

RESEARCH ARTICLE

Estimating information in time-varying signals

Sarah Anhala Cepeda-Humerez¹, Jakob Ruess^{2,3}, Gašper Tkačik^{1*}

1 Institute of Science and Technology Austria, A-3400 Klosterneuburg, Austria, **2** Inria Saclay – Ile-de-France, F-91120 Palaiseau, France, **3** Institut Pasteur, F-75015 Paris, France

* gkacik@ist.ac.at



Abstract

Across diverse biological systems—ranging from neural networks to intracellular signaling and genetic regulatory networks—the information about changes in the environment is frequently encoded in the full temporal dynamics of the network nodes. A pressing data-analysis challenge has thus been to efficiently estimate the amount of information that these dynamics convey from experimental data. Here we develop and evaluate decoding-based estimation methods to lower bound the mutual information about a finite set of inputs, encoded in single-cell high-dimensional time series data. For biological reaction networks governed by the chemical Master equation, we derive model-based information approximations and analytical upper bounds, against which we benchmark our proposed model-free decoding estimators. In contrast to the frequently-used k-nearest-neighbor estimator, decoding-based estimators robustly extract a large fraction of the available information from high-dimensional trajectories with a realistic number of data samples. We apply these estimators to previously published data on Erk and Ca²⁺ signaling in mammalian cells and to yeast stress-response, and find that substantial amount of information about environmental state can be encoded by non-trivial response statistics even in stationary signals. We argue that these single-cell, decoding-based information estimates, rather than the commonly-used tests for significant differences between selected population response statistics, provide a proper and unbiased measure for the performance of biological signaling networks.

OPEN ACCESS

Citation: Cepeda-Humerez SA, Ruess J, Tkačik G (2019) Estimating information in time-varying signals. *PLoS Comput Biol* 15(9): e1007290. <https://doi.org/10.1371/journal.pcbi.1007290>

Editor: Alexandre V. Morozov, Rutgers University, UNITED STATES

Received: January 4, 2019

Accepted: July 29, 2019

Published: September 3, 2019

Copyright: © 2019 Cepeda-Humerez et al. This is an open access article distributed under the terms of the [Creative Commons Attribution License](https://creativecommons.org/licenses/by/4.0/), which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

Data Availability Statement: Data used in this paper has been previously published in two other publications (Granados AA et al, *PNAS* 115: 6088, 2018; and Selimkhanov J et al, *Science* 346: 1370, 2014; both cited in the paper), and have been obtained by contacting the corresponding authors of those publications. We have produced no new data ourselves in this publication.

Funding: GT and SACH acknowledge the support of the Austrian Science Fund (<https://fwf.ac.at/en/>) grant FWF P28844. The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

Author summary

Cells represent changes in their own state or in the state of their environment by temporally varying the concentrations of intracellular signaling molecules, mimicking in a simple chemical context the way we humans represent our thoughts and observations through temporally varying patterns of sounds that constitute speech. These time-varying concentrations are used as signals to regulate downstream molecular processes, to mount appropriate cellular responses for the environmental challenges, or to communicate with nearby cells. But how precise and unambiguous is such chemical communication, in theory and in data? On the one hand, intuition tells us that many possible environmental changes could be represented by variation in concentration patterns of multiple signaling chemicals; on the other, we know that chemical signals are inherently noisy at the molecular scale. Here we develop data analysis methodology that allows us to pose and answer

Competing interests: The authors have declared that no competing interests exist.

these questions rigorously. Our decoding-based information estimators, which we test on simulated and real data from yeast and mammalian cells, measure how precisely individual cells can detect and report environmental changes, without making assumptions about the structure of the chemical communication and using only the amounts of data that is typically available in today's experiments.

Introduction

For their survival, reproduction, and differentiation, cells depend on their ability to respond and adapt to continually changing environmental conditions. Environmental information must be sensed and often transduced to the nucleus, where an appropriate response is initiated, usually by selectively up- or down-regulating the expression levels of target genes. This information flow is mediated by biochemical reaction networks, in which concentrations of various signaling molecules code for different environmental states or different response programs. This map between environmental input or response output and the internal chemical state is, however, highly stochastic, because typical networks operate with small absolute copy numbers of signaling molecules [1]; moreover, different environments can be encoded by the same signaling molecule, by differentially regulating the dynamics of its concentration [2]. This raises two fundamental questions: first, how much information the cells could, even in principle, encode in the combinatorial and possibly time-varying concentrations of multiple signaling molecules and how such information could be plausibly read out during “downstream” processing; and second, how can we quantify, in an unbiased and model-free fashion, the amount of information available to the cells from limited experimental data.

Information theory provides a framework within which the theoretical study of limits to communication as well as the empirical study of actual information flows can be addressed [3]. Applications of information theory to questions in biology and, in particular, neuroscience started already in the 1950s and continue to this day, with the main focus to understand how—and with what accuracy—neural activity encodes information about the environment [4–6]. Applications of analogous techniques to biochemical signaling only started recently and represent an active area of research at the interface of physics, biology, statistics, and engineering [7–10].

Recent theoretical work analyzed the reliability of information transmission through specific reaction systems in the presence of molecular noise, e.g., during ligand binding [11], in chemotaxis [12], gene regulation [13–19], biochemical signaling networks [20], etc., and asked how such transmission can be maximized by tuning the reaction rates. Generally, these studies focused on steady state, by considering the information encoded in a single temporal snapshot of the reaction network at equilibrium given the input signals. Rigorous extensions to dynamical signals have been either rare and only possible for simple cases, like the BIND channel [11], or required specific operating regimes that permitted linearization and Gaussianity assumptions [12, 21, 22]. At its core, the analysis of signal transduction through nonlinear noisy chemical systems requires one to have control over the distribution of concentration trajectories given the (possibly) time-varying inputs; even if it were possible to calculate this distribution in principle, the curse of dimensionality puts strong limits to the manipulations required to compute the information transmission. Consequently, problems of this kind are currently considered intractable in their full generality.

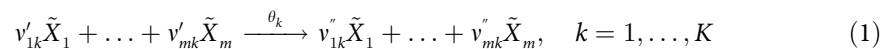
Empirical estimates of information transmission in biochemical networks similarly focused on the steady state [23, 24], or considered only specific, hand-picked dynamical features, such

as the amplitude or the frequency of the response, as information carriers [25]. Recent developments of fluorescent reporters and microfluidics have enabled the characterization of dynamical responses at a single cell resolution using large ($> 10^4$) numbers of sampled response trajectories, thereby permitting direct information estimates using generic estimators like the k -nearest-neighbors (knn) [26]. Existing approaches, however, suffer from severe limitations: they still require a prohibitive number of samples, especially when the response is distributed over multiple chemical species; or they necessitate uncontrolled assumptions about trajectory features that are supposed to be “relevant”. We recently proposed and applied decoding-based information estimators [27] as an alternative that draws on the past experiences in neuroscience [28–30] to dissect the yeast stress-response network. In this paper we provide a detailed account of the new methodology, show that it alleviates the most pressing problems of existing approaches, and benchmark it against synthetic and real data.

Models and methods

Biochemical reaction networks

At their core, cellular processes consist of networks of chemical reactions. A chemical reaction network consists of a set of m molecular species $\{\tilde{X}_1, \tilde{X}_2, \dots, \tilde{X}_m\}$ that interact through K coupled reactions of the form:



where v'_{1k}, \dots, v'_{mk} and $v''_{1k}, \dots, v''_{mk}$ are coefficients that determine how many molecules of each species are consumed and produced in the k -th reaction. $\theta_1 \dots \theta_k \in \mathbb{R}^+$ determine the rates at which the reactions occur and depend on binding affinities of chemical species, temperature and possibly the external conditions.

If we assume that the system is well-stirred, in thermal equilibrium and the reaction volume is constant, it can be shown that the probability that a reaction of type k takes place in an infinitesimal time interval $[t, t + dt]$ can be written as $a_k(\tilde{x}, \theta_k)dt = \theta_k g_k(\tilde{x})dt$, where $\tilde{x} = [\tilde{x}_1, \dots, \tilde{x}_m]^T \in \mathbb{N}_0^m$ contains the amounts of molecules of the m species that are present in the system at time t , and $g_k(\tilde{x}) = \prod_{i=1}^m \binom{\tilde{x}_i}{v'_{ik}}$ counts all possibilities of choosing the required reaction molecules out of all available molecules [31, 32]. θ_k is a constant that depends on the physical characteristics of the cell but also on the environmental conditions.

Let us denote the probability that \tilde{x} molecules of the m species are present in the system at time $t \in \mathbb{R}^+$ by $p(\tilde{x}, t)$ and define the stoichiometric change vectors

$v_k = [v_{1k}, \dots, v_{mk}]^T \in \mathbb{Z}^m$, $k = 1, \dots, K$, as the net changes in the amount of molecules in the reactions, i.e. $v_{ik} = v''_{ik} - v'_{ik}$, $i = 1, \dots, m$, $k = 1, \dots, K$. Then it can be shown [32] that the chemical master equation (CME) can be written as:

$$\dot{p}(\tilde{x}, t) = -p(\tilde{x}, t) \sum_{k=1}^K a_k(\tilde{x}, \theta_k) + \sum_{k=1}^K p(\tilde{x} - v_k, t) a_k(\tilde{x} - v_k, \theta_k), \quad (2)$$

or in a more compact form [32]

$$\dot{\mathbf{p}}(t) = \mathbf{M}\mathbf{p}(t), \quad (3)$$

where $\mathbf{p}(t)$ is a vector with components $p(\tilde{x}, t)$, which is, in principle, infinite dimensional, and \mathbf{M} contains the transition rates between all possible states, e.g. the transition rate from

state $\tilde{x}'_k = \tilde{x} - v_k$ to state \tilde{x} is given by

$$M_{\tilde{x}, \tilde{x}'_k} = a_k(\tilde{x}'_k, \theta_k) - \delta_{\tilde{x}, \tilde{x}'_k} \sum_q a_q(\tilde{x}, \theta_q), \tag{4}$$

where δ is the Kronecker delta, which is 1 when $\tilde{x} = \tilde{x}'_k$ and 0 otherwise.

The CME given in Eq (3) is an instance of a continuous-time discrete-state-space Markov Chain for a random process X that can be solved exactly only for a few simple cases. It is nevertheless possible to efficiently generate samples x of the random process X , which we will refer to as “trajectories” or “paths”, for a selected time interval, $t \in [0, T]$, according to the correct probability distribution p , by the Stochastic Simulation Algorithm (SSA, or the Gillespie algorithm) [33].

To study information transmission through the biochemical networks described by the CME, we need to define the input and output signals. In the simplest setup considered here, the input U is a discrete random variable that can take on one of the $q \geq 2$ possible values, $U \in \{u^{(1)}, u^{(2)}, \dots, u^{(q)}\}$. Each input in general corresponds to a distinct set of reaction rate constants θ , but in models of real biological networks, changing input often modulates only one or a few rates in the system, e.g., by representing the change in a key external ligand concentration, receptor activity, etc. Changes in the input are reflected in changes in the resulting trajectories of chemical species amounts, x . Typically, only a subset of chemical species could be considered as biologically-relevant “outputs” that encode the information about the environmental change: this would correspond to marginalizing p in Eq (3) over the unobserved (non-output) chemical species for the purposes of information transmission. While this is an interesting theoretical problem in its own right, here we work with simple toy examples where the output will be the trajectory, x , over the complete state space, i.e., we assume that all chemical species in the reaction network can be fully and perfectly observed. As we explain below, this allows us to define and compute the mutual information between a discrete input, U , and the output random process X given by the CME in a straightforward fashion. We later show that this computation can be carried out also when the continuous-time process X is sampled at uniform discrete times, as would be the case with experimental measurements.

Mutual information between discrete inputs and response trajectories

Information theory introduces the mutual information as the measure of fidelity by which changes in one random variable, e.g., the input U , can effect changes in another random variable, e.g., X . In this sense, mutual information is simply a measure of statistical dependency (i.e., any correlation, be it linear or not) between U and X , and can thus be written as a functional of the joint probability density function $p(x, u)$:

$$I(X; U) = \int_X \int_U p(x, u) \log_2 \left(\frac{p(x, u)}{p_X(x)p_U(u)} \right) du dx \tag{5}$$

where p_U and p_X are the marginal density functions for U and X , respectively, and we have generically written u and x as continuous variables; if they are discrete, integral signs are replaced by summations over the support for the corresponding probability distributions, as appropriate.

Mutual information is a non-negative symmetric quantity that is measured in bits, and is zero only if X and U are statistically independent. When studying information transmission through a channel $U \rightarrow X$ specified by $p(x|u)$, for which U serve as inputs drawn from an input

distribution $p_U(u)$, it is common to rewrite Eq (5) as

$$I(X; U) = H(U) - H(U|X) = H(X) - H(X|U), \tag{6}$$

where $H(X)$ is the differential entropy of X (and analogously for $H(U)$), defined as

$$H(X) = - \int_X p_X(x) \log_2 p_X(x) dx, \tag{7}$$

and the conditional entropy, $H(X|U)$, is

$$H(X|U) = \int_U H(X|u) p_U(u) du = - \int_U \int_X p_U(u) p(x|u) \log_2 p(x|u) dx du. \tag{8}$$

Eq (6) can be interpreted in two ways: information is either the difference between the total variability in the repertoire of responses X that the biochemical network can generate (measured by the *response entropy*, $H(X)$) and the average variability at fixed input that is due to noise in the network (measured by the *noise entropy*, $H(X|U)$); alternatively, information is also the entropy of the inputs, $H(U)$, minus equivocation $H(U|X)$, or the average uncertainty in what input was sent given that a particular response was observed. These interpretations make explicit the dependence of information both on the properties of the channel (the biochemical reaction network), as well as on the distribution of signals p_U that the network receives. In this work, we will consider discrete inputs and will assume uniform p_U . It is, however, also possible to compute the *channel capacity* C by maximizing the information flow at given $p(x|u)$ over all possible input distributions,

$$C = \max_{p_U} I(X; U); \tag{9}$$

Shannon’s classic work then proves that error-free transmission at rates higher than those given by capacity is impossible, while error-free transmission at rates below capacity can be achieved with the optimal use of the channel. Contrary to engineering, where the focus is on finding encoding and decoding schemes that best utilize a given channel, in biophysics and systems biology mutual information is used as a tool to quantify the limits to biological signal processing due to noise without needing to make assumptions about possible biochemical encoding and decoding mechanisms.

The setup we consider here is one in which inputs U are drawn independently from a uniform distribution and change rarely, i.e., at a rate that is much lower than the (inverse) timescale on which the reaction network in Eq (1) relaxes to its steady state. We assume that after an input change, we observe a fixed-time segment of the complete network dynamics, x , which is a sample path in m -dimensional discrete space, making direct calculation of information, $I(X; U)$, by integrating / summing over all possible trajectories as implied by Eq (5) intractable. We will nevertheless show that estimates of exact information are possible if the reaction network is known, by explicitly using the transition matrix M of the Markov Chain from Eq (3) and generating exact sample paths, that is, realizations of X , using SSA. We call this model-based approach *exact Monte Carlo approximation* and contrast it to uncontrolled model-free estimations such as those obtained by using Gaussian approximations or k -nearest-neighbors methodology. We then introduce various decoding estimators and establish a hierarchy through which these estimates upper and lower-bound the true information, as shown in Fig 1.

