

## Trade-offs between cost and information in cellular prediction

Age J. Tjalma,<sup>1</sup> Vahe Galstyan,<sup>1</sup> Jeroen Goedhart,<sup>1</sup> Lotte Slim,<sup>1</sup> Nils B. Becker,<sup>2</sup> and Pieter Rein ten Wolde<sup>1,\*</sup>

<sup>1</sup>AMOLF, Science Park 104, 1098 XG Amsterdam, The Netherlands

<sup>2</sup>Theoretical Systems Biology, German Cancer Research Center, 69120 Heidelberg, Germany

(Dated: January 10, 2023)

Living cells can leverage correlations in environmental fluctuations to predict the future environment and mount a response ahead of time. To this end, cells need to encode the past signal into the output of the intracellular network from which the future input is predicted. Yet, storing information is costly while not all features of the past signal are equally informative on the future input signal. Here, we show, for two classes of input signals, that cellular networks can reach the fundamental bound on the predictive information as set by the information extracted from the past signal: push-pull networks can reach this information bound for Markovian signals, while networks that take a temporal derivative can reach the bound for predicting the future derivative of non-Markovian signals. However, the bits of past information that are most informative about the future signal are also prohibitively costly. As a result, the optimal system that maximizes the predictive information for a given resource cost is, in general, not at the information bound. Applying our theory to the chemotaxis network of *Escherichia coli* reveals that its adaptive kernel is optimal for predicting future concentration changes over a broad range of background concentrations, and that the system has been tailored to predicting these changes in shallow gradients.

Keywords: prediction, information bottleneck, sensing, resource allocation

Single-celled organisms live in a highly dynamic environment to which they continually have to respond and adapt. To this end, they employ a range of response strategies, tailored to the temporal structure of the environmental variations. When these variations are highly regular, such as the daily light variations, it becomes beneficial to develop a clock from which the time and hence the current and future environment can be inferred [1, 2]. In the other limit, when the fluctuations are entirely unpredictable, cells have no choice but to resort to either the strategy of detect-and-respond or the bet-hedging strategy of stochastic switching between different phenotypes [3]. Yet arguably the most fascinating strategy lies in between these two extremes. When the environmental fluctuations happen with some regularity, then it becomes feasible to predict the future environment and initiate a response ahead of time. While it is commonly believed that only higher organisms can predict the future, experiments have vividly demonstrated that even single-cell organisms can leverage temporal correlations in environmental fluctuations in order to predict, e.g., future nutrient levels [4, 5].

The ability to predict future signals can provide a fitness benefit [6]. The capacity to anticipate changes in oxygen levels [4], or the arrival of sugars or stress signals [5], can increase the growth rate of single-celled organisms; modeling has revealed that prediction can enhance bacterial chemotaxis [7]. Yet, a predict-and-anticipate strategy is only advantageous if the cell can reliably predict the future on timescales that are longer than the time it takes to mount a response. What fundamentally limits the accuracy of cellular prediction remains, however, poorly understood.

While the cell needs to predict the future environment, it can only sense the present and remember the past (Fig. 1A). Consequently, for a given amount of information the cell can store about the present and past signal, there is a maximum amount of information it can possibly have about the future [6, 8] (Fig. 1C-I). This *information bound* is determined by the temporal structure of the environmental fluctuations [8, 9].

How close cells can come to this bound depends on the design of the intracellular biochemical network that senses and processes the environmental signals (Fig. 1B). To maximize the predictive power the cell must use its memory effectively: it should extract only those characteristics from the present and past signal that are most informative about the future [7]. Whether it can do so, is determined by the topology of the signaling network. Moreover, like any information processing device, biochemical networks require resources to be built and run. Molecular components are needed to construct the network, space is required to accommodate the components, time is needed to process the information, and energy is required to synthesize the components and operate the network [10]. These resources constrain the design and performance of any biochemical network, and the capacity to sense and process information is no exception (Fig. 1C-II).

Cellular signaling systems provide a unique opportunity for revealing the resource requirements for prediction. Cells live in a highly dynamic environment, with temporal statistics that are expected to vary markedly. Moreover, signaling networks have distinct topologies, which are likely tailored to the temporal statistics of the environment [7]. In addition, for cellular systems we can actually quantify the information processing capacity as a function of the resources that are necessary to build and run them—protein copies, time, and energy [10, 11].

\* p.t.wolde@amolf.nl

Cellular systems are thus ideal for elucidating the relationships between future and past information, system design (i.e. network topology) and resource constraints. Here, we derive the bound on the prediction precision as set by the information extracted from the past signal for two types of input signals. We will determine how close cellular networks can come to this bound, and how this depends on the topology of the network and the resources to build and run it.

