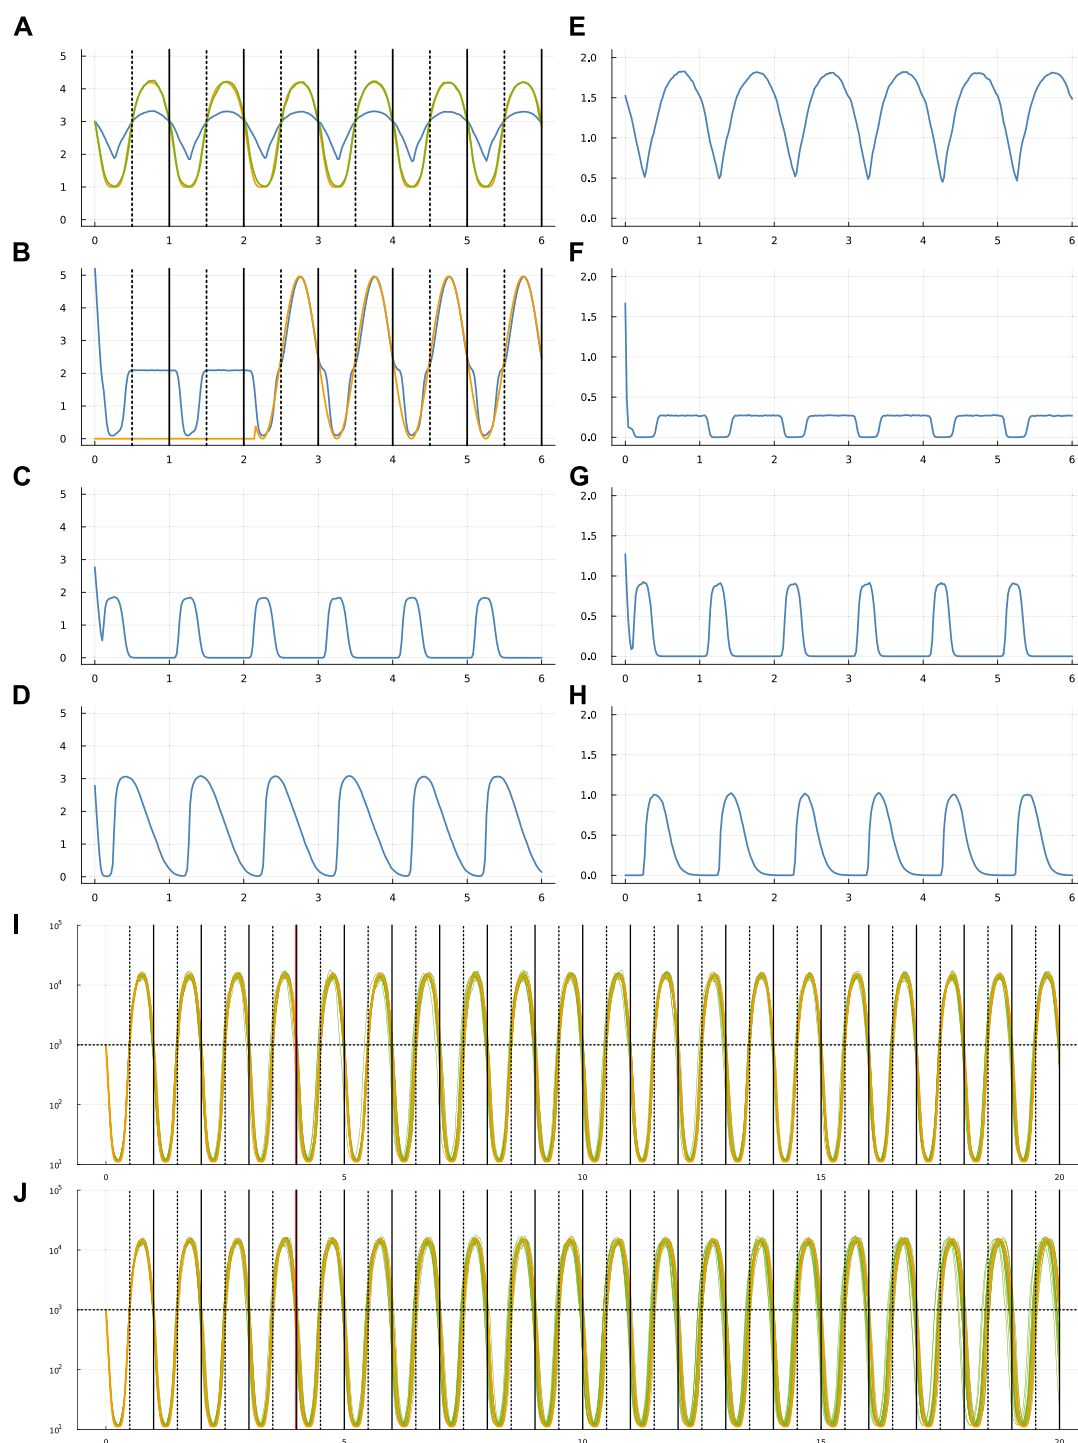
FIGURE 3
Circadian dynamics for run sde−4_8_t4.

Figure 4B. The deviation details in Figure 4B show the distributions induced by different random switching times for the external entrainment signal. Each set of three distributions again corresponds to runs with d = 10, 20, 30 days. The different sets correspond to mean waiting times for random exponential switching of 2, 4, 8, 16, and 1,000 days.