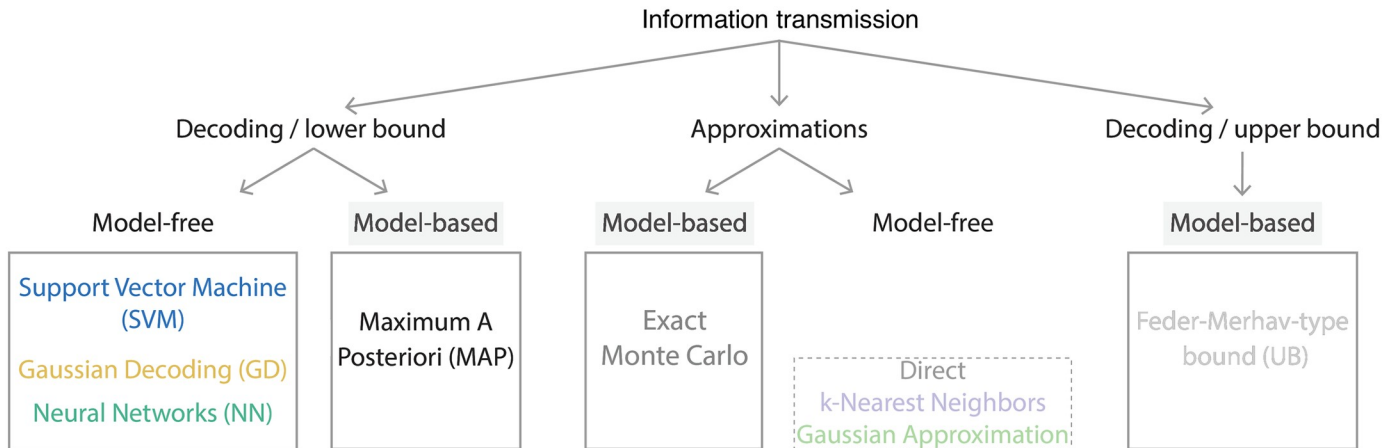


Fig 1. Information transmission between discrete inputs and response trajectories in biochemical networks. For fully-observed reaction networks whose dynamics are governed by a known chemical Master equation, information can be approximated to an arbitrary accuracy via Monte Carlo integration for either continuous-time or discrete-time response trajectories (model-based *exact Monte Carlo*, Section *Exact information calculations for fully observed reaction networks*). Since full knowledge of the reaction system is used, these approximations are tractable deep in the regimes where model-free estimations break down with uncontrolled errors (Section *Model-free information estimators*). True information estimates are lower-bounded by model-based maximum a posteriori (MAP) or Bayes optimal decoding (Section *Decoding-based information bounds*). This decoding gives the lowest average probability of error and the corresponding information lower bound can be used as a benchmark for information estimates derived from other model-free decoding approaches (that have at least the error probability of the MAP decoder); in Section *Decoding-based information estimators* we compare Support Vector Machine (SVM), Gaussian Decoding (GD) and Neural Network (NN) decoding approaches. Upper bounds like the Feder-Merhav bound [34] and our improvement on it [35] complete the picture by estimating the gap between optimal decoding-derived and exact information values (Section *Decoding-based information bounds*).

<https://doi.org/10.1371/journal.pcbi.1007290.g001>

Exact information calculations for fully observed reaction networks

Responses in continuous time. Given the specification of the biochemical reaction network in Eq (1), we sample N trajectories, x , using the Gillespie (SSA) algorithm. Each trajectory x can be represented as the sequence of consecutive states representing molecular species counts, $\mathbf{s} = [s_1, s_2, \dots, s_r]$, where $s_1 = \tilde{x}(t = 0)$, etc., and $s_r = \tilde{x}(t = \sum_{i=1}^{r-1} t_i)$, and the sequence of time intervals spent in each state, $\mathbf{t} = [t_1, \dots, t_r]$, $0 < t_i < T$, $i = 1, \dots, r$ and $T = \sum_i t_i$. We recall that SSA generates sample trajectories from the correct conditional distribution, $p(x|u)$. Given a particular trajectory x drawn from this distribution, we can compute the likelihood of x for a given input u exactly, by explicitly multiplying through all reaction events indexed by r :

$$p(x|u) = p(s_1) \exp(M_{s_1 s_1} t_1) \prod_{i=2}^r M_{s_i s_{i-1}} \exp(M_{s_i s_i} t_i) \tag{10}$$

where $p(s_1)$ is given by the initial conditions of the process, and the transition matrix \mathbf{M} depends on the input u . To get the marginal distribution, $p_X(x)$, we sum over all possible input values (since we are considering cases where the number of distinct inputs is finite, this summation can be done explicitly for each x):

$$p_X(x) = \sum_{i=1}^q p(x|u^{(i)}) p_U(u^{(i)}). \tag{11}$$

Since we are able to compute the exact likelihoods (Eqs (10) and (11)) for each path generated by the stochastic process, entropic quantities can be approximated without significant biases using Monte Carlo integration, where the integral over states in Eq (7) is replaced by an

average over N sampled trajectories:

$$\tilde{H}(X) = -\frac{1}{N} \sum_{i=1}^N \log_2 p_X(x_i). \tag{12}$$

Similarly, we can approximate $H(X|U)$:

$$\tilde{H}(X|U) = -\sum_{j=1}^q p_U(u^{(j)}) \left(\frac{1}{N} \sum_{i=1}^N \log_2 p(x_i|u^{(j)}) \right). \tag{13}$$

The exact Monte Carlo information approximation is finally obtained using Eq (6):

$$I_{\text{exact}}^* = \tilde{H}(X) - \tilde{H}(X|U), \tag{14}$$

where * reminds us that the paths are represented in continuous time. Taken together, this procedure relies on three key facts: first, that trajectories from the biochemical stochastic process determined by the Master equation can be drawn exactly using SSA; second, that the computation of the log likelihood for these trajectories is straightforward to evaluate from the known Master equation; and third, that a Monte Carlo approximation to the integrals for entropic terms is just an empirical average of the log likelihoods over the sampled trajectories.

Responses in discrete time. We can resample the continuous trajectories X on a grid of uniformly spaced time points to obtain a new discrete random variable, $\mathbf{X} = [X(t = \Delta t), \dots, X(t = i\Delta t), \dots, X(t = d\Delta t)] \in \mathbb{N}_0^{m \times d}$, where Δt is the discretization step, $d = T/\Delta t$ is the length of \mathbf{X} . For convenient notation we denote this random variable as $\mathbf{X} = [X^0, \dots, X^d]$, and its realizations, the discrete trajectories, as \mathbf{x} .

In the discrete case, the likelihood of \mathbf{x} for a given input u can be computed using the general solution to Eq (3):

$$\mathbf{p}(t) = e^{\mathbf{M}t} \mathbf{p}(0), \tag{15}$$

where $\mathbf{p}(t)$ is the probability distribution of states after time t , with the initial probability distribution $\mathbf{p}(t = 0) = \mathbf{p}(0)$. Using this formal solution we compute the transition matrix between discrete timesteps separated by Δt to get:

$$\mathbf{W} = e^{\mathbf{M}\Delta t}, \tag{16}$$

where \mathbf{M} and thus \mathbf{W} again depend on u . The likelihood of any discrete path can then be obtained by multiplying the transition probabilities between all the d consecutive states in the path for a given input u :

$$p(\mathbf{x}|u) = p(\mathbf{x}^0) \prod_{i=2}^d \mathbf{W}_{\mathbf{x}^i \mathbf{x}^{i-1}}. \tag{17}$$

We can now approximate the information between input U and a discretely sampled trajectory \mathbf{X} , I_{exact} , as in the continuous case: we get the marginal $p_X(\mathbf{x})$ with Eq (11) and use Eqs (12) and (13) in Eq (14). In general, temporal discretization loses information relative to the full (continuous-time) trajectory, where reaction events in the trajectory x are recorded with infinite temporal precision, so the information in discretely-sampled trajectories, I , must be bounded from above by the information in continuous-time trajectories, I^* :

$$I_{\text{exact}}(\mathbf{X}; U) \leq I_{\text{exact}}^*(X; U), \tag{18}$$

where equality is approached in the limit of ever finer temporal discretization, $\Delta t \rightarrow 0$.

Model-free information estimators

In the absence of a full stochastic model for the biochemical reaction network, mutual information estimation is tractable only if we make assumptions about the distribution of response trajectories given the input. We briefly summarize two approaches below: in the first, *k*-nearest-neighbor procedure, the space in which the response trajectories are embedded is assumed to have a particular metric; in the second, Gaussian approximation, we assume a particularly tractable functional form for the channel, $p(\mathbf{x}|u)$.

K-nearest-neighbors (knn) estimator. The idea of using the nearest neighbour statistics to estimate entropies is at least 70 years old [36, 37], while estimators for mutual information have been developed during the early 2000s [38, 39]. The cornerstone of the approach is to compute the estimate from the distances of *d*-dimensional real valued data points to their *k*-th nearest neighbour. Hence, the estimator depends on the metric chosen to define this distance. Furthermore, its performance is known to depend on the value of *k* (number of nearest neighbours), where small *k* increase the variance and decrease the bias [40]. This method has been used in several studies that estimated mutual information from single cell time series [24, 26, 41]. These studies used large numbers of response trajectories to provide the first evidence that the information available from the full timeseries of the response could be substantially higher than the information available from any response snapshot.

Gaussian approximation. A simplifying assumption in the Gaussian approximation is that the distribution of trajectories sampled at discrete times given input is approximately Gaussian, with the mean $\mu \in \mathbb{R}^d$ and covariance matrix $\Sigma \in \mathbb{R}^{d \times d}$ that may both depend on the input, *u*:

$$p(\mathbf{x}|u) = \mathcal{N}(\mathbf{x}; \mu(u), \Sigma(u)) = \frac{1}{\sqrt{\det(2\pi\Sigma)}} \exp\left(-\frac{1}{2}(\mathbf{x} - \mu)^T \Sigma^{-1}(\mathbf{x} - \mu)\right). \tag{19}$$

The entropy of the multivariate distribution in Eq (19) has an analytical expression that only depends on Σ :

$$H_G(\mathbf{X}|u) = \frac{1}{2} \log(\det(2\pi e \Sigma(u))), \tag{20}$$

which can be averaged over $p_U(u)$ to get the conditional entropy, $H_G(\mathbf{X}|U)$. To estimate the information, we further need $H(\mathbf{X})$ from Eq (6). This entropy of a Gaussian mixture has no closed form solution, but can be computed by Monte Carlo integration as in the previous section, following discrete analogs of Eqs (11) and (12): we draw random samples from each of the *q* different multivariate Gaussian distributions, Eq (19), one for each possible input *u*, and assign the marginal probabilities to each sample \mathbf{x} as

$$p_X(\mathbf{x}) = \sum_{i=1}^q p_U(u^{(i)}) \mathcal{N}(\mathbf{x}; \mu(u^{(i)}), \Sigma(u^{(i)})), \tag{21}$$

permitting us to use Eq (12) to approximate the total entropy of output trajectories in the Gaussian approximation, $H_G(\mathbf{X})$, and thus to obtain the Gaussian estimate for the information, $I_G(\mathbf{X}; U) = H_G(\mathbf{X}) - H_G(\mathbf{X}|U)$.

To apply this estimator, one must use real (or simulated) data to estimate the conditional mean, $\mu(u)$, and conditional covariance, $\Sigma(u)$ for every possible *u*, from a limited number of samples. While general caveats for such estimations have been detailed in many textbooks [42], we emphasize that information estimation is particularly sensitive due to the computation of the determinant in Eq (20) which can easily lead to ill-posed numerics when the

number of samples is small. This can be mitigated by various regularization methods (one of which, the diagonal regularization [64], we demonstrate later) that impose a prior structure on the estimated covariance. Yet even in the case of significant oversampling that we can explore using simulated data, the Gaussian approximation introduced here—in contrast to Gaussian decoding estimator introduced in the next section—can provide information values that deviate significantly from the true value and are not guaranteed to bound the true value from either above or below. This is because the true solutions of the CME live in the positive quadrant of the discrete space, and are thus essentially different from the Gaussian distributions assumed here. We nevertheless present this estimator because (i) it forms the basis for the Gaussian decoding estimator, introduced below, and (ii) real data itself often deviates from stochastic trajectories sampled from the CME in that it is continuous (since we measure, e.g., fluorescence proxy for a concentration of a protein of interest) and contains extra noise, making Gaussian approximation potentially applicable.

Decoding-based information bounds

Here and in the next section we introduce a class of decoding-based calculations that lower-bound the exact information, $I(\mathbf{X}; U)$, and can tractably be used as information estimators over realistically-sized data sets. Let \mathcal{D} consist of a set of N labeled paths, typically represented in discretely sampled time, $\mathcal{D} = \{(u_1, \mathbf{x}_1), (u_2, \mathbf{x}_2), \dots, (u_N, \mathbf{x}_N)\}$, where u_i and \mathbf{x}_i , for $i = 1, \dots, N$, are realizations of the random variables $U \in \{u^{(1)}, \dots, u^{(q)}\}$ and $\mathbf{X} \in \mathbb{R}^{m \times d}$, respectively. Here, \mathcal{D} can represent either real data (typically containing $N \sim 10^2 - 10^3$ trajectories) in case of model-free information estimates, or trajectories generated by exact simulation algorithms (in which case the sample size, N , is not limiting) from the full specification of the biochemical reaction network in case of model-based approximations.

The procedure of estimating the input \hat{u} from \mathbf{x} , such that the estimated \hat{u} is “as close as possible” to true u for a given trajectory \mathbf{x} , is known as decoding in information theory and neuroscience, and can equivalently be viewed as a classification task in machine learning or as an inference task in statistics. This procedure is implemented by a decoding function,

$$\hat{u} = F_\omega(\mathbf{x}); \tag{22}$$

F is typically parametrized by parameters ω that need to be learned from data for model-free approaches, or derived from biochemical reaction network specification in case of model-based approaches. F assigns to every \mathbf{x}_i in the dataset a corresponding “decode” \hat{u}_i from the same space over which the random variable U is defined; formally, these decodes are instances of a new random variable \hat{U} . The key idea of using decoding for information estimation starts with the observation that random variables

$$U \rightarrow X \xrightarrow{T_d} \mathbf{X} \xrightarrow{F_\omega} \hat{U}, \tag{23}$$

where T_d represents time discretization, form a Markov chain. In other words, the distribution of \hat{U} is conditionally independent of U and only depends on \mathbf{X} , $p(\hat{u}|\mathbf{x}, u) = p(\hat{u}|\mathbf{x})$, and so

$$p(\hat{u}, \mathbf{x}, u) = p_U(u)p(\mathbf{x}|u)p(\hat{u}|\mathbf{x}). \tag{24}$$

The data processing inequality [43] can be used to further extend the bounds in Eq (18):

$$I(U; \hat{U}) \leq I_{\text{exact}}(U; \mathbf{X}) \leq I_{\text{exact}}^*(U; X), \tag{25}$$

where equality between the first two terms holds only if $I(U; \mathbf{X}|\hat{U}) = 0$. Consequently,

$I(U; \hat{U})$ is a lower bound to the information between trajectories X and the input U [44]. Note that analogous reasoning holds for decoding directly from continuous-time trajectories X . Better decoders which increase the correspondence between the true inputs and the corresponding decoded inputs will typically provide a tighter lower bound on the information.