We find that for the two classes of input signals studied, cellular networks exists that can reach the information bound, yet reaching the bound is exceedingly costly. The first class of input signals consists of Markovian signals. Using the Information Bottleneck Method (IBM) [8, 12], we first show that the system that reaches the information bound copies the most recent input signal into the output from which the future input is predicted. Push-pull networks consisting of chemical modification or GTPase cycles, which are ubiquitous in prokaryotic and eukaryotic cells [13, 14], should be able to reach the information bound, because they are at heart copying devices [10, 11]. Yet, copying the most recent input into the output is extremely costly, because the operating cost, as set by the chemical power to drive the cycle, diverges at high copying speed. More surprisingly, our results show that the predictive and past information can be raised simultaneously by moving away from the information bound, even when the operating cost is negligible: the optimal system that maximizes the predictive information for a given protein synthesis cost is, in general, not at the information bound. The number of bits of past information per protein cost can be raised by increasing the integration time. While this decreases the predictive power per bit of past information, thereby moving the system away from the information bound, it can increase the total predictive information per protein cost. Our analysis thus highlights that not all bits of past information are equally costly, nor predictive.

Living cells that navigate their environment typically experience signals with persistence as generated by their own motion, which motivated us to study a simple class of non-Markovian signals. Moreover, these cells can typically detect changes in the concentration over a range of background concentrations that is orders of magnitude larger than the change in the concentration over the orientational correlation time of their movement. Our analysis reveals that in such a scenario the optimal kernel that allows the system to reach the information bound on predicting the future input derivative is a perfectly adaptive, derivative-taking kernel, precisely as the bacterium *E. coli* employs [15]. We again find, however, that reaching the information bound is prohibitively costly. The reason is that taking an instantaneous derivative, which is the characteristic of the input that is most informative about the future derivative, reduces the gain to zero because the system instantly adapts; the response becomes thwarted by biochemical noise. The optimal system that maximizes the predictive information under a resource

constraint thus emerges from a trade-off between taking a derivative that is recent and one that is reliable. Finally, our analysis reveals that the *E. coli* chemotaxis system has been optimally designed to predict future concentration changes in shallow gradients.

## RESULTS

We focus on cellular signaling systems that respond linearly to changes in the input signal [11, 16–19]. These systems not only allow for analytical results, but also describe information transmission often remarkably well [19–22]. The output of these systems can be written as

$$x(t) = \int_{-\infty}^t dt' k(t-t') \ell(t') + \eta_x(t), \quad (1)$$

where  $k(t)$  is the linear response function,  $\ell(t)$  the input signal, and  $\eta_x(t)$  describes the noise in the output. We will consider stationary signals with different temporal correlations, obeying Gaussian statistics.

Any prediction about the future state of the environment must be based on information obtained from its past (Fig. 1C-I). In particular, the cell needs to predict the input  $\ell_\tau \equiv \ell(t+\tau)$  at a time  $\tau$  into the future from the current output  $x_0 \equiv x(t)$ , which itself depends on the input signal in the past,  $\mathbf{L}_p \equiv (\ell(t), \ell(t'), \dots)$ , with  $t > t' > \dots$ . The (qualitative) shape of the integration kernel  $k(t)$ , e.g. exponential, adaptive or oscillatory, is determined by the topology of the signaling network [7]. The kernel shape describes how the past signal is mapped onto the current output, and hence which characteristics of the past signal the cell uses to predict the future signal. To maximize the accuracy of prediction, the cell should extract those features that are most informative about the future signal. These depend on the statistics of the input signal.

Deriving the upper bound on the predictive information as set by the past information is an optimisation problem, which can be solved using the IBM [8]. It entails the maximization of an objective function  $\mathcal{L}$ :

$$\max_{P(x_0|\mathbf{L}_p)} [\mathcal{L} \equiv I(x_0; \ell_\tau) - \gamma I(x_0; \mathbf{L}_p)]. \quad (2)$$

Here,  $I_{\text{pred}} \equiv I(x_0; \ell_\tau)$  is the predictive information, which is the mutual information between the system's current output  $x_0$  and the future ligand concentration  $\ell_\tau$ . The past information  $I_{\text{past}} \equiv I(x_0; \mathbf{L}_p)$  is the mutual information between  $x_0$  and the trajectory of past ligand concentrations  $\mathbf{L}_p$ . The Lagrange multiplier  $\gamma$  sets the relative cost of storing past over obtaining predictive information. Given a value of  $\gamma$ , the objective function in Eq. 2 is maximized by optimizing the conditional probability distribution of the output given the past input trajectory,  $P(x_0|\mathbf{L}_p)$ . For the linear systems considered here, this corresponds to optimizing the mapping of the past input signal onto the current output via the