The mean time of 2 days corresponds to the value used in during optimization and is the default in all other plots unless otherwise noted. The longer mean times illustrate how well the system can hold its internal circadian rhythm when receiving an external entrainment signal less reliably. When the mean is 1,000 days, the system almost never receives an external signal. The bigger deviations in that case show the greater
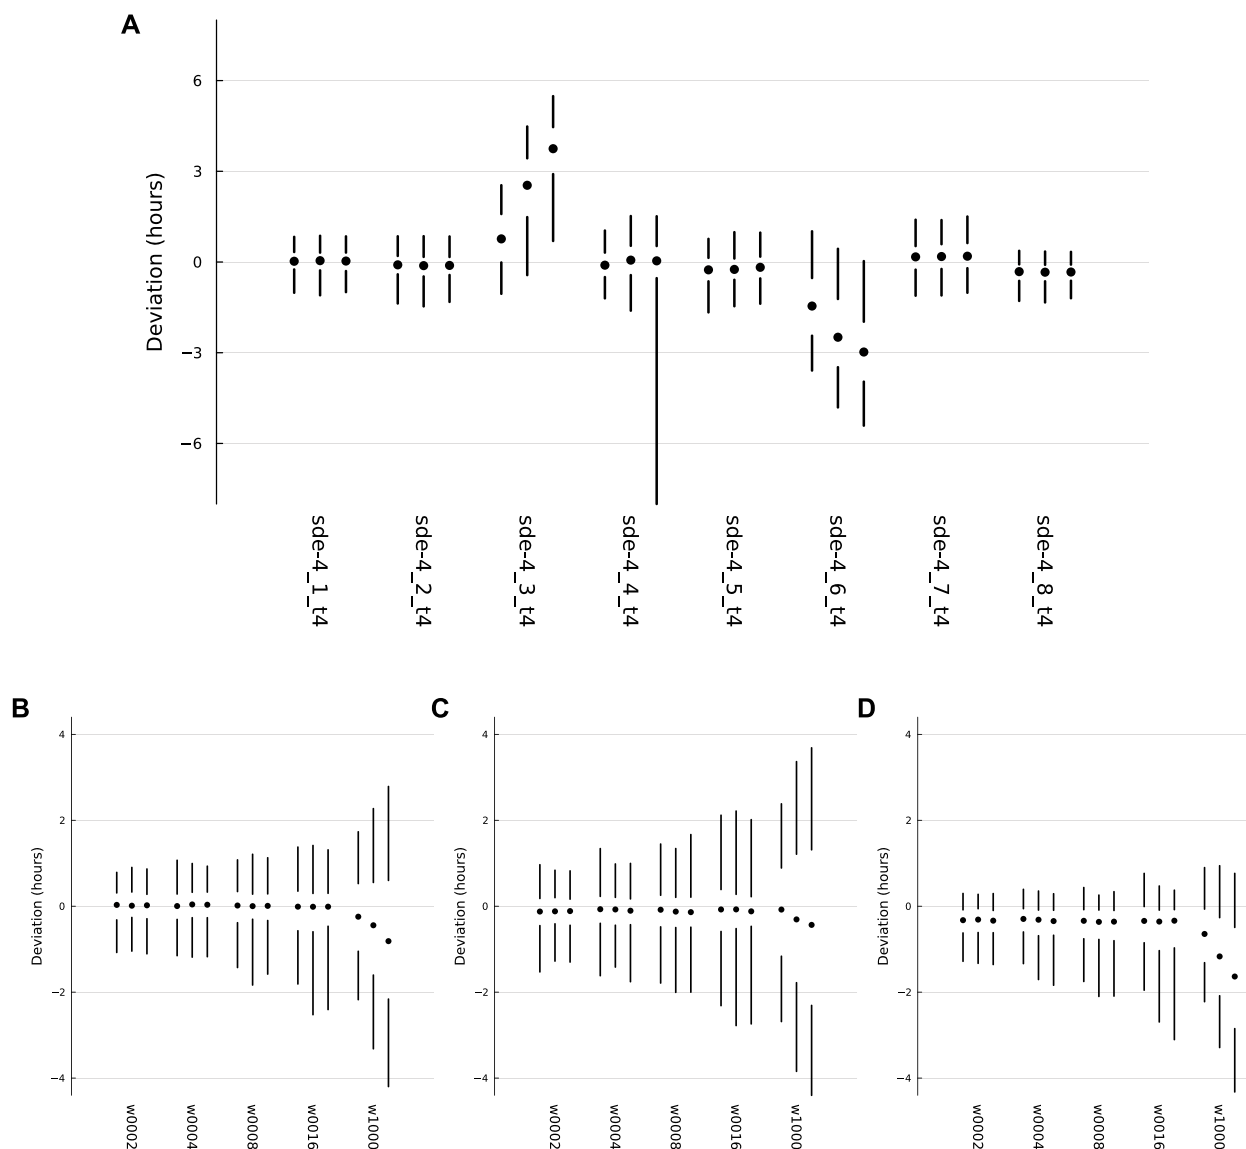
FIGURE 4
Probability distributions for deviations between the actual circadian pattern of the TF 1 protein and the target circadian pattern. Each set of two vertical lines and a circle describe a distribution, with the 5th to 25th percentiles of the distribution along the lower line, the 50th percentile at the circle, and the 75th to 95th percentiles along the upper line. **(A)** The overall patterns for the eight independent optimization runs, with the label for each run along the x-axis. **(B)** The distributions for sde−4_1_t4, for different values of w, the average waiting time for switching of the external circadian light signal. **(C)** The distributions for sde−4_2_t4. **(D)** The distributions for sde−4_8_t4. Further details in the text.

difficulty of maintaining a match to the target pattern without the opportunity for entrainment. Although the spreads are greater for larger d, the system nonetheless typically keeps a very good circadian pattern for 10 days and a reasonably good pattern for as along as 30 days.