To compute the information lower bound, we apply the decoding function to each trajectory in \mathcal{D} in model-based approximations or to each trajectory in the testing dataset for model-free estimators that need to be learned over training data first. We subsequently construct a $q \times q$ confusion matrix, also known as an error matrix, where each element ϵ_{ij} counts the fraction of realizations of \mathbf{x} generated by an input $u = u^{(i)}$ that decode into $\hat{u} = u^{(j)}$. This matrix provides an empirical estimate of the probability distribution $p(\hat{u}, u)$, which can thus be used to compute the information estimate:

$$I(\hat{U}; U) = \sum_{u, \hat{u}} p(\hat{u}, u) \log_2 \frac{p(\hat{u}, u)}{p_U(u)p_{\hat{U}}(\hat{u})} \approx \sum_{i=1}^q \sum_{j=1}^q \epsilon_{ij} \log_2 \frac{\epsilon_{ij}}{(\sum_k \epsilon_{kj})(\sum_l \epsilon_{il})}, \quad (26)$$

Crucially, in this estimation $O(N)$ data points are used to empirically estimate the elements of a $q \times q$ matrix ϵ , and information estimation involves a tractable summation over these matrix elements; in contrast, direct estimates of $I(U; \mathbf{X})$ would involve an intractable summation over (vastly undersampled) space for \mathbf{X} . For typical applications where q is small, decoding thus provides an essential dimensionality reduction prior to information estimation: in a simple but biologically relevant case of two distinct stimuli ($q = 2$), information estimation only requires us to empirically construct a 2×2 confusion matrix. If required, one can apply well-known debiasing techniques for larger q [5].

Maximum a posteriori (MAP) lower bound. In MAP lower bound, the decoding function F_ω is given by Bayesian inference of the most likely input u given that a response trajectory \mathbf{x} was observed, under the exact probabilistic model for the biochemical reaction network. MAP decoder is optimal in that it provides the lowest average probability of error, $\Pr(\hat{U} \neq U)$, among all decoders. Typically, this will lead to a high mutual information value $I(\hat{U}; U)$ compared to other (sub-optimal) decoders whose probability of error will likely be higher, making the information lower bound from MAP decoder a good benchmark for other decoder-based information estimates. We remind the reader, however, that even though MAP decoder achieves minimal error and typically high $I(\hat{U}; U)$ values, this does not mathematically guarantee that its information will *always* be higher or equal to the information of any other possible decoder, a fact that can be demonstrated explicitly using toy examples.

The MAP inference consists of finding the input that maximizes the posterior distribution [45]

$$p(u|\mathbf{x}) = \frac{p(\mathbf{x}, u)}{p_{\mathbf{X}}(\mathbf{x})} = \frac{p(\mathbf{x}|u)p_U(u)}{p_{\mathbf{X}}(\mathbf{x})}. \quad (27)$$

This corresponds to the following decoding function:

$$\hat{u} = F_\omega(\mathbf{x}) = \operatorname{argmax}_u [\log p(\mathbf{x}|u) + \log p_U(u)], \quad (28)$$

where ω represents the specification of the biochemical reaction network which determines $p(\mathbf{x}|u)$. Here, $p_U(u)$ is assumed to be known, and the likelihood $p(\mathbf{x}|u)$ can be calculated using Eqs (10) or (17), for the continuous or discrete time representations, respectively.

One can apply the MAP-decoding based calculation of $I_{\text{MAP}}(\hat{U}; U)$ in two ways. First, when applied over real data \mathcal{D} , one can think of the procedure as a proper statistical estimation assuming that the biochemical network model is the correct generative model of the data (with

estimation bias arising if it is not). Second, when applied, as we will do in the Results section, over trajectories \mathcal{D} generated using exact stochastic simulation from the biochemical network model in the large N limit, this procedure is a Monte Carlo approximation to the information lower bound.

Note that even though the MAP decoder is optimal, it does not follow that $I_{\text{MAP}}(\hat{U}; U) = I(\mathbf{X}; U)$. This is because optimal channel use that realizes $I(\mathbf{X}; U)$ may need to employ block codes, where a *sequence* of inputs is encoded jointly into a *sequence* of trajectories, which is later also jointly decoded. In contrast, the decoding bound $I_{\text{MAP}}(\hat{U}; U)$ relies on one-shot use of the channel: a single input u is transduced into \mathbf{x} which can immediately be decoded back into the estimate of the input, \hat{u} , on the basis of which the cell might make a decision. For many biological situations, this decoding setup should be more appropriate than the exact information calculation, as cells often need to react to stimuli as rapidly as possible in order to gain a selective advantage. Furthermore, it is difficult to conceive of biologically realistic encoders that would transform inputs into a block code in order to use the biochemical network channels optimally.

Maximum a posteriori upper bound (UB). Given that the optimal MAP decoding does not necessarily reach the exact mutual information, it is reasonable to ask how large the gap is between these two quantities. For discrete inputs, classic work in information theory proved a number of upper bounds on this gap when the channel is known [46], with the Feder-Merhav bound perhaps being the most well known [34]; Feder-Merhav provides an upper bound on the channel capacity given the overall probability of error in MAP decoding. In a separate work [35], we computed a new upper bound on information $I_{\text{UB}}(U; \mathbf{X})$ that is consistent with not just the overall probability of error as in Feder-Merhav bound, but with the full confusion matrix ϵ obtained from optimal MAP decoding, and showed that the new bound is tight.

Our self-contained derivation [35] gives the following result

$$I(U; \mathbf{X}) \leq I_{\text{UB}} = H(U) - \sum_{\hat{u}} p_{\hat{U}}(\hat{u}) \phi(\pi_{\hat{u}}), \tag{29}$$

where $\pi_{\hat{u}} = \Pr(U \neq \hat{U} | \hat{u}) = 1 - \Pr(U = \hat{u} | \hat{U} = \hat{u})$ and functions ϕ and α can be expressed with the help of the floor and ceiling functions as:

$$\phi(\pi) = \alpha(\pi) \log_2 \left\lfloor \frac{1}{1 - \pi} \right\rfloor + (1 - \alpha(\pi)) \log_2 \left\lceil \frac{1}{1 - \pi} \right\rceil \tag{30}$$

$$\alpha(\pi) = \left\lfloor \frac{1}{1 - \pi} \right\rfloor \left((1 - \pi) \left\lceil \frac{1}{1 - \pi} \right\rceil - 1 \right). \tag{31}$$

This bound applies irrespectively of how the response trajectory space is represented (continuous or discrete, possibly of dimensionality much larger than that of the random variable U), since it is stated solely in terms of the input variable U and its MAP decode, \hat{U} .

Decoding-based information estimators

Support Vector Machine (SVM) lower bound estimator. The first model-free decoding approach we consider is based on classifiers called Support Vector Machines (SVMs). To begin we consider two possible inputs, $q = 2$. We define a decoding function F_{ω} by means of a helper function $f_{\omega}(\mathbf{x})$, such that $F_{\omega}(\mathbf{x}) = u^{(1)}$ if $\text{sign } f_{\omega}(\mathbf{x}) = -1$ and $F_{\omega}(\mathbf{x}) = u^{(2)}$

otherwise. Here,

$$f_{\omega}(\mathbf{x}) = \sum_{i=1}^{N_t} \alpha_i k(\mathbf{x}_i, \mathbf{x}) + b \tag{32}$$

where $k : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}$ is the so-called “kernel function” to be defined below, b is the bias constant, N_t is the number of samples in $\mathcal{D}_{\text{train}}$ and $\alpha_1, \dots, \alpha_{N_t}$ are obtained by solving standard SVM equations:

$$\min_{\substack{\alpha_1, \dots, \alpha_{N_t} \in \mathbb{R}, \\ \xi_1, \dots, \xi_{N_t} \in \mathbb{R}^+}} \sum_{i,j=1}^{N_t} \alpha_i \alpha_j k(x_i, x_j) + \frac{C}{N_t} \sum_{i=1}^{N_t} \xi_i \tag{33}$$

subject to

$$y_i \sum_{j=1}^{N_t} \alpha_j k(\mathbf{x}_j, \mathbf{x}_i) \geq 1 - \xi_i, \quad \text{for } i = 1, \dots, N_t. \tag{34}$$

$y_i = -1$ whenever the input corresponding to the i -th trajectory in the training set, \mathbf{x}_i , is $u^{(1)}$, i.e., $u_i = u^{(1)}$; similarly $y_i = +1$ whenever the corresponding input is $u^{(2)}$, i.e., $u_i = u^{(2)}$. C is a positive regularization constant. Together, the parameters of the decoding function are $\omega = \{b, \alpha, \xi, C\}$.

To prevent overfitting and set the regularization parameter C using cross-validation, we split the full dataset \mathcal{D} into training data, $\mathcal{D}_{\text{train}}$, that consists of N_t (here $\sim 70\%$ of the total, N) of labeled sample trajectories, chosen randomly but balanced across different inputs u ; the remaining 30% of the data constitutes the testing data, $\mathcal{D}_{\text{test}}$. Parameters ω are estimated only over $\mathcal{D}_{\text{train}}$, after which the error matrix ϵ and the corresponding information estimate $I_{\text{SVM}}(\hat{U}; U)$ of Eq (26) are evaluated solely over $\mathcal{D}_{\text{test}}$. The test/train split procedure can be repeated multiple times to compute the mean and the bootstrapped error bar estimate for the information estimator, I_{SVM} [27].

When we apply SVM decoding, we are still free to choose the kernel function. Here, we focus on two possibilities:

- **Linear kernel**, $k(x, x') = x^T x'$. The information estimate is based on a linear classifier that can learn to distinguish responses that differ in their conditional means, $\mu(u)$, but will result in close to chance performance if they don't. This is the simplest model-free decoding estimator and is thus a useful benchmark for more complex, non-linear decoders.
- **Radial basis functions kernel**, $k(x, x') = \exp\left(-\frac{\|x-x'\|^2}{2\sigma^2}\right)$. This model-free decoder can be sensitive both to difference in the conditional means as well as higher-order statistics, e.g., the covariance matrix. Parameter σ is set via cross-validation to maximize the performance.

For multiclass classification we use a decision-tree SVM classification method [47], also called Dendrogram-SVM (DSVM) [48]. To translate the multi-class classification into the canonical binary classification problem, this method uses hierarchical bottom-up clustering to define the structure of the graph, on which a binary classification is performed using SVMs at each graph node.

Gaussian decoder (GD) lower bound estimator. In this model-free estimation, we revisit the assumption that the (discretely sampled) output trajectories \mathbf{x} given input u can be

approximated with a multivariate Gaussian distribution, Eq (19). The decoding function is then

$$\hat{u} = F_{\omega}(\mathbf{x}) = \operatorname{argmax}_u [\log \mathcal{N}(\mathbf{x}; \mu(u), \Sigma(u)) + \log p_U(u)]. \quad (35)$$

Here, parameters ω consist of conditional means and (possibly regularized) covariance matrices of the Gaussian distributions that need to be estimated from data, following the test/train procedure analogous to SVM decoding.

This method can be used with different parametric multivariate probability density functions replacing the multivariate Gaussian in Eq (35), with choices that approximate the statistics of the data (and thus the CME-derived distribution) better providing tighter lower bound estimates, $I_{GD}(\hat{U}; U)$, of the exact information. By analogy with the exact MAP decoding using CME-derived response distribution, this method can also be understood as maximum a posteriori decoder but using approximate response distributions that need to be estimated from data. Here we decided to use the Gaussian distributions because they are the most unstructured (random) distributions based on measured first- and second-order statistics of the data. GD decoder thus should be able to discriminate various inputs if their responses differs either in the response mean or response covariance.

Neural network (NN) lower bound estimator. Artificial neural networks, first introduced by the neurophysiologist Warren McCulloch and the mathematician Walter Pitts in 1943 [49], are nowadays the method of choice for classification that generally outperforms alternative machine learning techniques on very large and complex problems. Here we use one of the simplest neural networks, called the multi-layer perceptron (MLP). MLP is composed of layers of linear-threshold units (or LTUs), where each LTV computes a weighted sum of its inputs $z = \omega^T \mathbf{x}$, then applies an activation function to that sum and outputs the result $y = h(z) = h(\omega^T \mathbf{x} + \omega_0)$. Using a single LTV amounts to training a binary linear classifier by learning the weights ω . As with linear SVM, such classifier only has a limited expressive power [50], which can, however, be extended by stacking layers of LTUs so that outputs of the first layer are inputs to the second layer etc.

For illustrative purposes we choose for our decoding function $F_{\omega}(\mathbf{x})$ a fully connected neural network with two hidden layers (with 300 and 200 LTUs, respectively) that uses the exponential activation function with $\alpha = 1$:

$$h_x(z) = \begin{cases} \alpha(\exp(z) - 1) & \text{if } z \leq 0 \\ z & \text{if } z > 0 \end{cases}.$$

For training, we used He-initialization, which initializes the weights with a random number from a normal distribution with zero mean and standard deviation $\sigma = 2/\sqrt{n_{in}}$, where n_{in} is the number of inputs to units in a particular layer [51], and Adam optimization with batch normalization and drop-out regularization [51, 52]. As before, we trained the neural network on \mathcal{D}_{train} , followed by the evaluation of the error matrix ϵ and of the corresponding information estimate, $I_{NN}(\hat{U}; U)$, from Eq (26), over \mathcal{D}_{test} . We emphasize that the detailed architecture of the neural network we selected here is not relevant for other estimation cases; in general, the architecture is completely adjustable to the problem at hand and should be selected depending on the size of the training dataset. The only selection criterion is the network performance on test data, with better performing networks for a given dataset typically providing tighter information estimates.

Results

Information estimation on simulated data

We start by considering three simple chemical reaction networks for which we can obtain exact information values using the model-based approach outlined in Section *Exact information calculations for fully observed reaction networks*. This will allow us to precisely assess the performance of decoding-based model-free estimates, and systematically study the effects of time discretization, the number of sample trajectories, and the number of distinct discrete inputs, q .

The three examples are all instances of a simple molecular birth-death process, where molecules of \tilde{X} are created and destroyed with rates α and β , respectively:



The reaction rates, α and β , will depend in various ways on the input, U , and possibly time, as specified below. Given an initial condition, $x(t=0)$, the production and degradation reactions generate continuous-time stochastic trajectories, $x(t)$, recording the number of molecules of \tilde{X} at every time $t \in [0, T]$, according to the Chemical Master Eq (3). These trajectories, or their discretized representations, are considered as the “outputs” of the example reaction networks, defining the mutual information $I(X; U)$ that we wish to compute. In all three examples we start with the simplest case, where the random variable U can only take on two possible values, $u^{(1)}$ and $u^{(2)}$, with equal probability, $p_U(u^{(1)}) = p_U(u^{(2)}) = 0.5$.