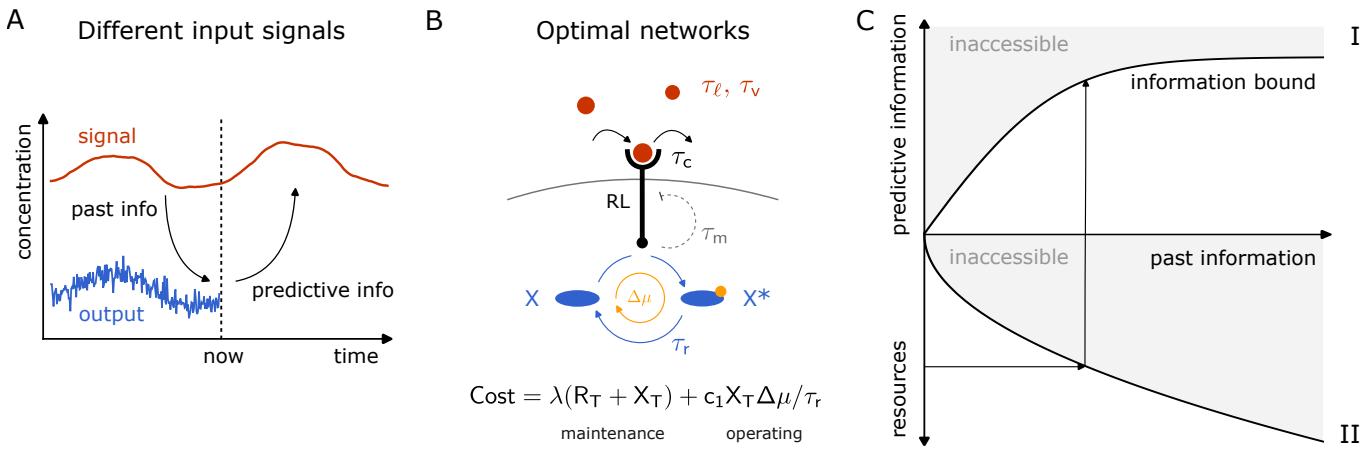


FIG. 1. Cells use biochemical networks to remember the past and predict the future. (A) Cells compress the past input into the dynamics of the signalling network from which the future input is then predicted. (B) The optimal topology of the network for predicting the future signal depends on the temporal statistics of the input signal. Push-pull networks, consisting of chemical modification cycles or GTPase cycles, can optimally predict the future value of Markovian signals, with correlation time  $\tau_\ell$ ; derivative-taking networks, like the *E. coli* chemotaxis system, can optimally predict the future derivative of non-Markovian signals, with correlation time  $\tau_v$ . The push-pull network consists of a receptor that drives a downstream phosphorylation cycle. The ligand binds the receptor with a correlation time  $\tau_c$ . The push-pull network, driven by ATP turnover, integrates the receptor with an integration time  $\tau_r$ . The chemotaxis system is a push-pull network, yet augmented with negative feedback on the receptor activity via methylation on a timescale  $\tau_m$ , as indicated by the dashed grey line. The total resource cost consists of a maintenance cost of receptor and readout synthesis at the growth rate  $\lambda$ , and an operating cost of driving the cycle. (C) The predictive information on the future signal  $I_{\text{pred}}$  is fundamentally bounded by how much information  $I_{\text{past}}$  it has about the past signal (panel I), which in turn is limited by the resources necessary to build and operate the biochemical network (panel II) [6].

integration kernel  $k(t)$ . Since our model obeys Gaussian statistics, we use the Gaussian IBM to derive the optimal kernel  $k^{\text{opt}}(t)$  and the *information bound*, defined to be the maximum predictive information as set by the past information [12] (see Appendix C).