Figure 4C corresponds to sde−4_2_t4, and Figure 4D corresponds to sde−4_8_t4. The sets do not match exactly between panel (a), which has a mean waiting time of 2 days, and the leftmost sets in panels (b-d), which also have mean waiting times of 2 days. The small mismatches arise because the different panels used different samples of the stochastic trajectories to calculate the distributions.

3.4 TF network input-output functions

Figure 5 shows the TF network input-output relations associated with sde−4_1_t4. Each quadrant presents the output as the relative mRNA production rate for a TF gene. For example, the upper-left block describes the output for TF 1. In each plot, the output level varies from off, in dark orange, to maximally on, in dark purple.

The output level depends on four inputs, the abundances of the four TF proteins. The x- and y-axes of a plot correspond to the abundances of TF 1 and TF 2, respectively, on a $\log_{10}(1 + z)$ scale for abundance z. The abundance of TF 3 rises across the rows of plots, from zero at the bottom of a quadrant to the maximum level at the
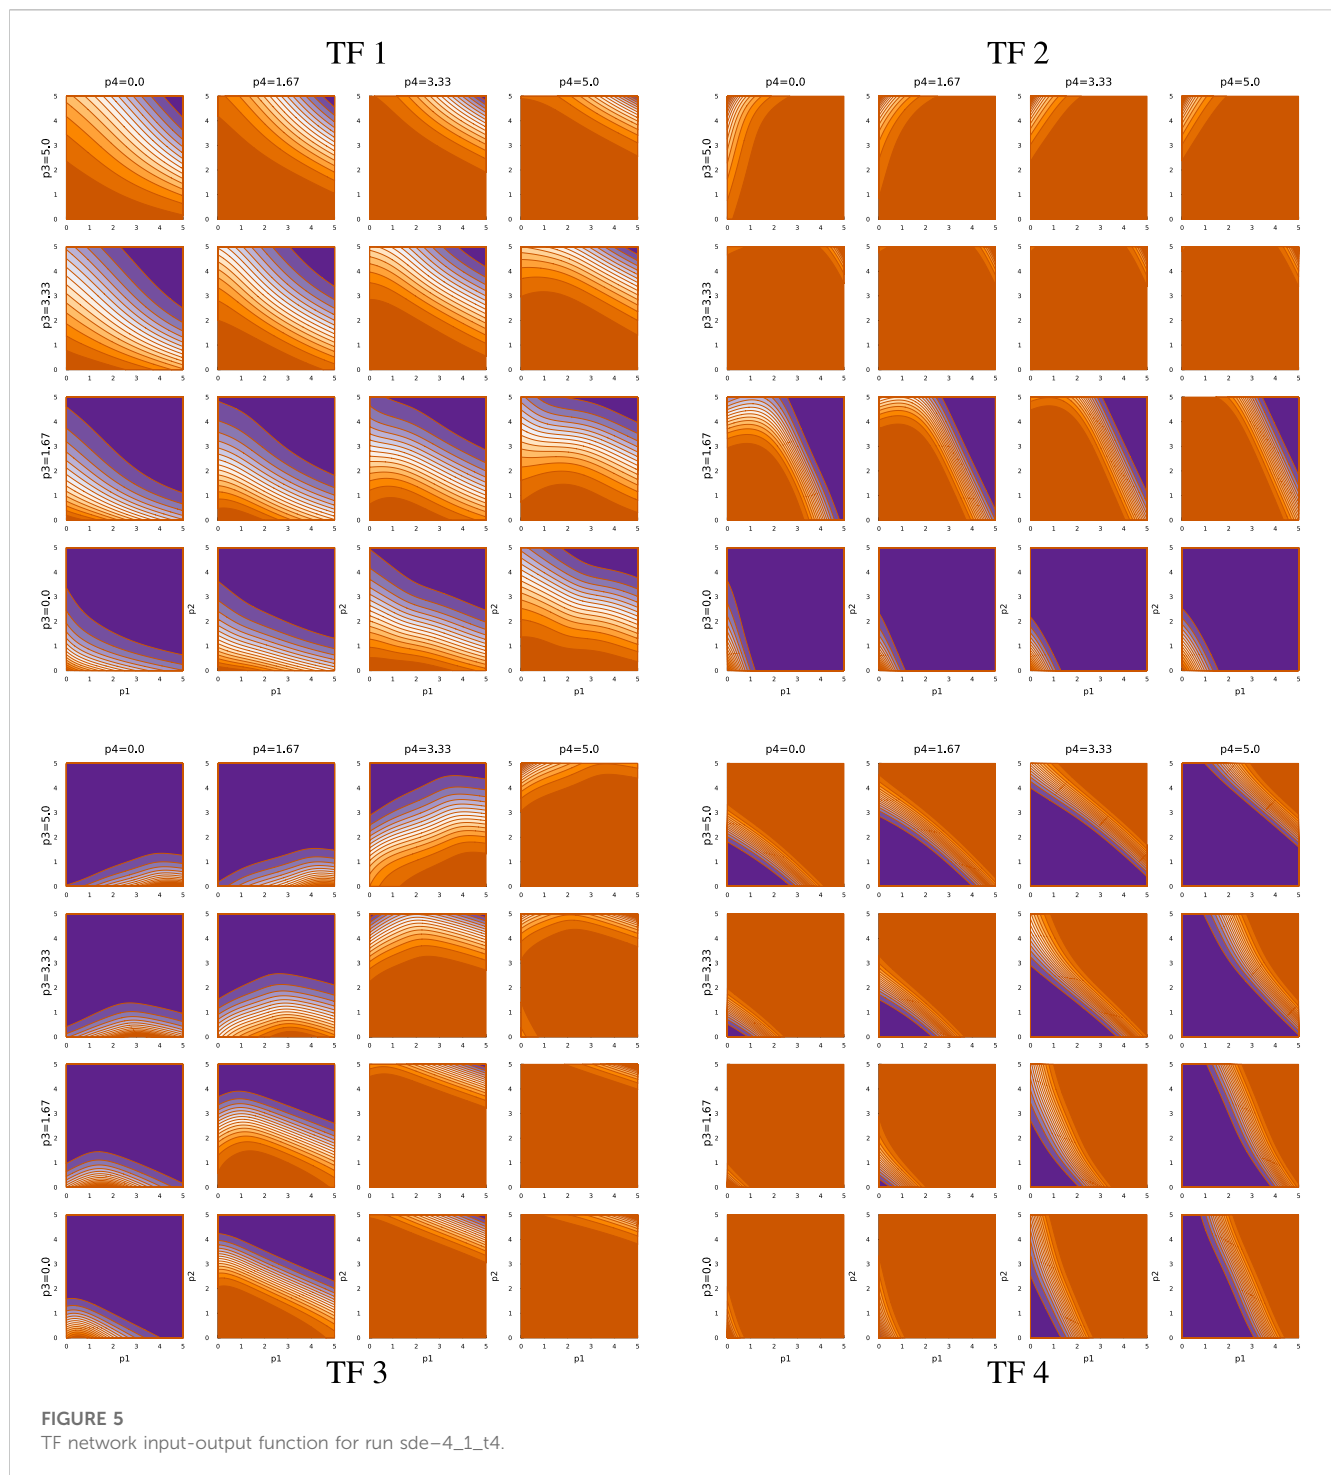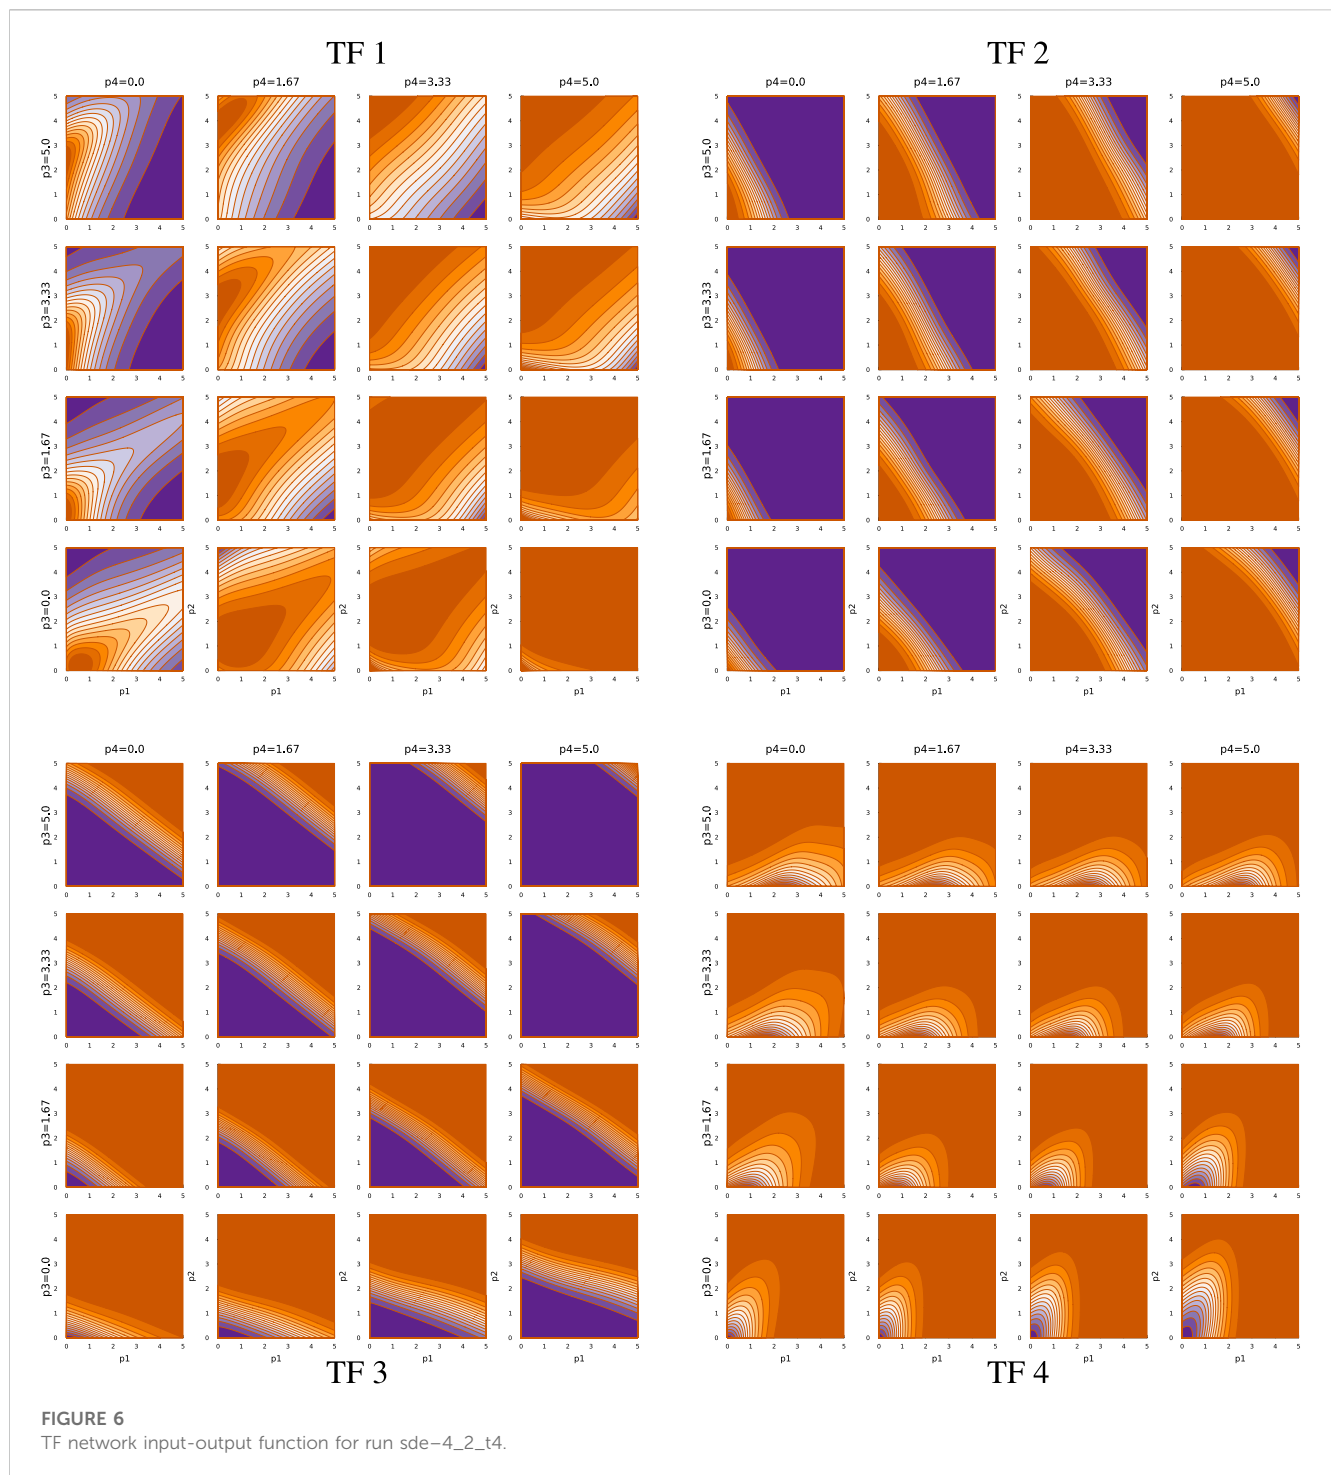
FIGURE 5
TF network input-output function for run sde−4_1_t4.