- Example 1.** In this case, $x(t=0) = 0$, $\beta = 0.01$, independent of the input U , and the production rate depends on the input as $\alpha(u^{(1)}) = 0.1$, $\alpha(u^{(2)}) = 0.07$. Here, the steady state is given by Poisson distribution with mean number of molecules $\langle x(t \rightarrow \infty) \rangle = \alpha/\beta$. Steady-state is approached exponentially with the timescale that is the inverse of the degradation rate, β^{-1} . These dynamics stylize a class of frequently observed biochemical responses where the steady-state mean expression level encodes the relevant input value. Even if the stochastic trajectories for the two possible inputs are noisy as shown in Fig 2A, we expect that the mutual information will climb quickly with the duration of the trajectory, T , since (especially in steady state) more samples provide direct evidence about the relevant input already at the level of the mean trajectories.
- Example 2.** In this case, $x(t=0) = 0$, $\beta = 0.01$, independent of the input U , and the production rate depends on the input as $\alpha(u^{(1)}, t) = 0.1$, $\alpha(u^{(2)}, t) = 0.05$ for all $t < 1000$, while for $t \geq 1000$ the production rate is very small and independent of input, $\alpha(u, t) = 5 \cdot 10^{-4}$. In the early period, this network approaches input-dependent steady state with means whose differences are larger than in Example 1, but the difference decays away for $t > 1000$ as the network settles towards vanishingly small activity for both inputs, as shown in Fig 2B. These dynamics stylize a class of transient biochemical responses that are adapted away even if the input state persists. In this case, lengthening the observation window T will not provide significant increases in information.
- Example 3.** In this case, $x(t=0) = 10$. All reaction rates depend on the input, $\alpha(u^{(1)}) = 0.1$, $\alpha(u^{(2)}) = 0.05$, $\beta(u^{(1)}) = 0.01$, $\beta(u^{(2)}) = 0.005$, and are chosen so that the mean $\langle x(t) \rangle = 10$ is constant across time and equal for both conditions, as shown in Fig 2C. In this difficult case, inputs cannot be decoded at the level of mean responses but require sensitivity to at least second-order statistics of the trajectories. Specifically, signatures of the input are present in the autocorrelation function for x : the timescale of fluctuations and mean-reversion is two-fold

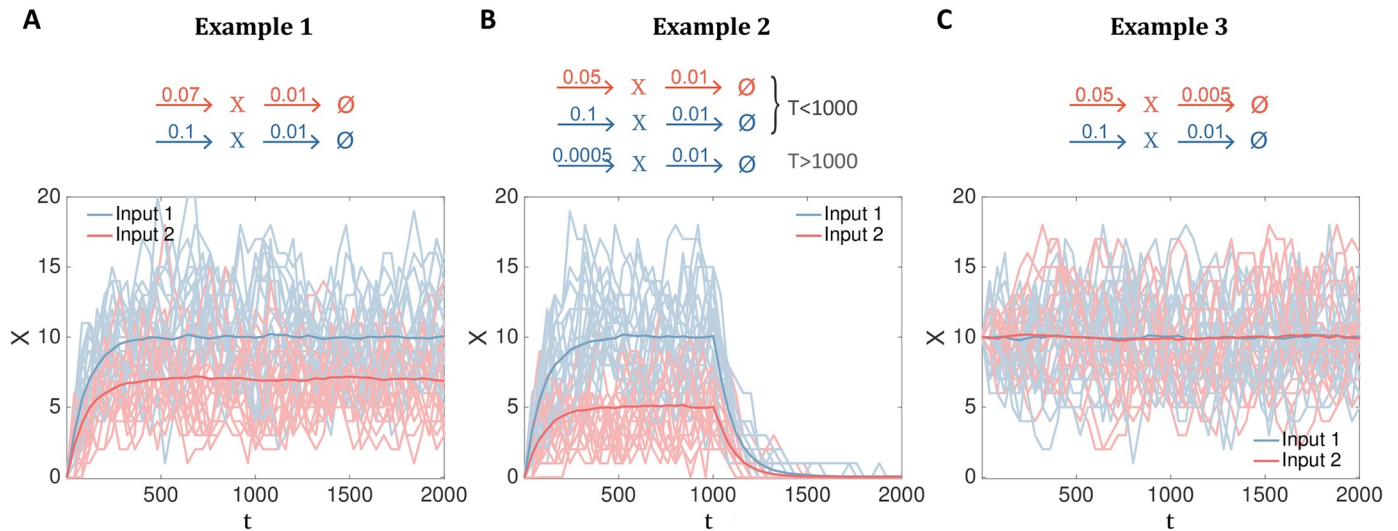


Fig 2. Example biochemical reaction networks and their behavior. Three example birth-death processes, specified by the reactions in the top row for each of the two possible inputs ($u^{(1)}$ in blue, $u^{(2)}$ in red), stylize simple behaviors of biochemical signaling networks. (A) Input is encoded in both the transient approach to steady state and the steady state value. (B) Input is encoded in the magnitude of the transient response which is subsequently adapted away. (C) Input is encoded only at the level of temporal correlations of the response trajectory. Bottom row shows example trajectories generated using the Stochastic Simulation Algorithm for the copy number of \bar{X} molecules, $t \in [0, 2000]$, for each network and the two possible inputs (light blue, light red); while plotted as a connected line for clarity, each trajectory represents molecular counts and is thus a step-wise function taking on only integer or zero values. Dark blue, red lines show the conditional means over $N = 1000$ trajectory realizations.

<https://doi.org/10.1371/journal.pcbi.1007290.g002>

faster for u^1 than u^2 . While this case is not frequently observed in biological systems, it represents a scenario where, by construction, no information about the input is present at the level of single concentration values and having access to the trajectories is essential. Because there is no difference in the mean response, we expect linear decoding methods to provide zero bits of information about the input. This case is also interesting because of the recent focus on pulsatile stationary-state dynamics in biochemical networks [53]. These pulses, reported for transcription factors such as Msn2, NF- κ B, p53, etc., occur stochastically and, when averaged over a population of desynchronized cells, can yield a flat and featureless mean response. Information about the stimulus could, nevertheless, be encoded in either the frequency, amplitude, or other shape parameters of the pulses. While a generative description of such pulsatile dynamics goes beyond a birth-death process considered here, from the viewpoint of decoding, both pulsatile signaling and our example present an analogous problem, where the mean response is not informative about the applied input.

Before proceeding, we note that our examples are not intended to be realistic models of intracellular biochemical networks, but are chosen here for their simplicity and analytical tractability, in order to benchmark model-free estimators against known “gold truth” standard. In particular, while our examples include intrinsic noise due to the stochasticity of biochemical reactions at low concentration, they do not include extrinsic noise or cell-to-cell variability which, in some systems, is known to importantly or even dominantly contribute to the total variability in the response [26, 54]. The presence of such additional sources of variation by no means makes the model-free estimators inapplicable, as we show in S1 Fig where we study estimator performance in the simplest Example 1 model that includes cell-to-cell variability; it solely prevents us from comparing their performance to a tractably-computable MAP decoder result.

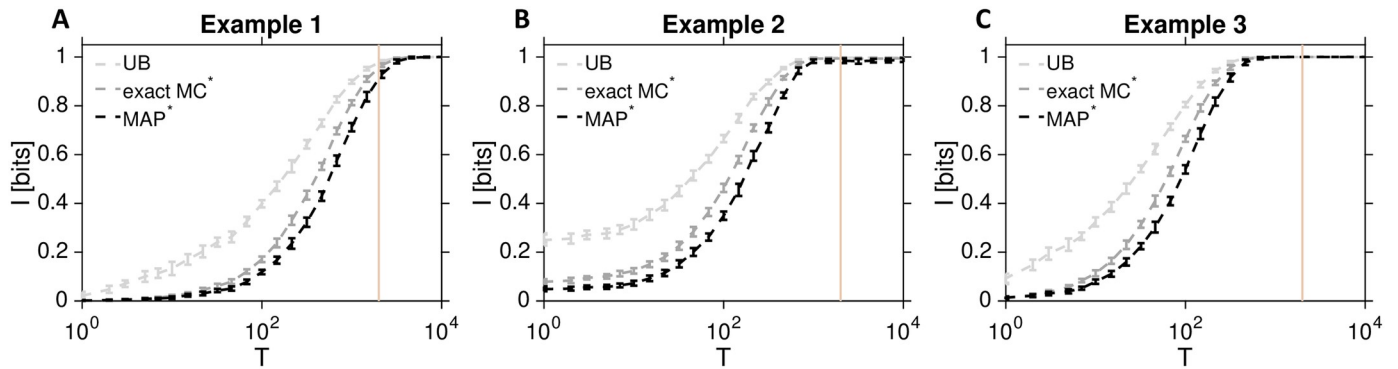


Fig 3. Information about inputs encoded by complete response trajectories of the example biochemical reaction networks. Exact Monte Carlo approximation for the information, $I_{\text{exact}}^*(X; U)$, is shown for Example 1 (A), Example 2 (B), and Example 3 (C) from Fig 2 in dashed dark gray line; error bars are standard deviations across 20 replicate estimations, each computed over $N = 1000$ independently generated sample trajectories per input condition. Information is plotted as a function of the trajectory duration, T ; yellow vertical line indicates $T = 2000$ as a representative duration used in further analyses below, at which most of the information about input is in principle available from the response trajectories of our systems. $I_{\text{MAP}}^*(\hat{U}; U)$ (dashed black line) is the optimal decoding lower bound, and I_{UB}^* (dashed light gray) is the upper bound on the information, computed by applying Eq (29).

<https://doi.org/10.1371/journal.pcbi.1007290.g003>

Exact information approximations and bounds for continuous and discrete trajectories. Armed with the full stochastic model for the three example reaction networks, we can compute the mutual information, $I_{\text{exact}}^*(X, U)$, between the continuous-time stochastic trajectories and the (binary) input variable U , following Eq (14). This result depends essentially on the length of the observed trajectory, $t \in [0, T]$, since T controls the number of observed reaction events and thus the accumulation of evidence for one or the other alternative input. As the approximation is implemented by Monte-Carlo averaging of exact log probabilities for the response trajectories, its variance will depend on the number of sample trajectories generated by the SSA. Because these information values will represent the “gold truth” against which to evaluate subsequent estimators, we choose a large number of $N = 1000$ trajectory realizations per input condition, and verify the tightness of the exact Monte Carlo approximation by computing the standard deviation over 20 independent re-runs of the approximation procedure.

Fig 3 shows how the exact Monte Carlo information computation depends on the trajectory duration, T , for each of the three example cases. As expected, the information increases monotonically with T towards the theoretically maximal value of 1 bit, corresponding to perfect information about two *a priori* equally likely input conditions. The exact shape of the information curve depends on the shape of the mean trajectory, as well as on its variance and higher-order statistics: for example, even though the two inputs for Example 1 are most distinct at the level of mean responses for later T values, the noise is higher compared to Example 2, such that at $T = 2000$ there is more total information in trajectories of Example 2 than Example 1. Conversely, even though the trajectories in Example 3 do not differ at the level of the mean at all, they still carry all information about the relevant input once sufficiently long trajectories can be observed (and assuming full knowledge of the reaction network is available).

One can similarly compute the Bayes-optimal or MAP decoding bound using Eq (28) for continuous trajectories. This quantifies the ultimate accuracy limit with which each single observed trajectory can be decoded into the input that gave rise to it. As demonstrated in Fig 3 in dashed black line and consistent with the Data Processing Inequality requirements outlined in the Methods, $I_{\text{MAP}}^*(\hat{U}; U) \leq I_{\text{exact}}^*(X; U)$. Equality is not reached because the optimal use of the channel requires block coding schemes, in contrast to our setting where different inputs

are sequentially sent through the biochemical network and immediately decoded. The observed gap between the MAP optimal decoding estimate and the true information appears to be small in each of the three cases; one can upper-bound the gap itself by an improvement over the standard Feder-Merhav calculation following Section *Maximum a posteriori upper bound*. While the resulting upper bound on information, I_{UB} , is not tight in this case, it nevertheless provides a control of how far optimal decoding could be from the true information estimate, a question that has repeatedly worried the neuroscience community facing similar problems [28]. It is worth noting that if MAP decoder can tractably be computed, so can the upper bound, irrespective of the dimensionality of the space of responses, X .

Fig 3 summarizes the absolute limits on information transmission and optimally decodable information, for each of our three example networks. These values are limits inasmuch as they assume that every reaction event can be observed and recorded with infinite temporal precision, and that the encoding stochastic process is perfectly known. Qualitatively, the curves in Fig 3 show the same sigmoidal behavior: as the time horizon T is increased, the trajectories contain signatures of more and more reaction events that are input-specific, leading to a monotonic increase in the information, which must saturate at 1 bit (the total entropy of a two-state equally likely stimulus). The fact that we near saturation as the time horizon T is increased suggests that later time points add zero (conditional) mutual information. When faced with real data, this could be used as a criterion to determine the relevant response time-scale T , by testing when the transfer entropy (i.e., mutual information between the subsequent trajectory piece and the input conditioned on the prior segment of the trajectory of duration T) becomes zero [55, 56].

Examining the information increase specific to each example, we see 1 bit is reached more quickly in Example 2 than in Example 1 because of a larger difference in reaction rate parameters for both inputs in Example 2 relative to Example 1 in the period $T < 1000$. In contrast, in the period from $T > 1000$ until the $T = 2000$ at vertical yellow line, reaction rates for both inputs in Example 2 are the same, leading to essentially no accumulation of new information and a flat information curve, in contrast to Example 1. Finally, Example 3 shows that a model-based decoder can also make optimal use of the higher-order statistics of response trajectories even when the means under two conditions are the same, to extract the full bit of information before the $T = 2000$ cutoff.

While it is interesting to contemplate whether biological systems themselves could compute with or act on singular, precisely-timed reaction events and thus make optimal use of the resulting channel capacity (mimicking the debate between spike timing code and spike rate code in neuroscience), our primary focus here is to estimate information flows from experimental data. Typically, experiments record the state of the system—e.g., concentration of signaling molecules—in discretely sampled time. To explore the effects of time discretization, we first fix the observation length for our trajectories to $T = 2000$, sufficiently long that the trajectories in principle contain more than 90% of the theoretically maximal information for each of the three example cases. We then resample the trajectories on a grid of d equally spaced time points, as illustrated in Fig 4A.

Fig 4B–4D compare the exact Monte Carlo information approximation for discrete trajectories, $I_{\text{exact}}(\mathbf{X}; U)$, MAP lower bound for discrete trajectories, $I_{\text{MAP}}(\hat{U}; U)$, and the corresponding upper bound, I_{UB} , to the theoretical limits from Fig 3 obtained using continuous trajectories. In line with the chain of inequalities in Eq (25), information in discretely resampled trajectories is lower than the true information in continuous trajectories, but converges to the true value as $d \rightarrow \infty$. In particular, once the discretization timestep T/d is much lower than the inverse of the fastest reaction rate in the system, discretization should incur no

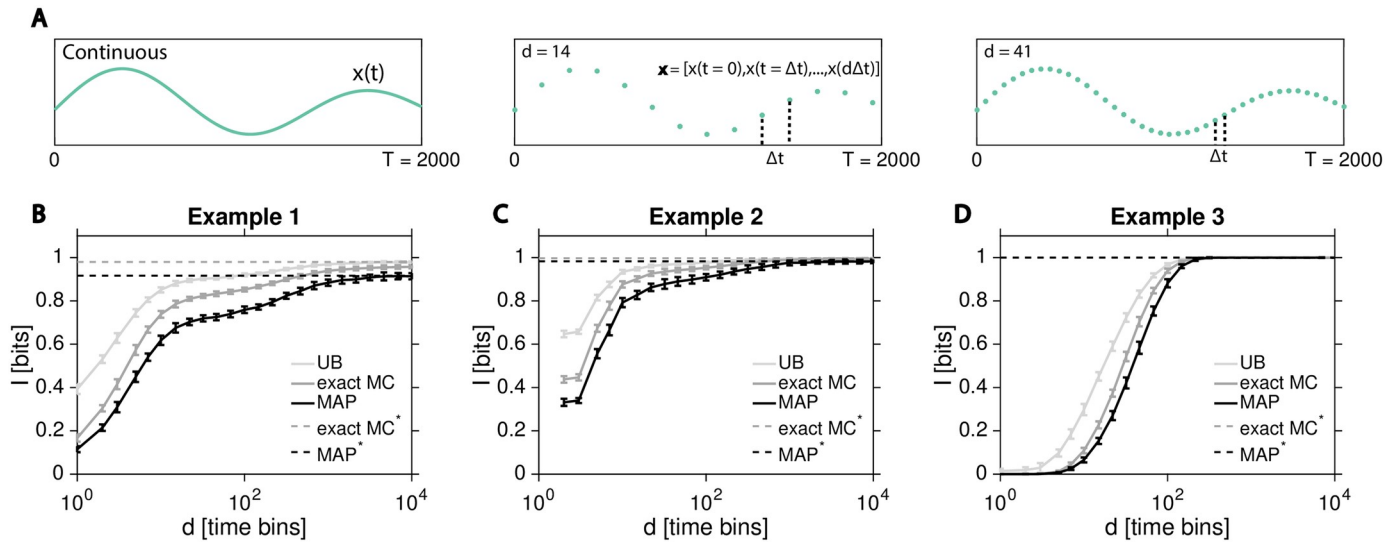


Fig 4. Information loss due to temporal sampling. (A) Schematic representation of the resampling of a continuous-time response trajectory (left) at $d = 14$ (middle) or $d = 41$ (right) equally spaced time points. Resampled response trajectories are represented as d -dimensional real vectors, $\mathbf{X} \in \mathbb{R}^d$, for the case of a single output chemical species. (B–D) Exact Monte Carlo information approximations for discrete trajectories, $I_{\text{exact}}(\mathbf{X}; U)$ (dark solid gray), optimal decoding lower bound, $I_{\text{MAP}}(\hat{U}; U)$ (dark solid black), and the upper bound, I_{UB} (light solid gray) are plotted as a function of d . Continuous-time limits from Fig 3 are shown as horizontal lines: $I_{\text{exact}}^*(X; U)$ (dashed dark gray), $I_{\text{MAP}}^*(\hat{U}; U)$ (dashed black). Error bars as in Fig 3.