### Markovian signals

#### Optimal prediction of Markovian signals: biochemical copying

Arguably the most elementary type of signal, albeit perhaps the hardest to predict, is a Markovian signal. We consider a Markovian signal  $\ell(t)$ , of which the deviations  $\delta\ell(t) = \ell(t) - \bar{\ell}$  from its mean  $\bar{\ell}$  follow an Ornstein-Uhlenbeck (OU) process:

$$\delta\dot{\ell} = -\delta\ell(t)/\tau_\ell + \eta_\ell(t), \quad (3)$$

where  $\tau_\ell$  is the correlation time of the fluctuations, and  $\eta_\ell(t)$  is Gaussian white noise,  $\langle \eta(t)\eta(t') \rangle = 2\sigma_\ell^2/\tau_\ell \delta(t - t')$ , with  $\sigma_\ell^2$  the amplitude of the signal fluctuations. This input signal obeys Gaussian statistics, characterized by  $\langle \delta\ell(0)\delta\ell(t) \rangle = \sigma_\ell^2 \exp(-t/\tau_\ell)$ . The optimal mapping is therefore a linear one. Utilizing the Gaussian IBM framework [12], we find that the optimal integration kernel is given by (see Appendix C2)

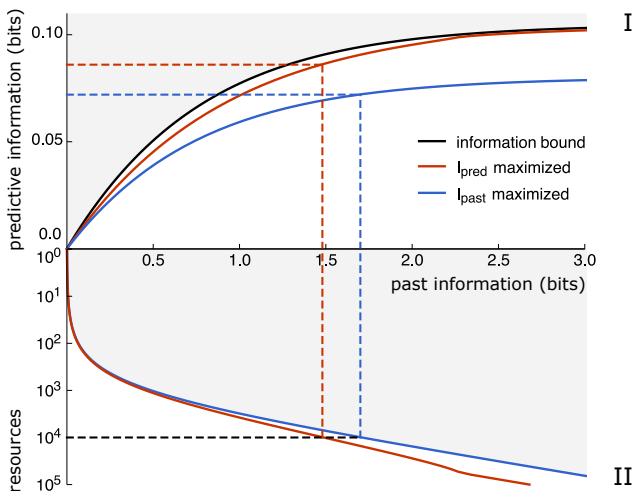
$$k^{\text{opt}}(t - t') = a\delta(t - t'). \quad (4)$$

This optimal integration kernel corresponds to a signaling system that copies the current input into the output. This is intuitive, since for a Markovian signal there is no additional information in the past signal that is not already contained in the present one. The prefactor  $a$  determines the gain  $\partial\bar{x}/\partial\bar{\ell}$ , which together with the noise strength  $\sigma_{\eta_x}^2$  (Eq. 1) and the signal amplitude  $\sigma_\ell^2$  set the magnitude of the past and predictive information,  $I_{\text{past}}$  and  $I_{\text{pred}}$ , respectively (see Appendix C1).

Fig. 2-I shows the maximum predictive information as set by the past information. This information bound applies to any linear system that needs to predict a Markovian signal. How close can biochemical systems come to this bound?

*Push-pull network can be at the information bound, yet increase the predictive and past information by moving away from it*

Although the upper bound on the accuracy of prediction is determined by the signal statistics, how close cells can come to this bound depends on the topology of the cellular signaling system, and the resources devoted to building and operating it. A network motif that could reach the information bound for Markovian signals is the push-pull network (Fig. 2), because it is at heart a copying device: it samples the input by copying the state of



**FIG. 2. The optimal push-pull network is not at the information bound.** Panel I: The black line is the information bound that maximizes the predictive information  $I_{\text{pred}} = I(x_0; \ell_\tau)$  for a given past information  $I_{\text{past}} = I(x_0; \mathbf{L}_p)$ . The red curve shows  $I_{\text{pred}}$  against  $I_{\text{past}}$  for systems in which  $I_{\text{pred}}$  has been maximized for a given resource cost  $C = R_T + X_T$ . The blue curve shows  $I_{\text{pred}}$  versus  $I_{\text{past}}$  for systems where  $I_{\text{past}}$  has been maximized for a given  $C$ . Panel II shows  $I_{\text{past}}$  against  $C$  for the corresponding systems. The forecast interval is  $\tau = \tau_\ell$ . The optimization parameters are the ratio  $X_T/R_T$ ,  $\tau_r$ ,  $p$  and  $f$  (see Appendix E). Parameter values:  $(\sigma_\ell/\bar{\ell})^2 = 10^{-2}$ ,  $\tau_c/\tau_\ell = 10^{-2}$ .

the input, e.g. the ligand-binding state of a receptor or the activation state of a kinase, into the activation state of the output, e.g. phosphorylation state of the readout [10, 11, 23].