top. The abundance of TF 4 rises across the columns, from zero at the left to maximal at the right.

The TF network logic can be read from the figure. Consider, for example, the TF 2 block in the upper right that shows the TF 2 mRNA output rate. An increase in TF 3 abundance, going up the rows, associates with repression of TF 2, because increasing orange means lower output. In the TF 3 block (lower left), an increase in TF 4, going right across columns, associates with repression of TF 3. In the TF 4 block (lower right), an increase in

TF 2, going right along the *x*-axis of each plot, associates with repression of TF 4.

Thus, TF 3 represses TF 2, and TF 4 represses TF 3, and TF 2 represses TF 4. That cyclic repression defines the basic repressilator circuit, which can cause periodic oscillation of the associated TF proteins (Elowitz and Leibler, 2000). However, in this case, that core repressilator is embedded within a more complex network of modulating inputs. For example, TF 2 output is stimulated by TF 1 when both TF 2 and TF 3 abundances are low.

FIGURE 6
TF network input-output function for run sde−4_2_t4.

In this optimization problem, one major function of the TF network is the entrainment to an external circadian signal when available, illustrated by panels (a) and (b) of Figure 1. Returning to the upper-left quadrant of Figure 5, note that increasing TF 2 abundance (*y*-axis of plots) stimulates TF 1 expression. Thus, when an external light signal increases TF 2, that rise in TF 2 stimulates a rise in TF 1 expression. A decline TF 2 abundance during the night represses TF 1 expression.

Other inputs modulate the relation between TF 2 and TF 1. For example, high TF 3 and TF 4 prevent TF 1 expression, unless both

TF 1 and TF 2 are very high. TF 1 stimulates its own expression on the way up and represses its own expression on the way down.

The second major function of the TF network is to maintain the circadian pattern by buffering the intrinsic biochemical stochasticity. That buffering arises from the various interactions of the TF network logic. One aspect may be the tendency to maintain both TF 1 and TF 2 cycling in synchrony with the target rhythm. Joint rhythmicity provides a partially redundant cycle that can be used to smooth out fluctuations in either TF protein.