<https://doi.org/10.1371/journal.pcbi.1007290.g004>

significant loss of information. In practice, however, high sampling rate (large d) limit has significant drawbacks: first, it is technically difficult to take snapshots of the system at such high rates (e.g., due to fluorophore bleaching); second, the fast dynamics of the reaction network may be low-pass filtered by the readout process (e.g., due to fluorophore maturation time, or slower downstream reaction kinetics); and third, for model-free approaches high d implies that decoders need to be learned over input spaces of high dimensionality, which could be infeasible given a limited number of experimentally recorded response trajectories. In previous work [25, 41, 57], trajectories were typically represented as $d \approx 1 \sim 100$ dimensional vectors, which in our examples would capture $\sim 80\%$ or more of the theoretically available information. It is likely that this can be improved further with smart positioning of the sampling points and that not all theoretically available information could actually be accessed by the organism itself, suggesting that typically used discretization approaches have the potential to capture most of the relevant information in the responses. What is important for the analysis at hand is that given the dimensionality d of the discretized response trajectories, MAP decoder is guaranteed to reach the minimal decoding error among all possible decoders, and will turn out to be a relevant benchmark, by yielding the highest information, $I_{\text{MAP}}(\hat{U}; U)$, in Fig 4B–4D among all decoders considered. In what follows, we will examine how various model-free decoding estimators approach this limit, as a function of d and the number of sample trajectories, N .

Performance of decoding-based estimators. After establishing our model-based “gold standard” for decoding-based estimators acting on trajectories represented in discretized time, $I_{\text{MAP}}(\hat{U}; U)$, we turn our attention to the performance comparison between various model-free algorithms. The results are summarized in Fig 5, which shows how estimator accuracy depends on the dimensionality of the problem, d , and the number N , of sample trajectories per input condition.

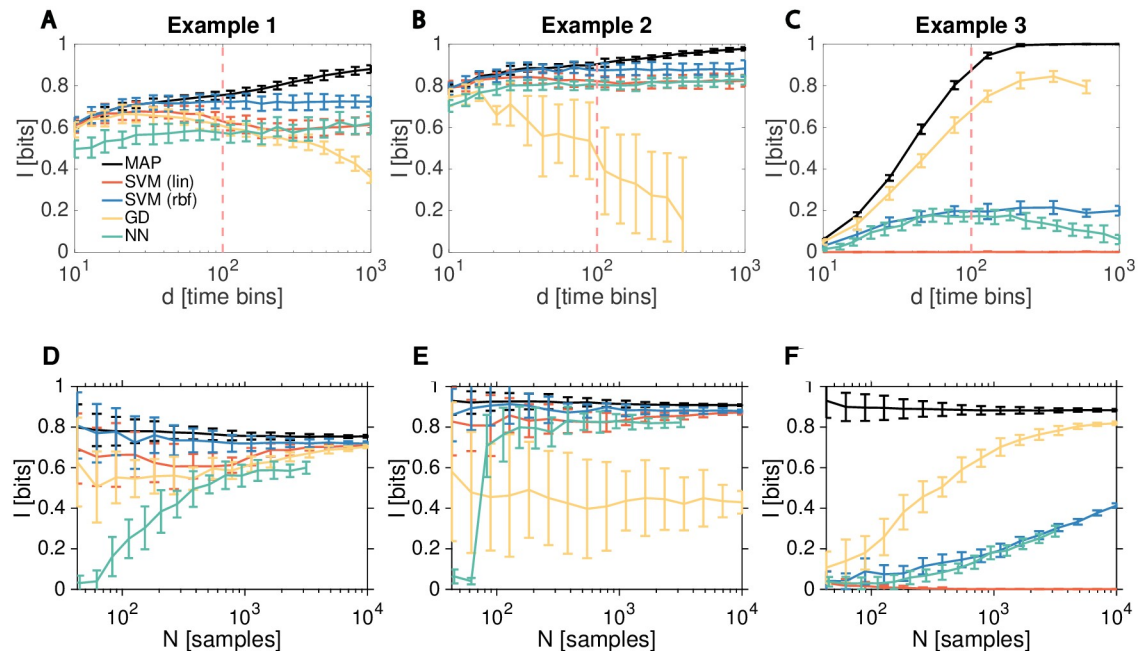


Fig 5. Performance of decoding-based estimators depends on the dimensionality of the response trajectories and on the number of response trajectory samples. Performance of various model-free decoding estimators (colored lines) for Examples 1 (A, D), 2 (B, E), 3 (C, F), respectively, compared to the MAP bound, I_{MAP} (black line), as a function of input trajectory dimension, d (at fixed $N = 1000$) in A, B, C; or as a function of the number of samples, N , per input condition (at fixed $d = 100$) in D, E, F. Error bars are std over 20 replicate estimations. Decoding estimators: linear SVM, $I_{SVM(lin)}$ (orange); radial basis functions SVM, $I_{SVM(rbf)}$ (blue); the Gaussian decoder with diagonal regularization (see S2 Fig for the effects of covariance matrix regularization and signal filtering on Gaussian decoder estimates), I_{GD} (yellow); multi-layer perceptron neural network, I_{NN} (green). Dashed vertical orange line marks the $d \leq 100$ regime typical of current experiments. Note that while the amount of information must in principle increase monotonically with d , the amount that decoders can actually extract given a limited number of samples, N , has no such guarantee. The decrease, at high d , in Gaussian decoder information estimate in A, B, C and neural network information estimate in C, happens because the number of parameters of the decoder grows with d albeit at fixed number of samples, leading to overfitting that regularization cannot fully compensate, and thus to the consequent loss of performance on the test data.

<https://doi.org/10.1371/journal.pcbi.1007290.g005>

Fig 5 leads us to the following conclusions:

- **Nonlinear SVM** using the radial basis functions (rbf) kernel performs best for Examples 1 and 2. Regardless of the number of samples, N , or the number of time bins, d , its estimates are very close to I_{MAP} , especially for the relevant regime $d \sim 10 - 100$. Even for higher d , the estimator shows hardly any overfitting and thus stable performance, a feature we have observed commonly in our numerical explorations. The estimator is sample efficient, typically providing estimates with smallest error bars.
- **Linear SVM** slightly underperforms kernelized SVM on Examples 1 and 2, and—as expected—completely fails on the linearly inseparable Example 3. Interestingly, even though more expressive, kernelized SVM seems to incur no generalization cost relative to linear SVM even at low number of samples. For all examples we tested, kernelized SVM thus appears to be a method of choice; linear SVM, however, is still useful as a benchmark to measure what fraction of the information is linearly decodable from the signal.
- **The Gaussian decoder** has the best performance on Example 3, is competitive for low d for Example 1, and does not perform satisfactory for Example 2. As shown in S2 Fig, regularized estimation of covariance matrix appears crucial for good performance, but smoothing of the

originally discrete trajectory does not help. Even with regularization, this estimator is not sample efficient for Example 1: the trajectories are linearly separable without a full estimate of the covariance (as evidenced by the success of the linear SVM), yet the Gaussian decoder requires one to two orders of magnitude more samples to match the linear decoder performance. This drawback turns into a benefit for Example 3: the Gaussian assumption can be viewed as a prior that second-order statistics are important for decoding (which is correct in this case). Kernelized SVM and the neural network, while more general, need to learn from many more training samples to zero in on these features, and fail to reach the Gaussian decoder performance even for $N = 10^4$. We hypothesized that the failure of the Gaussian decoder on Example 2 is due to the difficulty of the Gaussian approximation to capture the period $T > 1000$ when the mean number of \tilde{X} is close to zero: here, first, the Gaussian assumption must be strongly violated, and, second, the estimation of (co)variance from finite number of samples is close to singular due to the small number of reaction events in this period. Even though the $T > 1000$ epoch is not informative about the input, a badly conditioned decoder for this epoch can actually adversely affect performance. We confirmed this hypothesis by building the Gaussian decoder restricted to $T < 1000$ that reliably extracted ≥ 0.8 bits of information in Example 2, close to the MAP decoding bound and the performance of SVM-based estimators. We also examined the performance of the Gaussian decoder when the mean steady state number of molecules in Example 2 in $T < 1000$ period is increased from 10 to 20, 50, or 100 by scaling up the production rates, to see consistent increases in decoder performance (which can approach 1 bit, the MAP decoder limit, for $d \leq 100$ with the same parameters as shown in the main figure). At higher production rates, the signal-to-noise ratio is higher and trajectories in the $T < 1000$ period become more distinguishable, increasing the information; simultaneously, in the $T > 1000$ period, the mean number, although decaying, is also higher, making Gaussian approximation more applicable and preventing the covariance matrix from becoming poorly-conditioned. In sum, Gaussian decoder can provide competitive performance if the mean signal is sufficiently high so that the discreteness and positivity of molecular counts is not an issue, and when the covariance matrix is not close to ill-conditioned.

- Neural network decoder** reaches a comparable performance on Examples 2 and 3 to the SVMs, but fails to be competitive for the simple Example 1. This is most likely because this estimator is sample inefficient, as implied by its continual increase in performance with N that did not saturate at highest N shown in the figure; we confirmed further growth in performance of neural network decoder reaching (but not saturating) $I_{\text{NN}} \approx 0.65$ bits at $N = 10^5$. Given their expressive power, neural network decoders should be viewed as the opposite benchmark to the linear decoders: they have the ability to pick up complex statistical structures but only with a sufficient number of samples. Indeed, as we will see subsequently for applications to real data, neural networks can match and exceed the performance of SVMs. We emphasize that we used a neural network with a fixed architecture for all three examples on purpose, to make results comparable across examples; the performance can likely be improved by optimizing the architecture separately for each estimation problem. We did examine the issue of network architecture in greater detail in [S3 Fig](#), where we compared 14 architectures (1 one-layer, 12 two-layer, and 1 three-layer) on the difficult Example 3. We found that the simple one-layer network cannot extract any information; the three-layer network and several different two-layer networks can reach, but not exceed, the performance of the architecture used in the main figure given only $N = 1000$ samples, confirming our conclusion that higher performance requires significant increases in training set size independent of the neural network architecture.

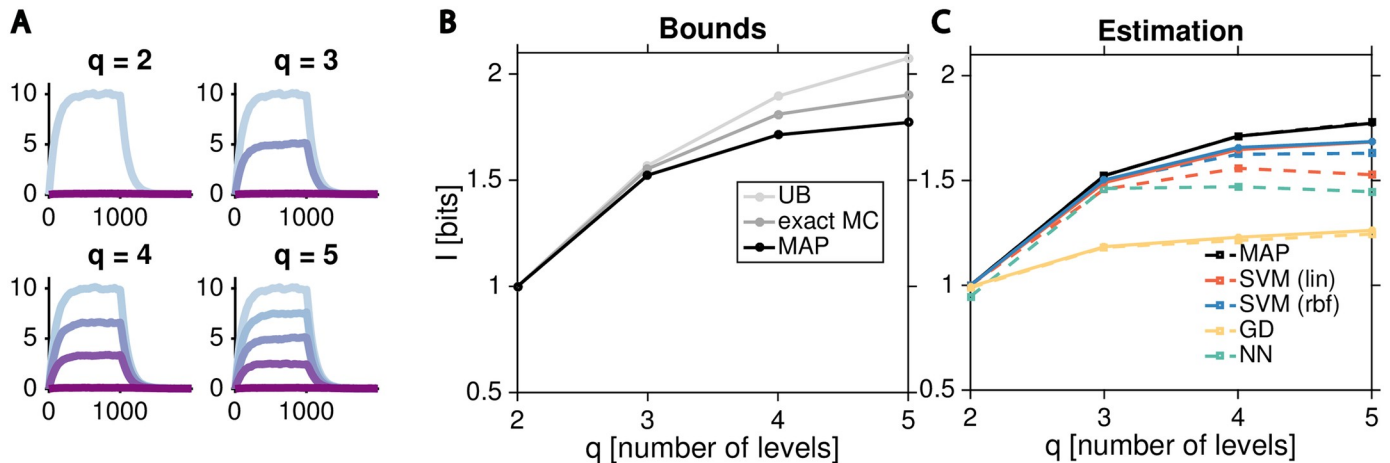


Fig 6. Information estimation for multilevel inputs. (A) Extension of Example 2 from Fig 2B to $q = 2, \dots, 5$ discrete inputs. We chose the inputs such that the response for the system at $T < 1000$ converges towards q equally spaced levels with the same dynamic range as the original example; dynamics at $T \geq 1000$ remain unchanged from the original Example 2. (B) Model-based information bounds as a function of the number of input levels for trajectories represented as $d = 100$ dimensional vectors: exact Monte Carlo calculation (dark gray), MAP decoding bound (black), upper bound (light gray). (C) Performance of model-free estimators, as indicated in the panel, compared to the MAP bound (black). Dashed lines show estimations using $N = 10^3$ sample trajectories per condition, solid lines using $N = 10^4$ samples per condition; in both cases, we show an average over 20 independent replicates, error bars are suppressed for readability.

<https://doi.org/10.1371/journal.pcbi.1007290.g006>

Multilevel information estimation. We next asked whether our conclusions hold also when the space of possible inputs is expanded beyond binary, assuming that U can take on q distinct values with equal probability, i.e., $p_U(u) = 1/q$. We focused on Example 2, and constructed cases for $q = 2, \dots, 5$ such that the production rate α for $0 < T < 1000$ takes on q uniformly spaced values between 0 and the maximal rate equal to $\alpha = 0.1$ used in Fig 2B. In effect, this “tiles” the original, two-state-input dynamic range uniformly with q input states, as illustrated in Fig 6A.

Our expectation is that with increasing q , the information should increase, but slowly saturate as reliable distinctions between nearby input levels can no longer be made due to the intrinsic biochemical stochasticity. This is indeed what we see in Fig 6B, which shows the exact information, the MAP lower bound and the upper bound. Consistent with our findings for two-input case, SVM using radial basis functions remains the estimator of choice for all values of q , followed by the linear SVM and then the neural network decoder, as shown in Fig 6C.