We model the push-pull network in the linear-noise approximation:

$$\delta \dot{R}L = b\delta\ell(t) - \delta RL(t)/\tau_c + \eta_{RL}(t), \quad (5)$$

$$\delta \dot{x}^* = \gamma \delta RL(t) - \delta x^*(t)/\tau_r + \eta_x(t). \quad (6)$$

Here,  $\delta RL$  represents the number of ligand-bound receptors and  $\delta x^*$  the number of modified readout molecules, defined as deviations from their mean values;  $b$  and  $\gamma$  are parameters that depend on the number of receptor and readout molecules,  $R_T$  and  $X_T$  respectively, the fraction of ligand-bound receptors  $p$  and active readout molecules  $f$ ;  $\eta_{RL}$  and  $\eta_x$  are Gaussian white noise terms (see Appendix E). Key parameters are the correlation time of receptor-ligand binding,  $\tau_c$ , and the relaxation time of  $x^*$ ,  $\tau_r$ . The latter determines for how long  $x^*$  carries information on the ligand-binding state of the receptor and thus sets the integration time. The readout-modification dynamics yield an exponential integration kernel  $k(t) \propto \exp(-t/\tau_r)$ , which in the limit  $\tau_r \rightarrow 0$  reduces to a  $\delta$ -function, hinting that the system may reach the information bound.

How much information cells can extract from the past signal depends on the resources devoted to building and

operating the network (Fig. 2-II). We define the total resource cost to be:

$$C = \lambda(R_T + X_T) + c_1 X_T \Delta\mu/\tau_r \quad (7)$$

The first term expresses the fact that over the course of the cell cycle all components need to be duplicated, which means that they have to be synthesized at a speed that is at least the growth rate  $\lambda$ . The second term describes the chemical power that is necessary to run the push-pull network [10, 11]; it depends on the flux through the network,  $X_T/\tau_r$ , and the free-energy drop  $\Delta\mu$  over a cycle, e.g. the free energy of ATP hydrolysis in the case of a phosphorylation cycle. The coefficient  $c_1$  describes the relative energetic cost of synthesising the components during the cell cycle versus that of running the system. For simplicity, we first consider the scenario that the cost is dominated by that of protein synthesis, setting  $c_1 \rightarrow 0$ . While in this scenario  $R_T + X_T$  is constrained,  $X_T/R_T$  and other system parameters are free for optimization.

The available resources put a hard bound on the information  $I_{\text{past}}$  that can be extracted from the past signal, which in turn sets a hard limit on the predictive information  $I_{\text{pred}}$  (Fig. 1C). To maximize the predictive information, it therefore seems natural to maximize the past information  $I_{\text{past}}$  for a given resource cost  $C$ . The blue line in Fig. 2-II shows the result for the push-pull network. We then compute the corresponding predictive information for the systems along this line, which is the blue line in Fig. 2-I. Strikingly, the resulting information curve lies far below the information bound, i.e. the upper bound on the predictive information as set by the past information (black line, Fig. 2-I). This shows that systems that maximize past information under a resource constraint, do not in general also maximize predictive information. It implies that not all bits of past information are equally predictive about the future.

Precisely because not all bits of past information are equally predictive about the future, it is paramount to directly maximize the predictive information for a given resource cost in order to obtain the most efficient prediction device. This yields the red lines in panels I and II in Fig. 2. It can be seen that the predictive information is higher while the past information is lower, as compared to the information curves of the systems optimized for maximizing the past information under a resource constraint (blue lines). It reflects the idea that not all bits are equally predictive. More surprisingly, while the bound on the predictive information as set by the resource cost (red line panel I) is close to the bound on the predictive information as set by the past information (black line), it does remain lower. This is surprising, because the push-pull network is a copying device [10, 23], which can, as we will also show below, reach the latter bound. These two observations together imply that not all bits of past information are equally costly. If they were, the cell would select under the two constraints the same bits based on their predictive information content, and the bound on the predictive information as set by

the resource cost would overlap with that as set by the past information.

We thus find that not all bits of past information are equally predictive, nor equally costly. As we show next, it implies that the optimal information processing system faces a trade-off between using those bits of past information that are most informative about the future and those that are cheapest.

#### *Trade-off between cost and predictive power per bit*

To understand the connection between predictive and past information, and resource cost, we map out the region in the information plane that can be reached given a resource constraint  $C$  (Fig. 3A, green region). We immediately make two observations. Firstly, the system can indeed reach the information bound. Secondly, the system can increase both the past and the predictive information by moving away from the bound. To elucidate these two observations, we investigate the system along the isocost line of  $C = 10^4$ , which together with the information bound envelopes the accessible region for the maximum resource cost  $C \leq 10^4$ .