FIGURE 7
TF network input-output function for run sde−4_8_t4.
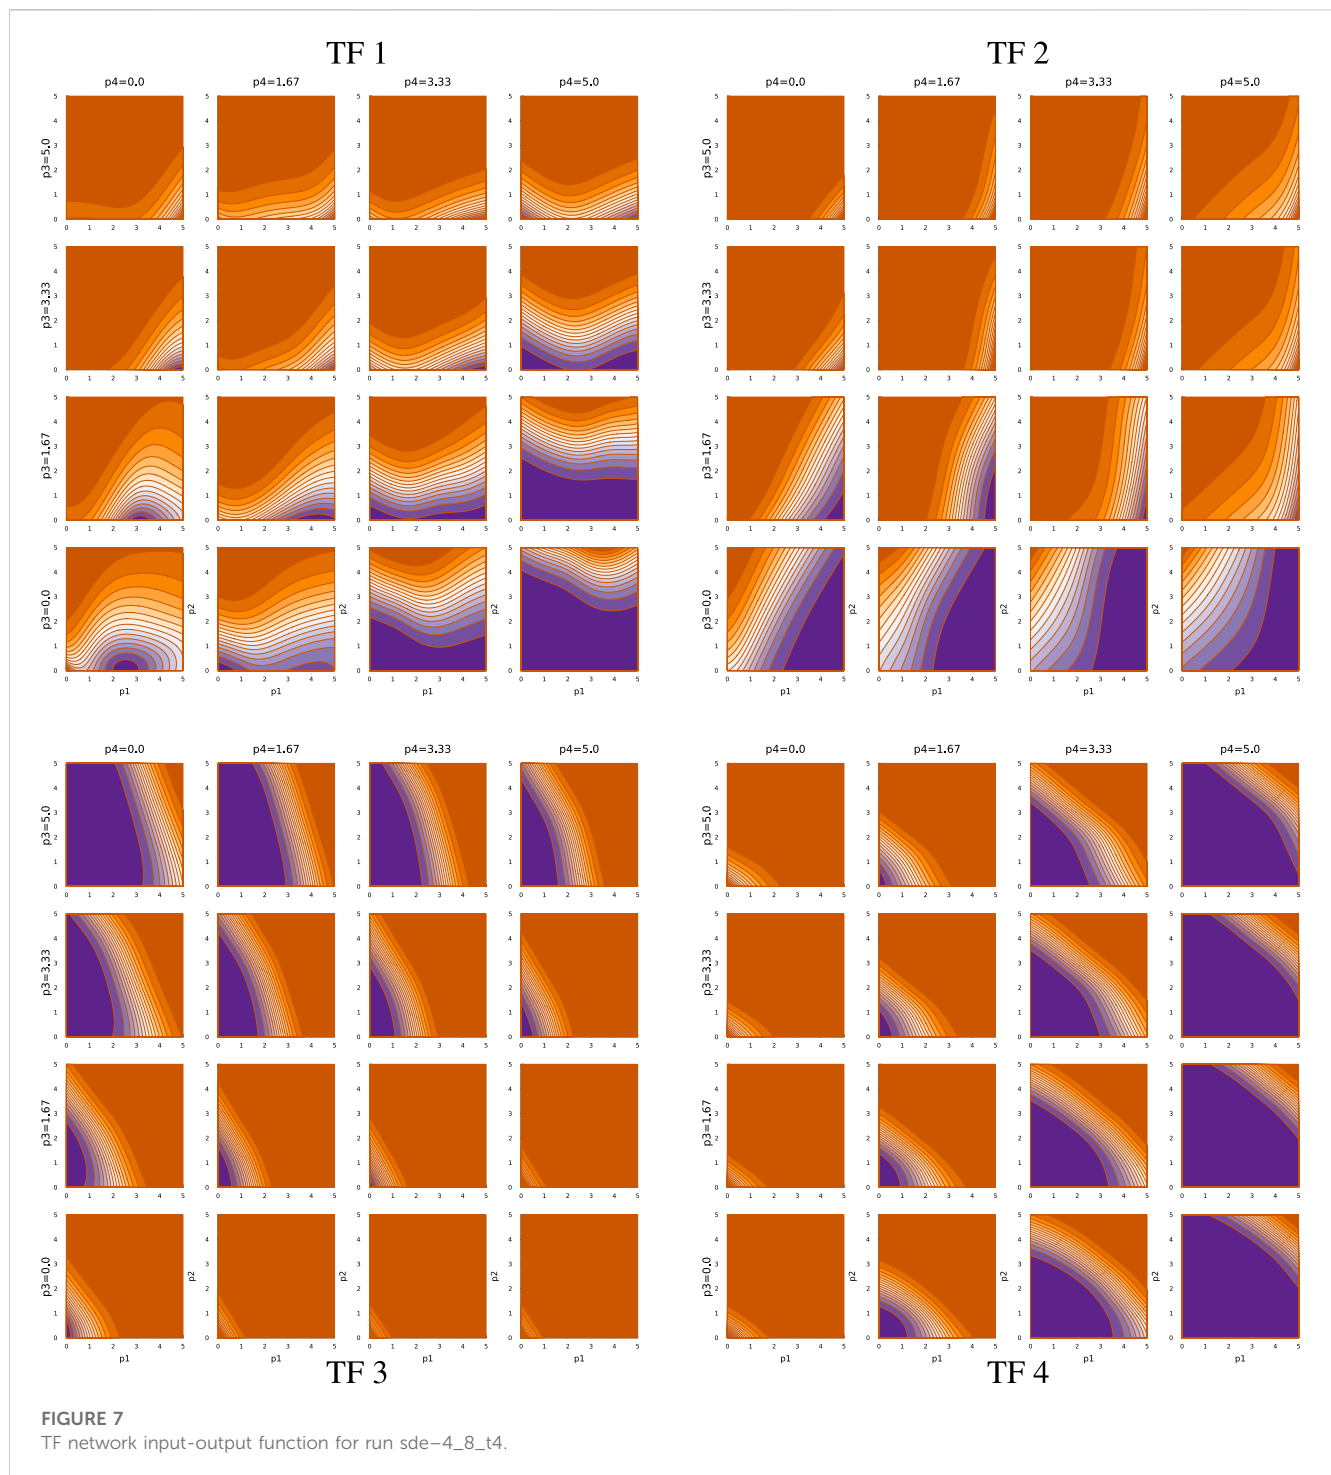
Figures 6, 7 show the TF network logic for sde−4_2_t4 and sde−4_8_t4, respectively. Looking at Figure 6, one can see that the quantitative relations differ significantly from Figure 5. For example, most effects of changing abundance in Figure 5 are monotonic. By contrast, in Figure 6, the upper-left and lower-right quadrants show reversals in the direction of change in outputs with changes in particular inputs.