Performance comparisons with model-free information approximations. There exist many algorithms for estimating information directly, without making use of the decoding lower bound. The best known estimator for continuous signals is perhaps the k -nearest-neighbor (knn) estimator [39]. We have also introduced estimators based on parametric assumptions about the response distribution, such as the Gaussian approximation (Section *Model-free information estimators*); both belong in the family of binless approximations, which act directly on real-valued response vectors. In contrast, binning approximations first discretize the responses \mathbf{X} . The simplest such approach is perhaps the direct estimator of information or entropy [5], and a good review is provided in Ref [4]. We evaluated the performance of the Gaussian approximation to find that it can systematically overshoot the true information with a bias that is difficult to assess (S4 Fig); this appears to happen also in the regime where the biochemical noise should be small (relative to the mean), and the stochastic dynamics should be describable in terms of Langevin approximations with the resulting Gaussian response distributions. These approximations converge to the true solution in terms of their first and second

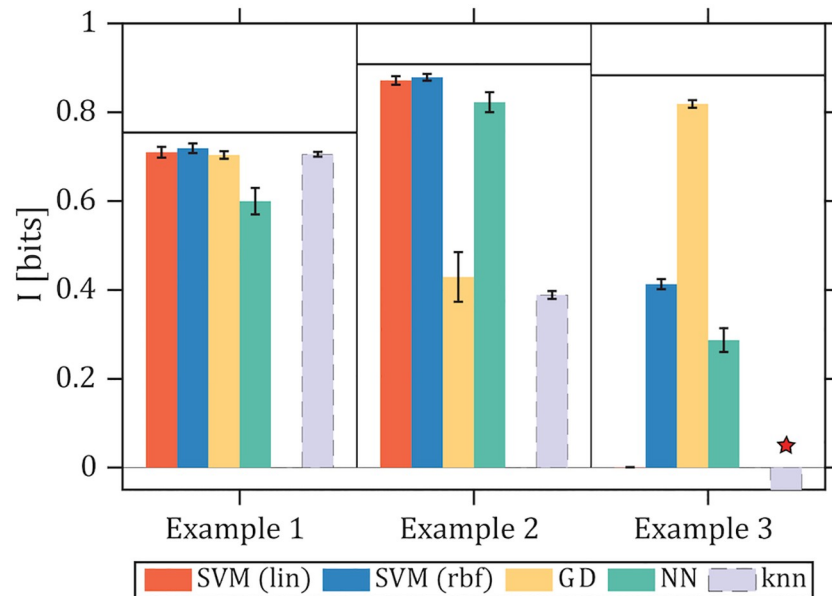


Fig 7. Comparison of decoding-based and knn information estimators. Information estimates for decoding-based (color bars) and knn (gray bar) algorithms (here we set $k = 1$, for further details on knn estimation, including varying k , see S5 Fig). Note that knn is not a decoding estimator and thus could exceed $I_{\text{MAP}}(\hat{U}; U)$ (shown as a horizontal black line for each of the three example cases) to approach the exact $I_{\text{exact}}(\mathbf{X}; U)$. Here we use trajectories discretized over $d = 100$ time bins, and $N = 10^4$ trajectory samples per input. The performance of knn can be substantially improved by adding a small amount of gaussian noise to the trajectory samples; its resulting performance as a function of N and d is shown in S5 Fig. Red star denotes the failure of knn on Example 3 where substantially negative information values are returned (exact value not plotted).

<https://doi.org/10.1371/journal.pcbi.1007290.g007>

moments, yet do not seem to lead to unbiased estimate for the entropies and thus the mutual information. In contrast to the Gaussian decoder, Gaussian approximation should not be used without a better understanding of its bias and applicability.

We therefore decided to focus on the comparison of decoding estimators with knn, which has been used previously on data from biochemical signaling networks [26]. The results are shown in Fig 7. K-nearest-neighbors performs well on the easy Example 1, and suffers drastic performance drop for Example 2, while crashing catastrophically by reporting negative values in Example 3. We reasoned that part of the difficulty may be the fact that synthetic trajectories for our Examples are defined over non-negative whole numbers only, whereas the knn assumes real valued vectors. This is confirmed by S5 Fig which shows that the knn performance can be substantially improved by adding a small amount of gaussian iid noise to every component of the response trajectory vectors, \mathbf{X} . This restores the knn performance in Example 2 close to that of the SVM-based estimators, but still produces close-to-zero bits of information for Example 3.

Applications to real data

To illustrate the use of our estimators in a realistic context, we analyzed data from two previously published papers. The first paper focused on the representation of environmental stress in the nuclear localization dynamics of several transcription factors (here we focus on data for Msn2, Dot6, and Sfp1) in budding yeast [27]. The second paper studied information transmission in biochemical signaling networks in mammalian cells (here we focus on data for ERK

and Ca^{2+}) [26]. In both cases, single-cell trajectory data were collected in hundreds or thousands of single cells sampled at sufficient resolution to represent the trajectories discretized at tens to hundreds of timepoints. Similarly, both papers estimate the information transmission in trajectories about a discrete number of environmental conditions: Ref [27] uses the linear SVM approach presented here, while Ref [26] uses the knn estimator. This makes the two datasets perfectly suited for estimator comparisons. We further note that in both datasets the trajectories can be divided into two response periods: the early “transient” response period when the external condition changes, and the late “near steady-state” response period. Typically, the transient dynamics exhibit clear differences in the trajectory means between various conditions, reminiscent of our Example 1 or early Example 2; in contrast, in the late period the response may have been adapted away, or the stimulus could be encoded only in higher-order statistics of the traces, reminiscent of the late period in Example 2 or Example 3.

Fig 8 shows the raw data and summarizes our estimation results for the early and late response periods for the three translocating factors in yeast that report on the change from 2% glucose rich medium to 0.1% glucose poor stress medium. Fig 9 similarly shows the raw data and estimation results for the early and late response periods for the signaling molecules in mammalian cells responding to multilevel inputs.

Consistent with the published report [27], transient response in yeast nuclear localization signal can be decoded well with the linear SVM estimator that yields about 0.6 bits of information per gene about the external condition. Kernelized SVM outperforms the linear method slightly by extracting an extra 0.1–0.2 bits of information, while knn underperforms the linear method significantly for *Msn2* and *Dot6* (but not for *Sfp1*). The Gaussian decoder estimate shows a mixed performance and the neural network estimate is the worst performer, most likely because the number of samples here is only $N = 100$ per input condition and neural network training is significantly impacted.

It is interesting to look at the stationary responses in yeast which have not previously been analyzed in detail. First, low estimates provided by linear SVM for *Msn2* and *Dot6* imply that information in the stationary regime, if present, cannot be extracted by the linear classifier. Second, the Gaussian decoder also performs poorly in the stationary regime, potentially indicating that the relevant features are encoded in higher-than-pairwise order statistics of the response (e.g., pulses could be “sparse” features as in sparse coding [58]); it is, however, hard to exclude small number of training samples as the explanation for the poor performance of the Gaussian decoder. Third, K-nearest-neighbor estimator also yields low estimates, either due to small sample number or low signal-to-noise ratio, the regime for which knn method has been observed to show reduced performance [40]. A particularly worrying feature of the knn estimates is their non-robust dependence on the length of the trajectory T . As S6 Fig shows, the performance of knn peaks at $T \approx 50$ min and then drops, even well into unrealistic negative estimates for $T \approx 400$ min (corresponding to the highest dimensionality $d = 170$ of discrete trajectories). While it is possible to make an *ad hoc* choice to always select trajectory duration at which the estimate peaks, the performance of kernelized SVM is, in comparison, extremely well behaved and increases monotonically with T , as theoretically expected. Finally, nonlinear SVM estimator extracts up to 0.4 bits of information about condition per gene, more than half of the information in the early transient period. This is even though on average the response trajectories for the two conditions, 2% glucose and 0.1% glucose, for *Msn2* and *Dot6* are nearly identical. For *Sfp1* there is a notable difference in the mean response, which the linear estimator can use to provide a ~ 0.15 bits of information, yet still significantly below ~ 0.4 bits extracted by the nonlinear SVM. For both transient and stationary responses in yeast, our results are qualitatively in line with the expectations from the synthetic example cases—given the small number of trajectories, tightest and most robust estimates are provided

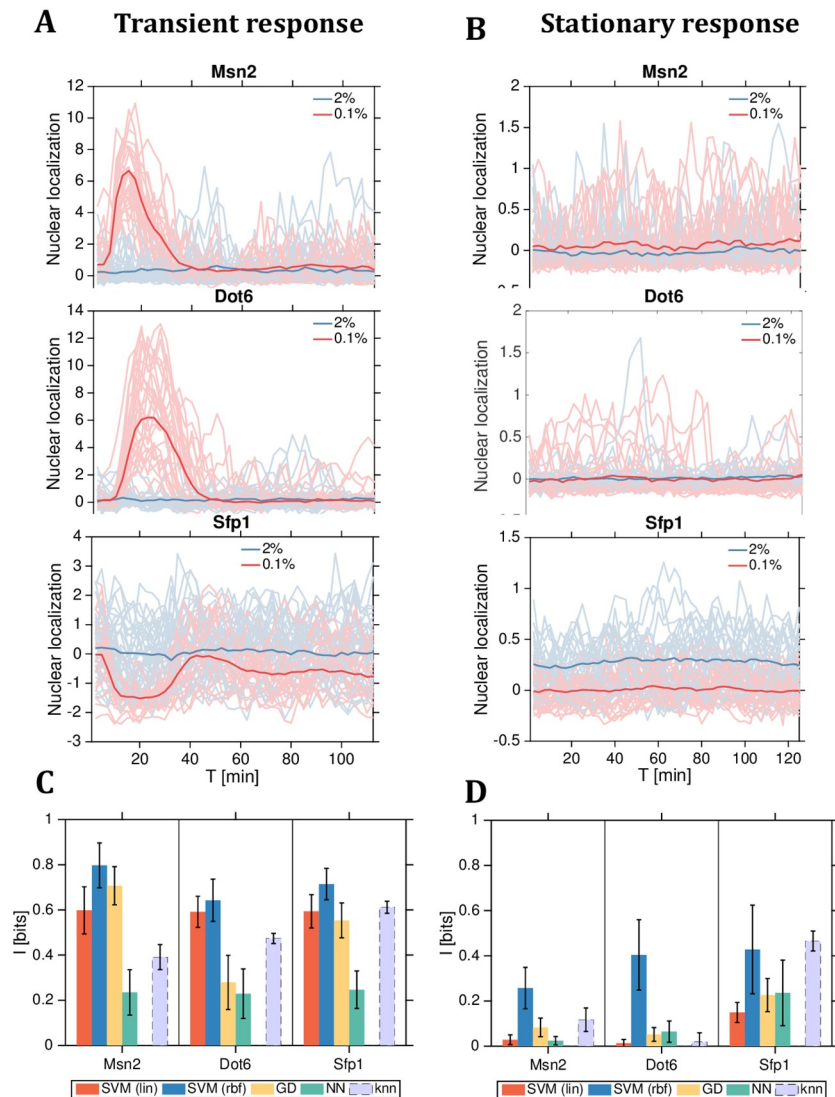


Fig 8. Two-level mutual information estimates from single-cell time-series data for nuclear translocation of yeast transcription factors. (A, B) Data replotted from Ref [27] for Msn2 (top row), Dot6 (middle row), and Sfp1 (bottom row); early transient responses (A) after nutrient shift at $t = 0$ min from glucose rich (2%, blue traces) to glucose poor (0.1%, red traces) medium are shown in the left column, stationary responses (B) are collected after cells are fully adapted to the new medium. Sampling frequency is 2.5 min, $d = 45$, and the number of sample trajectories per nutrient condition is $N = 100$. Thin lines are individual single cell traces, solid lines are population averages. (C, D) Information estimates for the transient (left, C) and stationary (right, D) response periods. Colored bars use model-free decoding-based estimators as indicated in the legend, gray bar is the knn estimate; error bars computed from estimation bootstraps by randomly splitting the data into testing and training sets.

<https://doi.org/10.1371/journal.pcbi.1007290.g008>

by the decoding information estimator based on nonlinear (kernelized) SVM. Regardless of the decoding methodology and even without small sample corrections at $N = 100$ trajectories per input, our estimates are not significantly impacted by the well-known information estimation biases thanks to the dimensionality reduction that decoding provides by mapping high dimensional trajectories X back into the space for inputs U which is low dimensional; this is verified in S7 Fig by estimating the (zero) information in trajectories whose input labels have been randomly assigned.

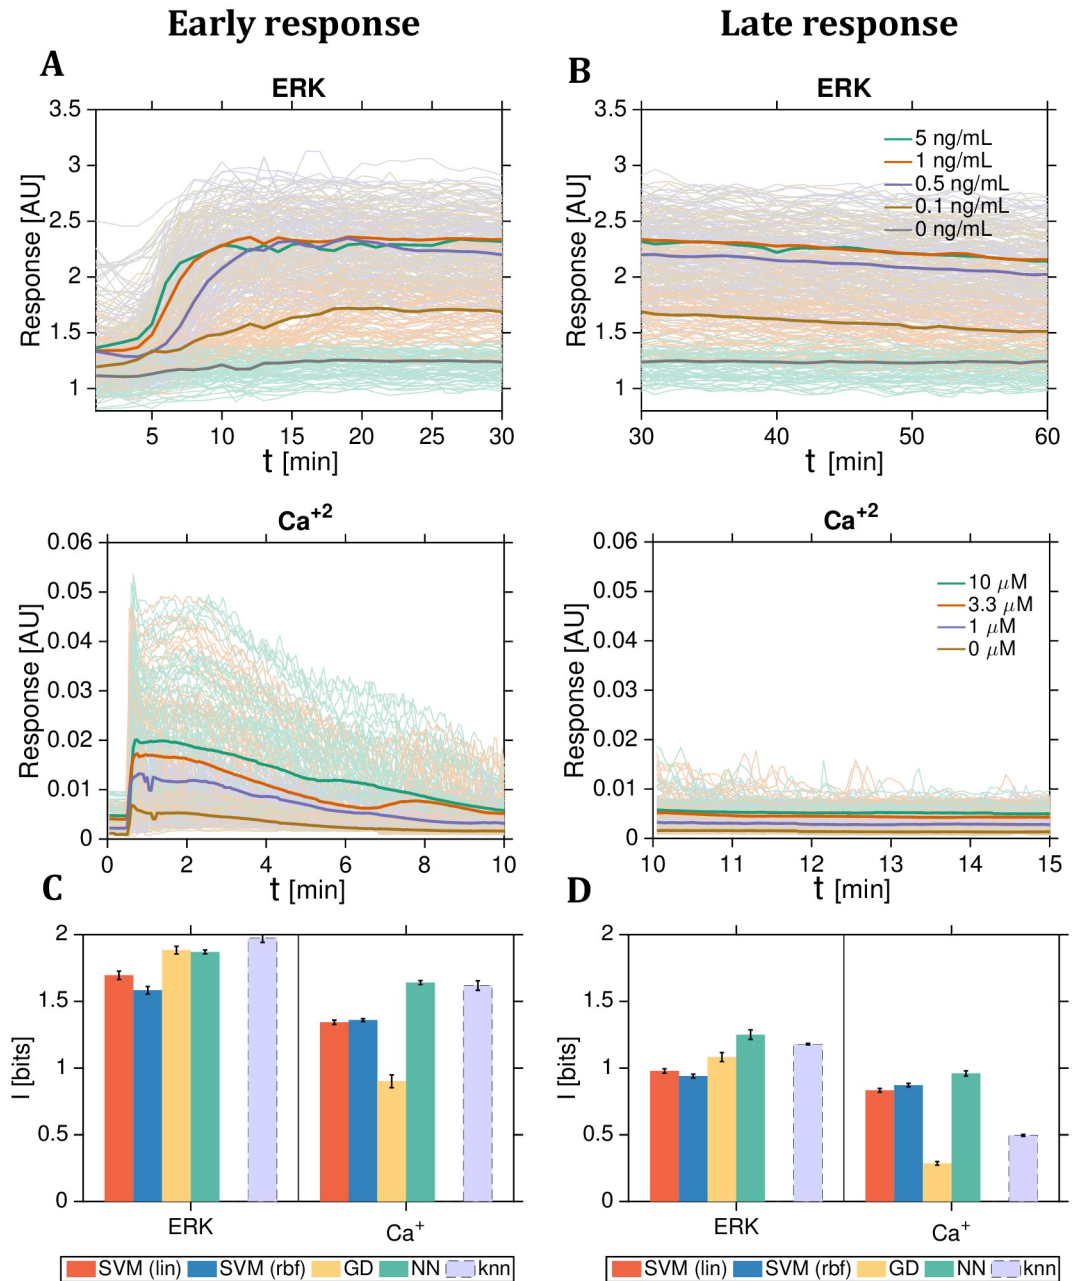


Fig 9. Multilevel mutual information estimates from single-cell time-series data for mammalian intracellular signaling. Data replotted from Ref [26] for ERK (top row) and Ca^{2+} (bottom row). (A) Early transient responses after addition of 5 different levels of EGF for ERK (or 4 different levels of ATP for Ca^{2+} , respectively) at $t = 0$ min, as indicated in the legend. (B) In the late response most, but not all, of the transients have decayed. (C,D) Information estimation using different methods (legend) in the early (C) and late (D) period, for ERK (left half of the panels) and Ca^{2+} (right half). Data for ERK: $N = 1678$ per condition, $T = 30$ min ($d = 30$) for early response and $T = 30$ min ($d = 30$) for late response. Data for Ca^{2+} : $N = 2995$ per condition, $T = 10$ min ($d = 200$) for early response and $T = 5$ min ($d = 100$) for late response. Plotting conventions as in Fig 8.