Along the isocost line, the ratio of the number of readout over receptor molecules is  $X_T/R_T = 2\sqrt{p/(1-p)}\sqrt{1+\tau_r/\tau_c}$  (see Appendix E3). This can be understood intuitively using the optimal resource allocation principle [10]. It states that in a sensing system that employs its proteins optimally, the total number of independent concentration measurements at the level of the receptor during the integration time  $\tau_r$ ,  $R_T(1+\tau_r/\tau_c)$ , equals the number of readout molecules  $X_T$  that store these measurements, so that neither the receptors nor the readout molecules are in excess. This design principle specifies, for a given integration time  $\tau_r$ , the ratio  $X_T/R_T$  at which the readout molecules sample each receptor molecule roughly once every receptor correlation time  $\tau_c$ .

While the optimal allocation principle gives the optimal ratio  $X_T/R_T$  of the number of readouts over receptors for a *given* integration time  $\tau_r$ , it does not prescribe what the optimal integration time  $\tau_r^{\text{opt}}$ , and hence (globally) optimal ratio  $X_T^{\text{opt}}/R_T^{\text{opt}}$ , is that maximizes  $I_{\text{pred}}$  for a given resource constraint  $C = R_T + X_T$ . Fig. 3B shows that as the distance  $\theta$  along the isocost line is increased,  $\tau_r$  and hence  $X_T/R_T$  increase monotonically. Near the information bound, corresponding to  $\theta = 0$ , the integration time  $\tau_r$  is zero and the number of readout molecules equals the number of receptor molecules:  $X_T = R_T$ . In this limit, the push-pull network is an instantaneous responder, with an integration kernel given by Eq. 4; only the finite receptor correlation time  $\tau_c$  prevents the system from fully reaching the information bound. Yet, as  $\theta$  increases and the system moves away from the bound, the predictive and past information first rise along the contour, and thus with  $X_T/R_T$  and  $\tau_r$ , before they eventually both fall.

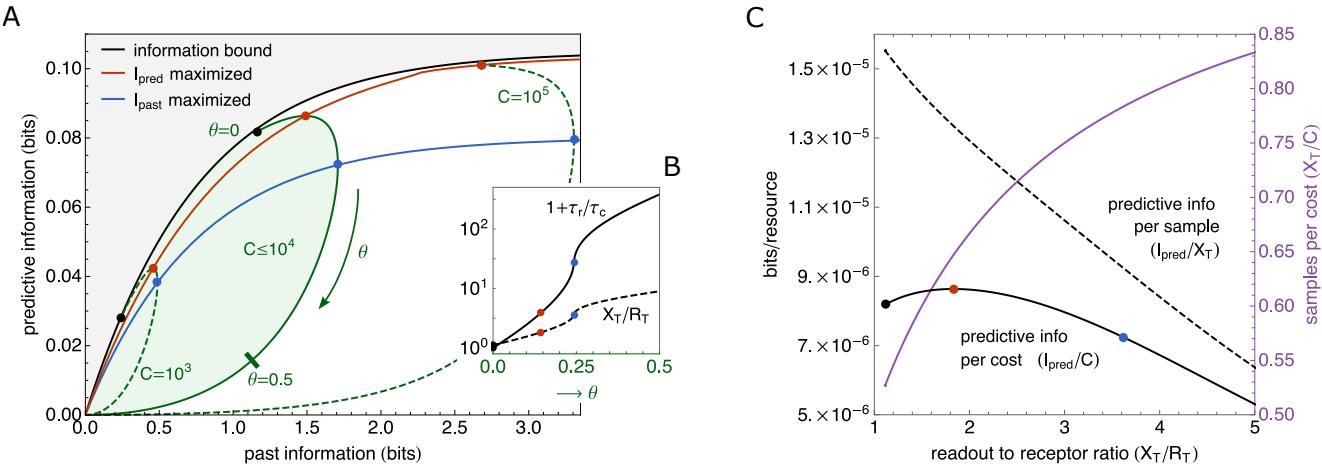
To understand why the predictive and past information first rise and then fall with  $X_T/R_T$  and  $\tau_r$ , we note that each readout molecule constitutes 1 physical bit and that its binary state (phosphorylated or not) encodes at most 1 bit of information on the ligand concentration. The number of readout molecules  $X_T$  thus sets a hard upper bound on the sensing precision and hence the predictive information. To raise this bound,  $X_T$  must be increased. For a given resource constraint  $C = R_T + X_T$ ,  $X_T$  can only be increased if the number of receptors  $R_T$  is simultaneously decreased. However, the cell infers the concentration not from the readout molecules directly, but via the receptor molecules: a readout molecule is a sample of the receptor that provides at most 1 bit of information about the ligand-binding state of a receptor molecule, which in turn provides at most 1 bit of information about the input signal. To raise the lower bound on the predictive information, the information on the input must increase at both the receptor and the readout level.