Overall, Figures 5–7 demonstrate that TF networks can produce good performance in different ways. Finding the necessary and sufficient characteristics of the TF networks remains an open problem.

4 Discussion and conclusion

How do cellular regulatory networks solve the challenges of life? Ultimately, that is an empirical question that must be answered by observational and experimental study. However, it

may be a difficult question to answer without some conceptual guidance.

Suppose, for example, that the network connections regulating cellular response are dense, deep, and complex. It might be that the network is overparameterized. In other words, the dimensionality of the molecular parameters controlling cellular response may be greater than the dimensionality of the challenges posed by the environment.

Computational deep learning models are typically highly overparameterized artificial neural networks (Bartlett et al., 2020; Poggio et al., 2020; Radhakrishnan et al., 2020). It turns out that overparameterized networks are particularly good at solving complex problems. And, for such overparameterized networks, it is particularly difficult to figure out how the network actually succeeds at solving its challenge (Amey et al., 2021).

One might find some network nodes that work together to achieve a particular component of the broader solution. For example, some nodes might together identify an environmental state or trigger a particular response that feeds into other parts of the network. But, overall, the relations between the individual network components and the response of the network to particular inputs typically remains opaque.

In the same way, it may be difficult to build up an understanding of cellular function by starting with small molecular subnetworks and then expanding analysis to increasingly broader subcomponents. If so, then understanding in theory how such networks might solve problems may provide some guidance.

For example, what are the variety of ways in which cells might buffer against stochastic perturbations to circadian periodicity? How might cells entrain their rhythm to erratic external signals? In overparameterized networks, there is rarely just one solution. Theory can help to identify the range of possible solutions, which could then guide where to look within the complexity of actual cellular regulation.

This article provides an early step on the path to developing theoretical approaches. By using artificial neural networks for the transcription factor network, the computational models can explore the variety of potential solutions to life's challenges. The methods presented here show how relatively simple computer code can discover potential solutions.

5 Afterword: more on potential applications

The computational method here seeks TF networks that do well at responding to particular challenges. I chose the challenge of maintaining an internal circadian rhythm as a good illustration of the method. The first steps in this kind of modeling require sorting out how to do the computations to find good candidate solutions. No prior method did this in a simple and general way for the intrinsic stochastic fluctuations in molecular abundances that inevitably happen in real cells.

As such methods improve, what sorts of insights might arise? Initially, synthetic biology provides the most likely first applications (He et al., 2016; Gibson et al., 2017; Glass and Alon, 2018; Lim, 2022). Following the highly influential repressilator model that gave rise to the title of this article (Elowitz and Leibler, 2000), the computational approach here may help to find new candidate designs for synthetically engineered cells. With these methods,

one can easily change the design goals to solve problems other than circadian patterns. Any problem that can be expressed as a differentiable loss function should fit within this article's approach.

Understanding nature's designs presents different challenges. Optimization methods can often help (Maynard Smith, 1978; Parker and Maynard Smith, 1990). For example, in theory what kinds of environmental challenges lead to regulatory network designs that match what we see? How does changing the broad nature of the challenge lead to different designs? How much do details of the environmental challenge and biophysical constraints influence favored designs? When comparing organisms that solve similar challenges in different ways, what sorts of evolutionary forces might explain those observed differences?

Data availability statement

The original contributions presented in the study are available in the Zenodo repository, https://doi.org/10.5281/zenodo.7178508, further inquiries can be directed to the corresponding author.

Author contributions

SF: Conceptualization, Formal Analysis, Funding acquisition, Investigation, Software, Visualization, Writing–original draft, Writing–review and editing.

Funding

Conflict of interest

The author declares that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Publisher's note

Supplementary material

The Supplementary Material for this article can be found online at: https://www.frontiersin.org/articles/10.3389/fsysb.2023.1276734/full#supplementary-material

References

Amey, J. L., Keeley, J., Choudhury, T., and Kuprov, I. (2021). Neural network interpretation using descrambler groups. *Proc. Natl. Acad. Sci. U. S. A.* 118, e2016917118. doi:10.1073/pnas.2016917118

Bartlett, P. L., Long, P. M., Lugosi, G., and Tsigler, A. (2020). Benign overfitting in linear regression. *Proc. Natl. Acad. Sci. U. S. A.* 117, 30063–30070. doi:10.1073/pnas.1907378117

Baydin, A. G., Pearlmutter, B. A., Radul, A. A., and Siskind, J. M. (2018). Automatic differentiation in machine learning: a survey. *J. Mach. Learn. Res.* 18, 1–43.