<https://doi.org/10.1371/journal.pcbi.1007290.g009>

Random pulses that encode stationary environmental signals have been observed for at least 10 transcription factors in yeast [53] and for tens of transcription factors in mammalian cells [59]. Recent studies investigated the role of the pulsatile dynamics in cellular decision-making [57, 60]. Nevertheless, methods for quantifying the information encoded in stochastic pulses are still in their infancy. Our nonlinear SVM decoding estimates convincingly show that there is information to be learned at the single cell level from the stationary stochastic pulsing. An interesting direction for future work is to ask whether hand-crafted features of the response trajectories (pulse frequency, amplitude, shape, etc) can extract as much information from the trajectories as the generic SVM classifier: for that, one would construct for each response trajectory a “feature vector” by hand, compute the linear SVM decoding bound information estimate from the feature vectors, and compare that to the kernelized SVM estimate over the original trajectories. This approach is a generic and operationally-defined path for finding “sufficient statistics” of the response trajectories—or a compression of the original signal to the relevant set of features—in the information-theoretic sense.

A different picture emerges from the mammalian signaling network data shown in Fig 9. The key difference here is the order of magnitude larger number of sample trajectories per condition compared to yeast data. Most of the information seems linearly separable in both the early and late response periods, as evidenced by the success of the linear SVM based estimator whose performance is not improved upon by the kernelized SVM (indeed, for early ERK response period linear SVM gives a slightly higher estimate than the nonlinear version). The big winner on this dataset is the neural-network-based estimator that yields the best performance in all conditions among the decoding-based estimators, likely owing to sufficient training data. As before, the Gaussian decoder shows mixed performance which can get competitive with the best estimators under some conditions. Lastly, knn appears to do well except on the late Ca^{2+} data (perhaps due to low signal-to-noise ratio). It also shows counter-intuitive non-monotonic behavior with trajectory duration T in S8 Fig (cf. with Fig 2C of Ref [27], where the analysis of information conveyed in dynamical signals as a function of trajectory duration was also very revealing about signaling in yeast). Once again it is worth keeping in mind that knn is estimating the full mutual information which could be higher than the information decodable from single responses.

Discussion

Increasing availability of single-cell time-resolved data should allow us to address open questions regarding the amount of information about the external world that is available in the time-varying concentrations, activation or localization patterns, and modification state of various biochemical molecules. Do full response trajectories provide more information than single temporal snapshots, as early studies suggest? Is this information gain purely due to noise averaging enabled by observing multiple snapshots, or—more interestingly—due to the ability of these intrinsically high-dimensional signals to provide a richer representation of the cellular environment? Can we isolate biologically relevant features of the response trajectories, e.g., amplitude, frequency, pulse shape, relative phase or timing, without *a priori* assuming what these features are? How can cells read out the environmental state from these response trajectories and how close to the information-theoretic bounds is this readout process? More broadly, a framework for analyzing complete response trajectories in signaling or genetic regulatory networks at the single cell level could lead to architectural and functional constraints on the biological network [27], and allow us to further pursue the ideas of optimal information representation in biological systems [9].

Here, we made methodological steps towards answering these questions by focusing on two related problems: first, if we are given a full stochastic description of a biochemical reaction network, under what conditions can we theoretically compute information transmission through this network and various related bounds; second, if we are given real data with no description of the network, what are tractable schemes to estimate the information transmission. We show that when the complete state of the reaction network is observed and the inputs are discrete sets of reaction rates, there exist tractable Monte Carlo approximation schemes for the information transmission. These exact results that we compute for three simple biological network examples then serve to benchmark a family of decoding-based model-free estimators and compare their performance to the commonly-used knn estimator. We show that decoding-based estimators can closely approach the optimal decoder performance and in many cases perform better than knn, especially with typical problem dimensions ($d \sim 1 - 100$) and typical number of sample trajectories ($N \sim 10^2 - 10^3$). This is especially true when we ask about the combinatorial representation of the environmental state in the time trajectories of several jointly observed chemical species, as in our previous work [27], where alternative information estimation methods usually completely fail due to the high dimensionality of the input space.

It is necessary to emphasize the flexibility of the decoding approach: decoding-based information estimation is based directly on the statistical problems of classification (for discrete input variable, U) or regression (for continuous input variable, U), so any classification / regression algorithm with good performance can provide the basis for information estimation. Concretely, for problems in the low data regime (small N), linear or kernelized SVM approaches appear powerful, while at larger N neural-network-based schemes can provide a better performance and thus typically a tighter information lower bound. In contrast to information approximations for which it is often impossible to assess their precision or bias (or even its sign) when the dimension, d , of the problem is large, the decoding approach yields a conservative estimate of the true information. Statistical algorithms underlying decoding-based estimations have the extra advantage that, (i), we may be able to gain biological insight by inspecting which features of the response carry the relevant stimulus information (e.g., by looking at the linear kernels or features that neural networks extract in their various layers); (ii), pick a decoding algorithm based on features previously reported as relevant (e.g., the Gaussian decoder for second-order statistics as in Example 3); (iii), estimate the information as a function of trajectory duration; and (iv), gain confidence in our estimates by testing their performance on withheld data. While we tested these estimators on a very restricted set of toy examples in order to be able to compare to analytically computed results, model-free decoding-based approaches are applicable more generally, e.g., to complex, partially-observed reaction systems, or networks with significant contribution of cell-to-cell variability or extrinsic noise.

By construction, decoding-based estimators only provide a lower bound to the true information. This, however, could turn out to be a smaller problem in practice than it appears in theory, especially for biochemical reaction networks. First, our extension to the Feder-Merhav bound provides us with an estimate of how large the gap between the true information and the decoded estimation can be. The bound is not tight on our examples, and can only be applied when the optimal MAP decoder can be constructed [61, 62]. Second, and perhaps more importantly, information that can be decoded after single input presentations is the quantity that is likely more biologically relevant than the true channel capacity, if the organisms are under constraint to respond to the environmental changes quickly. Typically, organisms across the complexity scale operate under speed-accuracy tradeoffs [63]: faster decisions based on noisy information lead to more errors and, conversely, with enough time to integrate sensory information errors can be reduced. When speed is at a premium or relevant inputs are sparse,

decisions need to be taken after single input presentations. In this case, decoding-based estimation should not be viewed as an approximate but rather as the correct methodology for the biological problem at hand. Of course, there is still the question of whether the model-free decoders that we use on real data can achieve a performance that is close to the optimal MAP decoder that represents the absolute performance limit. While there is no general way to answer this question, it appears that simple SVM decoding schemes work well when the response trajectories differ in their conditional mean, and neural networks as general approximators can be used to check for more complicated encoding features when data is plentiful. Unlike in neuroscience, there is much less clarity about what kind of read-out or decoding operations biochemical networks can mechanistically realize to mimic the functioning of our *in silico* decoders, and it may be challenging to biochemically implement even arbitrary linear classification of response trajectories. Until experimentally shown otherwise, it thus appears reasonable to proceed with the assumption that environmental signals can be read out from the time-dependent internal chemical state with a simple repertoire of computations.

We also mention a caveat when using decoding-based estimators that rely on classification or regression methods with large expressive power, such as neural networks. While it is possible to successfully guard against overfitting within the same dataset using cross-validation, scientific insights into biological function often require generalization beyond one particular dataset. Typically, we ask for generalization at least over independent experimental replicates, but sometimes even over similar (but not same) external conditions, strains, or experimental setups. This can present a serious issue if e.g., neural networks overfit to such systematic variations between replicates or conditions even when such variations are not biologically relevant. Regularization alone will not necessarily guard against this, unless the networks are actually trained over a subset of all data on which they will be tested. A pertinent recommendation here is to evaluate the difference in performance of expressive decoding-based estimators when trained over a subset or over all replicates, and to compare that to the generalization of less-expressive methods for which the sufficient statistics are known (e.g., linear or Gaussian decoders).

We conclude by emphasizing a simple yet important point. The decoding-based approach that we introduced here should also motivate us to look beyond methodological problems of significance and estimation, to truly biological problems of cellular decision making. Currently, data on biological regulatory processes is often analyzed by looking for “statistically significant differences” in the network response for, say, two possible network inputs. For example, one may report that the steady-state mean expression level of a certain gene is significantly larger in the stimulated vs unstimulated condition, with the statistical significance of the mean difference established through an appropriate statistical test that takes into account the number of collected population samples. While statistical significance is a necessary condition to validly report *any* difference in the response, it is very different from the question of whether *a single cell* could discriminate the two conditions given access only to its own expression levels. In caricature, population-level statistics tell us with what confidence we, as scientists having access to N samples, can discriminate between conditions given some biological readout; decoding based information estimates, on the other hand, are relevant to the $N = 1$ case of individual cells. We hope that further work along the latter path can clarify and quantify better the difficult constraints and conditions under which real cells need to act based on individual noisy readouts of their stochastic biochemistry.

Supporting information

S1 Fig. Decoding-based estimators applied to Example 1 trajectories with extrinsic noise.

(A) Sample trajectories as in Example 1 of Fig 2A, using the same parameters and plotting

conventions. Extrinsic noise was introduced by perturbing the degradation rate (fixed at 0.01 for all cells in the main paper), by adding a Gaussian IID distributed random variable with $\sigma = 0.001$ (shown in A and B) or $\sigma = 0.0002$ (shown in C); the random variable is drawn separately for each cell at the beginning of the simulation and is held fixed through time. (B,C) Estimator performance on test data (where extrinsic noise is also resampled for each cell in the test set) for both extrinsic noise levels. SVM and knn estimators perform best, but unlike in the main paper, here we do not have a reference comparison of the MAP decoder. While decoding-based estimators are giving a conservative lower-bound, we have no guarantees of whether knn extracts more information (and thus has better performance), or actually overestimates decodable information.

(TIF)

S2 Fig. Effects of covariance matrix regularization and signal smoothing on Gaussian-decoder-based estimation. (A) At left. Diagonal covariance regularization following Ref 64 of the main paper. Briefly, λ times the identity matrix is added to the empirical covariance matrix with the hyperparameter λ set so that the likelihood on test data is maximized. Shown is the empirical (left) and regularized (right) covariance matrix for Example 3, using $d = 20$ and $N = 30$ sample trajectories. At right. Information estimates for Example 3: I_{MAP} decoding bound (black), Gaussian decoder estimate, $I_{\text{GD}(\text{reg})}$, with optimal diagonal regularization for each d (yellow, as in Fig 5C), Gaussian decoder estimate, $I_{\text{GD}(\text{no reg})}$ (brown). Without regularization, the estimate suffers an abrupt drop as d increases and the empirically estimated covariance matrix becomes close to singular. N and plotting conventions are as in Fig 5. (B) The effects of trajectory filtering on information estimates. At left. A raw integer-valued stochastic trajectory for \tilde{X} (blue) can be filtered by a low-pass exponential decay filter with adjustable timescale, $\tau = 1 - 10^3$, here $\tau = 50$ (red) to yield real-valued trajectory. At right. Regularized Gaussian-decoder information estimates with (brown) and without (yellow) filtering. Filtering does not improve but can decrease the estimation performance, even when the filtering timescale is adjusted.

(TIF)

S3 Fig. The effect of neural network architecture on the performance of information estimators for Example 3. Shown are I_{NN} estimates on Example 3, analogous to Fig 5C and 5F of the main paper. Decoders are trained on $N = 1000$ sample trajectories per input condition, and trajectories are represented as $d = 50$ dimensional vectors. The performance of the network architecture used in the main paper is shown as the horizontal red line. Alternative architectures are denoted on the x-axis label. “1L, 2L, 3L” stands for 1, 2 or 3 layers, respectively. The number in front of letter “N” represents the total number of neurons on the hidden layers. Other technical details of the networks as reported in the Methods.

(TIF)

S4 Fig. Gaussian approximation to the information can lead to an uncontrolled overestimation of the true information. Gaussian approximation is evaluated for Example 3 in Fig 5C, using $N = 1000$ per condition. Exact Monte Carlo approximation of the information, $I_{\text{exact}}(\mathbf{X}; U)$, is shown in dark gray. Information estimates following Section *Model-free information estimators* are shown in violet (Gaussian approximation for raw, integer-valued response trajectories) or in cyan (Gaussian approximation for filtered trajectories), as in S2 Fig. In both cases the Gaussian approximation overshoots the true information value. Further numerical analyses indicated that the difference is hard to predict and that it persists even when the reaction rates are chosen such that the mean expression level is ten-fold higher (and the intrinsic stochasticity correspondingly lower). This makes direct Gaussian approximation

risky to use, in contrast to the Gaussian-decoder based estimate, which is guaranteed to stay below I_{exact} .

(TIF)

S5 Fig. Behavior of the knn information estimator. Compared to knn results in Fig 7, the results in A, B and D are estimated following the same procedure, while adding a small amount of IID zero-mean Gaussian noise to each response trajectory at every time bin; the noise variance must be $\ll 1$ but otherwise does not affect the results much. This results in good estimates even at low sample number, N , and provides nearly stable estimation as a function of the trajectory dimension, d , for Example 1 and Example 2. It, however, does not resolve the estimator failure for Example 3. (A) Dependence of the knn estimator performance on the number of samples. Yellow plot symbols indicate the number of samples per condition, $N = 10^3$, used in Fig 7. (B) Dependence of the knn estimator performance on the trajectory dimension. Yellow plot symbols indicate the dimension, $d = 10^2$, used in Fig 7. (C, D) Dependence of the knn estimates on the number of nearest neighbors, k , at $N = 10^3$ and $d = 10^2$, without the addition of noise (C) or with the addition of noise (D).

(TIF)

S6 Fig. Estimator behavior for longer trajectory data for Dot6. When the samples are limited, here to $N = 100$ samples per input glucose level condition as in Fig 8A (middle), radial-basis-function SVM estimate (blue) is well-behaved with no observable overfitting and consequent drop in information estimate as the trajectory duration, T , is increased (maximal T corresponds to $d = 170$ dimensional trajectory vectors). In contrast, knn estimate (brown) shows a collapse in the estimation performance, even yielding strongly negative numbers, as the dimensionality of input vectors is increased at fixed number of trajectory samples.

(TIF)

S7 Fig. Assessing information estimation bias due to small sample size. By randomly shuffling the binary labels assigned to different response trajectories, we break all response-input correlations leading to zero information. Here we test whether our estimators correctly report zero information within error bars given a finite number of samples, or are subject to positive information estimation bias. Decoding-based estimates (linear SVM, red; kernelized SVM, blue; Gaussian decoder, yellow) and knn (gray). First three sets of bars correspond to synthetic examples of Fig 3; estimations are done with $d = 100$ and $N = 1000$ per input condition as in Fig 5, following the same plotting conventions. Last two sets of bars are estimated with $N = 100$ per input condition using real data for Sfp1 yeast TF from Fig 8A. In all cases, even without explicit small-sample debiasing for Eq (26) (which may be required for multilevel estimation), the estimates are consistent with zero.