To elucidate how this can be achieved, we note that the maximum number of independent receptor samples and hence concentration measurements is given by  $N_I^{\text{max}} = \min(X_T, R_T(1 + \tau_r/\tau_c))$  [10]. For  $\theta > 0$ , the system can increase  $N_I^{\text{max}}$  if, and only if,  $X_T$  and  $R_T(1 + \tau_r/\tau_c)$  can be raised simultaneously. This can be achieved, while obeying the constraint  $C = X_T + R_T$ , by decreasing  $R_T$  yet increasing  $\tau_r$  (Fig. 3B). This is the mechanism of time averaging, which makes it possible to increase the number of independent receptor samples [11], and explains why both the predictive and the past information initially increase (Fig. 3C). However, as  $\tau_r$  is raised further, the receptor samples become older: the readout molecules increasingly reflect receptor states in the past that are less informative about the future ligand concentration. The collected bits of past information have become less predictive about the future (Fig. 3C). For a given resource cost, the cell thus faces a trade-off between maximizing the number of physical bits of past information (i.e. the receptor samples  $X_T$ ) and the predictive information per bit. This antagonism gives rise to an optimal integration time  $\tau_r^{\text{opt}}$  that maximizes the total predictive information  $I_{\text{pred}}$  (Fig. 3C).

Interestingly, while  $I_{\text{pred}}$  decreases beyond  $\tau_r^{\text{opt}}$ , the past information  $I_{\text{past}}$  first continues to rise because  $N_I^{\text{max}}$  still increases. However, when the integration time becomes longer than the input signal correlation time, the correlation between input and output will be lost and  $I_{\text{past}}$  will fall too.

#### *Chemical power prevents the system from reaching the information bound*

So far, we have only considered the cost of maintaining the cellular system, the protein cost  $C = R_T + X_T$ . Yet, running a push-pull network also requires energy. As Eq. 7 shows, the running cost scales with the flux



**FIG. 3. The push-pull network maximizes the predictive power under a resource constraint by moving away from the information bound.** (A) The region of accessible predictive information  $I_{\text{pred}} = I(x_0; \ell_\tau)$  and past information  $I_{\text{past}} = I(x_0; \mathbf{L}_p)$  in the push-pull network under a resource constraint  $C \leq (R_T + X_T)$ , for the Markovian signals specified by Eq. 3 (green). The black line is the information bound at which  $I_{\text{pred}}$  is maximized for a given  $I_{\text{past}}$ . The push-pull network can be at the information bound (black points), but maximizing  $I_{\text{pred}}$  for a resource constraint  $C$  moves the system away from it. The red and blue lines connect, respectively, the points where  $I_{\text{pred}}$  and  $I_{\text{past}}$  are maximized along the green isocost lines (the contourlines of constant  $C$ ); they correspond to the red and blue lines in Fig. 2, respectively. The accessible region of  $I_{\text{pred}}$  and  $I_{\text{past}}$  for a given  $C$  has been obtained by optimizing over  $\tau_r$ ,  $p$ ,  $f$ , and  $X_T/R_T$ . The forecast interval is  $\tau = \tau_\ell$ . (B) The integration time  $\tau_r$  over the receptor correlation time  $\tau_c$ ,  $\tau_r/\tau_c$ , and the ratio of the number of readout and receptor molecules,  $X_T/R_T$ , as a function of the distance  $\theta$  along the isocost line corresponding to  $C = 10^4$  in panel A; the red and blue points denote where  $I_{\text{pred}}$  and  $I_{\text{past}}$  are maximized along the contourline, respectively. For  $\theta \rightarrow 0$ ,  $\tau_r \rightarrow 0$ : the system is an instantaneous responder, which is essentially at the information boundary; as predicted by the optimal resource allocation principle,  $X_T = R_T$ . The system can increase  $I_{\text{pred}}$  and  $I_{\text{past}}$  by increasing  $\tau_r$  and  $X_T/R_T$ . (C) While this decreases the predictive information  $I_{\text{pred}}$  per physical bit of past information,  $I_{\text{pred}}/X_T$  (dashed line), increasing  $X_T/R_T$  does increase the number of physical bits per resource cost,  $X_T/C$  (purple line). This trade-off gives rise to an optimal predictive information per resource cost,  $I_{\text{pred}}/C$  (red dot on solid black line). Parameter values unless specified:  $(\sigma_\ell/\ell)^2 = 10^{-2}$ ,  $\tau_c/\tau_\ell = 10^{-2}$ .

around the phosphorylation cycle, which is proportional to the inverse of the integration time,  $\tau_r^{-1}$ . The power thus diverges for  $\tau_r \rightarrow 0$ . Since the information bound is reached precisely in this limit, it is clear that the chemical power prevents the push-pull network from reaching the bound (see Fig. 3 in the appendix).