Bezanson, J., Edelman, A., Karpinski, S., and Shah, V. B. (2017). Julia: a fresh approach to numerical computing. *SIAM Rev.* 59, 65–98. doi:10.1137/141000671

Bintu, L., Buchler, N. E., Garcia, H. G., Gerland, U., Hwa, T., Kondev, J., et al. (2005a). Transcriptional regulation by the numbers: models. *Curr. Opin. Genet. Dev.* 15, 116–124. doi:10.1016/j.gde.2005.02.007

Bintu, L., Buchler, N. E., Garcia, H. G., Gerland, U., Hwa, T., Kondev, J., et al. (2005b). Transcriptional regulation by the numbers: applications. *Curr. Opin. Genet. Dev.* 15, 125–135. doi:10.1016/j.gde.2005.02.006

Bolouri, H., and Davidson, E. H. (2002). Modeling transcriptional regulatory networks. *BioEssays* 24, 1118–1129. doi:10.1002/bies.10189

De Jong, H., Gouzé, J.-L., Hernandez, C., Page, M., Sari, T., and Geiselmann, J. (2004). Qualitative simulation of genetic regulatory networks using piecewise-linear models. *Bull. Math. Biol.* 66, 301–340. doi:10.1016/j.bulm.2003.08.010

D. G. Gibson, C. A. HutchisonIII, O. Hamilton, and J. C. Venter (Editors) (2017). *Synthetic biology: tools for engineering biological systems* (Woodbury, NY: Cold Springs Harbor Laboratory Press).

Elowitz, M. B., and Leibler, S. (2000). A synthetic oscillatory network of transcriptional regulators. *Nature* 403, 335–338. doi:10.1038/35002125

Frank, S. A. (2017). Puzzles in modern biology. V. Why are genomes overwired? *F1000Research* 6, 924. doi:10.12688/f1000research.11911.1

Frank, S. A. (2018). *Control theory tutorial: basic concepts illustrated by software examples*. Cham, Switzerland: Springer.

Frank, S. A. (2022a). Automatic differentiation and the optimization of differential equation models in biology. *Front. Ecol. Evol.* 10, 1010278. doi:10.3389/fevo.2022.1010278

Frank, S. A. (2022b). An enhanced transcription factor repressilator that buffers stochasticity and entrains to an erratic external circadian signal: Julia software code. *bioRxiv.* doi:10.5281/zenodo.7178508

Frank, S. A. (2022c). Optimization of transcription factor genetic circuits. *Biology* 11, 1294. doi:10.3390/biology11091294

Glass, D. S., and Alon, U. (2018). Programming cells and tissues. *Science* 361, 1199–1200. doi:10.1126/science.aav2497

Goodfellow, I., Bengio, Y., and Courville, A. (2016). *Deep learning*. Cambridge, MA: MIT Press.

He, F., Murabito, E., and Westerhoff, H. V. (2016). Synthetic biology and regulatory networks: where metabolic systems biology meets control engineering. *J. R. Soc. Interface* 13, 20151046. doi:10.1098/rsif.2015.1046

Kauffman, S., Peterson, C., Samuelsson, B., and Troein, C. (2003). Random Boolean network models and the yeast transcriptional network. *Proc. Natl. Acad. Sci. U. S. A.* 100, 14796–14799. doi:10.1073/pnas.2036429100

Lim, W. A. (2022). The emerging era of cell engineering: harnessing the modularity of cells to program complex biological function. *Science* 378, 848–852. doi:10.1126/science.add9665

Marbach, D., Prill, R. J., Schaffter, T., Mattiussi, C., Floreano, D., and Stolovitzky, G. (2010). Revealing strengths and weaknesses of methods for gene network inference. *Proc. Natl. Acad. Sci. U. S. A.* 107, 6286–6291. doi:10.1073/pnas.0913357107

Margossian, C. C. (2019). A review of automatic differentiation and its efficient implementation. *WIREs Data Min. Knowl. Discov.* 9, e1305. doi:10.1002/widm.1305

Maynard Smith, J. (1978). Optimization theory in evolution. *Annu. Rev. Ecol. Syst.* 9, 31–56. doi:10.1146/annurev.es.09.110178.000335

Mishra, B., Sun, Y., Howton, T., Kumar, N., and Mukhtar, M. S. (2018). Dynamic modeling of transcriptional gene regulatory network uncovers distinct pathways during the onset of arabidopsis leaf senescence. *NPJ Syst. Biol. Appl.* 4, 35–44. doi:10.1038/s41540-018-0071-2

Misra, D. (2019). *Mish: a self regularized non-monotonic activation function*. arXiv 1908. doi:10.48550/arXiv.1908.08681

Parker, G. A., and Maynard Smith, J. (1990). Optimality theory in evolutionary biology. *Nature* 348, 27–33. doi:10.1038/348027a0

Patel, N., and Bush, W. S. (2021). Modeling transcriptional regulation using gene regulatory networks based on multi-omics data sources. *BMC Bioinforma.* 22, 200–219. doi:10.1186/s12859-021-04126-3

Poggio, T., Banburski, A., and Liao, Q. (2020). Theoretical issues in deep networks. *Proc. Natl. Acad. Sci. U. S. A.* 117, 30039–30045. doi:10.1073/pnas.1907369117

Rackauckas, C., Ma, Y., Martensen, J., Warner, C., Zubov, K., Supekar, R., et al. (2020). *Universal differential equations for scientific machine learning*. arXiv:2001. 04385.

Rackauckas, C., and Nie, Q. (2017). DifferentialEquations.jl—a performant and feature-rich ecosystem for solving differential equations in Julia. *J. Open Res. Softw.* 5, 15. doi:10.5334/jors.151

Radhakrishnan, A., Belkin, M., and Uhler, C. (2020). Overparameterized neural networks implement associative memory. *Proc. Natl. Acad. Sci. U. S. A.* 117, 27162–27170. doi:10.1073/pnas.2005013117