(TIF)

S8 Fig. Information estimates for mammalian signaling networks as a function of the trajectory duration. Shown are information estimates as a function of the total trajectory duration, T , for the early response period for ERK (A) and Ca^{2+} (B). Plotting conventions, procedures, and data set sizes same as in Fig 9.

(TIF)

Acknowledgments

We thank Alejandro Granados, Mihal Hledik, Julian Pietsch, and Christoph Zechner for stimulating discussions.

Author Contributions

Conceptualization: Jakob Ruess, Gašper Tkačik.

Formal analysis: Sarah Anhala Cepeda-Humerez, Gašper Tkačik.

Funding acquisition: Gašper Tkačik.

Investigation: Sarah Anhala Cepeda-Humerez, Jakob Ruess, Gašper Tkačik.

Methodology: Sarah Anhala Cepeda-Humerez, Jakob Ruess, Gašper Tkačik.

Software: Sarah Anhala Cepeda-Humerez.

Supervision: Gašper Tkačik.

Visualization: Sarah Anhala Cepeda-Humerez.

Writing – original draft: Sarah Anhala Cepeda-Humerez, Jakob Ruess, Gašper Tkačik.

Writing – review & editing: Sarah Anhala Cepeda-Humerez, Jakob Ruess, Gašper Tkačik.

References

1. Eldar A, Elowitz MB. Functional roles for noise in genetic circuits. *Nature*. 2010; 467(7312):167–173. <https://doi.org/10.1038/nature09326> PMID: 20829787
2. Purvis JE, Lahav G. Encoding and Decoding Cellular Information through Signaling Dynamics. *Cell*. 2013; 152(5):945–956. <https://doi.org/10.1016/j.cell.2013.02.005> PMID: 23452846
3. Shannon CE, Weaver W. *The Mathematical Theory of Communication*. vol. 27. Urbana: University of Illinois Press; 1949.
4. Paninski L. Estimation of Entropy and Mutual Information. *Neural Computation*. 2003; 15:1191–1253. <https://doi.org/10.1162/089976603321780272>
5. Strong S, Koberle R, de Ruyter van Steveninck R, Bialek W. Entropy and Information in Neural Spike Trains. *Physical Review Letters*. 1998; 80:197–200. <https://doi.org/10.1103/PhysRevLett.80.197>
6. Quiroga RQ, Panzeri S. Extracting information from neuronal populations: information theory and decoding approaches. *Nature Reviews Neuroscience*. 2009; 10(3):173. <https://doi.org/10.1038/nrn2578>
7. Bowsher CG, Swain PS. Environmental sensing, information transfer, and cellular decision-making. *Current Opinion in Biotechnology*. 2014; 28:149–155. <https://doi.org/10.1016/j.copbio.2014.04.010> PMID: 24846821
8. Bialek W. *Biophysics: Searching for Principles*. Princeton University Press; 2012.
9. Tkačik G, Bialek W. Information Processing in Living Systems. *Annual Review of Condensed Matter Physics*. 2016; 7(1):89–117. <https://doi.org/10.1146/annurev-conmatphys-031214-014803>
10. Tkacik G, Aleksandra W. Information transmission in genetic regulatory networks: a review. *Journal of Physics: Condensed Matter*. 2011; 23. PMID: 21460423
11. Thomas PJ, Eckford AW. Capacity of a Simple Intercellular Signal Transduction Channel. *IEEE Transactions on Information Theory*. 2016; 62(12):7358–7382. <https://doi.org/10.1109/TIT.2016.2599178>
12. Tostevin F, ten Wolde PR. Mutual Information between Input and Output Trajectories of Biochemical Networks. *Phys Rev Lett*. 2009; 102:218101. <https://doi.org/10.1103/PhysRevLett.102.218101> PMID: 19519137
13. Tkačik G, Callan CG, Bialek W. Information capacity of genetic regulatory elements. *Phys Rev E*. 2008; 78:011910. <https://doi.org/10.1103/PhysRevE.78.011910>
14. Sokolowski TR, Tkačik G. Optimizing information flow in small genetic networks. IV. Spatial coupling. *Physical Review E*. 2015; 91(6):062710. <https://doi.org/10.1103/PhysRevE.91.062710>
15. Sokolowski TR, Walczak AM, Bialek W, Tkačik G. Extending the dynamic range of transcription factor action by translational regulation. *Physical Review E*. 2016; 93(2):022404. <https://doi.org/10.1103/PhysRevE.93.022404> PMID: 26986359
16. Tkacik G, Walczak AM, Bialek W. Optimizing information flow in small genetic networks. III. A self-interacting gene. *Phys Rev E*. 2012; 85(4):041903. <https://doi.org/10.1103/PhysRevE.85.041903>
17. Walczak AM, Tkačik G, Bialek W. Optimizing information flow in small genetic networks. II. Feed-forward interactions. *Physical Review E*. 2010; 81(4):041905. <https://doi.org/10.1103/PhysRevE.81.041905>

18. Tkačik G, Walczak AM, Bialek W. Optimizing information flow in small genetic networks. *Physical Review E*. 2009; 80(3):031920. <https://doi.org/10.1103/PhysRevE.80.031920>
19. Rieckh G, Tkačik G. Noise and Information Transmission in Promoters with Multiple Internal States. *Biophysical Journal*. 2014; 106(5):1194–1204. <https://doi.org/10.1016/j.bpj.2014.01.014> PMID: 24606943
20. Cheong R, Rhee A, Wang CJ, Nemenman I, Levchenko A. Information Transduction Capacity of Noisy Biochemical Signaling Networks. *Science*. 2011; 334(6054):354–358. <https://doi.org/10.1126/science.1204553> PMID: 21921160
21. Tostevin F, Ten Wolde PR. Mutual information in time-varying biochemical systems. *Physical Review E*. 2010; 81(6):061917. <https://doi.org/10.1103/PhysRevE.81.061917>
22. de Ronde W, Tostevin F, ten Wolde PR. Multiplexing Biochemical Signals. *Phys Rev Lett*. 2011; 107:048101. <https://doi.org/10.1103/PhysRevLett.107.048101> PMID: 21867046
23. Dubuis JO, Tkačik G, Wieschaus EF, Gregor T, Bialek W. Positional information, in bits. *Proceedings of the National Academy of Sciences*. 2013; 110(41):16301–16308. <https://doi.org/10.1073/pnas.1315642110>
24. Voliotis M, Perrett RM, McWilliams C, McArdle CA, Bowsher CG. Information transfer by leaky, heterogeneous, protein kinase signaling systems. *Proceedings of the National Academy of Sciences*. 2014; 111(3):E326–E333. <https://doi.org/10.1073/pnas.1314446111>
25. Hansen AS, Shea EKO. Limits on information transduction through amplitude and frequency regulation of transcription factor activity. *eLife*. 2015; p. 1–19.
26. Selimkhanov J, Taylor B, Yao J, Pilko A, Albeck J, Hoffmann A, et al. Accurate information transmission through dynamic biochemical signaling networks. *Science (New York, NY)*. 2014; 346(6215):1370–3. <https://doi.org/10.1126/science.1254933>
27. Granados AA, Pietsch JMJ, Cepeda-Humerez SA, Farquhar IL, Tkačik G, Swain PS. Distributed and dynamic intracellular organization of extracellular information. *Proceedings of the National Academy of Sciences of the United States of America*. 2018; 115(23):6088–6093. <https://doi.org/10.1073/pnas.1716659115> PMID: 29784812
28. Borst A, Theunissen FE. Information theory and neural coding. *Nature neuroscience*. 1999; 2:947–957. <https://doi.org/10.1038/14731> PMID: 10526332
29. Marre O, Botella-Soler V, Simmons KD, Mora T, Tkačik G, Berry MJ II. High accuracy decoding of dynamical motion from a large retinal population. *PLoS computational biology*. 2015; 11(7):e1004304. <https://doi.org/10.1371/journal.pcbi.1004304> PMID: 26132103
30. Rieke F, Warland D, Bialek W. Coding Efficiency and Information Rates in Sensory Neurons. *EPL (Europhysics Letters)*. 1993; 22(2):151.
31. Gillespie DT. A rigorous derivation of the chemical master equation. *Physica A: Statistical Mechanics and its Applications*. 1992; 188(1-3):404–425. [https://doi.org/10.1016/0378-4371\(92\)90283-V](https://doi.org/10.1016/0378-4371(92)90283-V)
32. Van Kampen NG. *Stochastic Processes in Physics and Chemistry*; 2007.
33. Gillespie DT. Exact stochastic simulation of coupled chemical reactions. *The Journal of Physical Chemistry*. 1977; 81(25):2340–2361. <https://doi.org/10.1021/j100540a008>
34. Feder M, Merhav N. Relations between entropy and error probability. *IEEE Transactions on Information Theory*. 1994; 40(1):259–266. <https://doi.org/10.1109/18.272494>
35. Hledik M, Sokolowski T, Tkačik G. A tight upper bound on mutual information. *arXiv.org*. 2018; p. 1812.01475.
36. Dobrushin RL. A simplified method of experimental estimation of the entropy of a stationary distribution. *Tear Veroyatnost i Primenen; English transl Theory Probab Appl*. 1958; 3:462–464.
37. Vasicek OA. A note on using cross-sectional information in bayesian estimation of security betas. *The Journal of Finance*. 1973; 28(5):1233–1239. <https://doi.org/10.1111/j.1540-6261.1973.tb01452.x>
38. Kaiser A, Schreiber T. Information transfer in continuous processes. *Physica D: Nonlinear Phenomena*. 2002; 166(1-2):43–62. [https://doi.org/10.1016/S0167-2789\(02\)00432-3](https://doi.org/10.1016/S0167-2789(02)00432-3)
39. Kraskov A, Stögbauer H, Grassberger P. Estimating mutual information. *Physical Review E—Statistical, Nonlinear, and Soft Matter Physics*. 2004; 69.
40. Khan S, Bandyopadhyay S, Ganguly AR, Saigal S, Erickson DJ, Protopopescu V, et al. Relative performance of mutual information estimation methods for quantifying the dependence among short and noisy data. *Physical Review E*. 2007; 76(2):026209. <https://doi.org/10.1103/PhysRevE.76.026209>
41. Potter GD, Byrd TA, Mugler A, Sun B. Dynamic Sampling and Information Encoding in Biochemical Networks. *Biophysical Journal*. 2017; 112(4):795–804. <https://doi.org/10.1016/j.bpj.2016.12.045> PMID: 28256238
42. Anderson TW. *An introduction to multivariate statistical analysis*. vol. 2. Wiley New York; 1958.

43. Cover TM, Thomas JA. Elements of Information Theory. John Wiley and Sons; 2005.
44. Brunel N, Nadal JP. Mutual information, Fisher information, and population coding. *Neural computation*. 1998; 10(7). <https://doi.org/10.1162/089976698300017115> PMID: 9744895
45. Murphy KP. Machine Learning: A Probabilistic Perspective. The MIT Press; 2012.
46. Samengo I. Information loss in an optimal maximum likelihood decoding. *Neural Computation*. 2002; 14(4):771–779. <https://doi.org/10.1162/089976602317318947> PMID: 11936960
47. Benabdeslem Khalid and Bennani Y. Dendrogram-based SVM for Multi-Class Classification. *Journal of Computing and Information Technology*. 2006; 14(4):283–289. <https://doi.org/10.2498/cit.2006.04.03>
48. Lajnef T, Chaibi S, Ruby P, Aguera PE, Eichenlaub JB, Samet M, et al. Learning machines and sleeping brains: Automatic sleep stage classification using decision-tree multi-class support vector machines. *Journal of Neuroscience Methods*. 2015; 250:94–105. <https://doi.org/10.1016/j.jneumeth.2015.01.022> PMID: 25629798
49. McCulloch WS, Pitts W. A logical calculus of the ideas immanent in nervous activity. *The Bulletin of Mathematical Biophysics*. 1943; 5(4):115–133. <https://doi.org/10.1007/BF02478259>
50. Rosenblatt F. The Perceptron—A Perceiving and Recognizing Automaton. Cornell Aeronautical Laboratory; 1957.
51. Géron A. Hands-On Machine Learning with Scikit-Learn and TensorFlow: Concepts, Tools, and Techniques to Build Intelligent Systems. O'Reilly; 2017.
52. Abadi M, Agarwal A, Barham P, Brevdo E, Chen Z, Citro C, et al. TensorFlow: Large-Scale Machine Learning on Heterogeneous Systems; 2015.
53. Dalal CK, Cai L, Lin Y, Rahbar K, Elowitz MB. Pulsatile Dynamics in the Yeast Proteome. *Current Biology*. 2014; 24(18):2189–2194. <https://doi.org/10.1016/j.cub.2014.07.076> PMID: 25220054
54. Taniguchi Y, Choi PJ, Gene-Wei L, Chen H, Babu M, Hearn J, et al. Quantifying E. coli proteome and transcriptome with single-molecule sensitivity in single cells. *Science*. 2010; 30:533. <https://doi.org/10.1126/science.1188308>
55. Schreiber T. Measuring information transfer. *Phys Rev Lett*. 2000; 85:461. <https://doi.org/10.1103/PhysRevLett.85.461> PMID: 10991308
56. Lindner M, Vicente R, Priesemann V, Wibral M. TRENTOOL: A Matlab open source toolbox to analyse information flow in time series data with transfer entropy. *BMC Neuroscience*. 2011; 12:119. <https://doi.org/10.1186/1471-2202-12-119> PMID: 22098775
57. Hafner A, Stewart-Ornstein J, Purvis JE, Forrester WC, Bulyk ML, Lahav G. p53 pulses lead to distinct patterns of gene expression albeit similar DNA-binding dynamics. *Nature Structural & Molecular Biology*. 2017; 24(10):840–847. <https://doi.org/10.1038/nsmb.3452>
58. Olshausen BA, Field DJ. Sparse coding of sensory inputs. *Current Opinion in Neurobiology*. 2004; 14(4):481–487. <https://doi.org/10.1016/j.conb.2004.07.007> PMID: 15321069
59. Levine JH, Lin Y, Elowitz MB. Functional Roles of Pulsing in Genetic Circuits. *Science (New York, NY)*. 2013; 342(December):1193–1200. <https://doi.org/10.1126/science.1239999>
60. Albeck JG, Mills GB, Brugge JS. Frequency-modulated pulses of ERK activity transmit quantitative proliferation signals. *Molecular cell*. 2013; 49(2):249–61. <https://doi.org/10.1016/j.molcel.2012.11.002> PMID: 23219535
61. Tkačik G, Dubuis JO, Petkova MD, Gregor T. Positional Information, Positional Error, and Readout Precision in Morphogenesis: A Mathematical Framework. *Genetics*. 2015; 199(1):39–59. <https://doi.org/10.1534/genetics.114.171850> PMID: 25361898
62. Petkova M, Tkacik G, Bialek W, Wieschaus EF, Gregor T. Optimal decoding of cellular identities in a genetic network. *Cell*. 2019; 176:844. <https://doi.org/10.1016/j.cell.2019.01.007> PMID: 30712870
63. Heitz RP. The speed-accuracy tradeoff: history, physiology, methodology, and behavior. *Frontiers in neuroscience*. 2014; 8:150. <https://doi.org/10.3389/fnins.2014.00150> PMID: 24966810
64. Yatsenko D, Josić K, Ecker AS, Froudarakis E, Cotton RJ, Tolias AS. Improved Estimation and Interpretation of Correlations in Neural Circuits. *PLOS Computational Biology*. 2015; 11(3):e1004083. <https://doi.org/10.1371/journal.pcbi.1004083> PMID: 25826696