### Non-Markovian signals

#### Predicting the future change

The push-pull network can optimally predict Markovian signals, yet not all signals are expected to be Markovian. Especially organisms that navigate through an environment with directional persistence will sense a non-Markovian signal, as generated by their own motion. Moreover, when these organisms need to climb a concentration gradient, as *E. coli* during chemotaxis, then knowing the change in the concentration is arguably more useful than knowing the concentration itself. Indeed, it is well known that the kernel of the *E. coli* chemotaxis system detects the (relative) change in the ligand concentration by taking a temporal derivative of the concen-

tration [15]. However, as we will show here, the converse statement is more subtle. If the system needs to predict the (future) change in the signal, then the optimal kernel is not necessarily one that is based on the derivative only: in general, the optimal kernel uses a combination of the signal value and its derivative. However, the *E. coli* chemotaxis system can respond to concentrations that vary between the dissociation constants of the inactive and active state of the receptors, which differ by several orders of magnitude [24]. This range of possible background concentrations is much larger than the typical concentration change over the orientational correlation time of the bacterium. As our analysis below reveals, in this regime the optimal kernel is a perfectly adaptive, derivative-taking kernel that is insensitive to the current signal value, precisely like that of the *E. coli* chemotaxis system [15, 25–28]. Our analysis thus predicts that this system has an adaptive kernel, because this is the optimal kernel for predicting concentration derivatives over a broad range of background concentrations.

To reveal the signal characteristics that control the shape of the optimal integration kernel, we will consider the family of signals that are generated by a harmonic

oscillator:

$$\delta\dot{\ell} = v(t), \quad (8)$$

$$\dot{v} = -\omega_0^2 \delta\ell(t) - v(t)/\tau_v + \eta_v(t), \quad (9)$$

where  $\delta\ell$  is the deviation of ligand concentration from its mean  $\bar{\ell}$ ,  $v$  its derivative,  $\tau_v$  a relaxation time,  $\eta_v$  a Gaussian white noise term, and the frequency  $\omega_0^2 = \sigma_v^2/\sigma_\ell^2$  controls the variance  $\sigma_\ell^2$  of the concentration and that of its derivative  $\sigma_v^2$ .

Using the IBM framework it can be shown that the optimal encoding that allows the system to reach the information bound, is based on a linear combination of the current concentration  $\ell(t)$  and its derivative  $v(t)$ , such that the output  $x(t)$  is given by (Appendix C3):

$$x(t) = a \frac{\delta\ell(t)}{\sigma_\ell} + b \frac{v(t)}{\sigma_v} + \eta_x(t). \quad (10)$$

This can be understood by noting that while the signal of Eqs. 8 and 9 is non-Markovian in the space of  $\ell$ , it is Markovian in  $\ell$  and  $v$ : all the information on the future signal is thus contained in the current concentration and its derivative. To maximize the predictive information  $I_{\text{pred}} = I(x_0; v_\tau)$  between the current output  $x_0$  and the future derivative of the input  $v_\tau$  for a given amount of past information  $I_{\text{past}} = I(x_0; \mathbf{L}_p)$ , i.e. to reach the information bound for predicting the future signal derivative, the coefficients must obey

$$a^{\text{opt}} = G \frac{\langle \delta\ell(0) \delta v(\tau) \rangle}{\sigma_\ell \sigma_v} \equiv G \rho_{\ell_0 v_\tau}, \quad (11)$$

$$b^{\text{opt}} = G \frac{\langle \delta v(0) \delta v(\tau) \rangle}{\sigma_v^2} \equiv G \rho_{v_0 v_\tau}. \quad (12)$$

Here,  $G$  is the gain, which together with the noise  $\sigma_{\eta_x}^2$  sets the scale of  $I_{\text{pred}}$  and  $I_{\text{past}}$ ,  $\rho_{\ell_0 v_\tau}$  is the cross-correlation coefficient between the current concentration value  $\ell_0$  and the future concentration derivative  $v_\tau$  and  $\rho_{v_0 v_\tau}$  that between the current and future derivative (Appendix C3). These expressions can be understood intuitively: if the future signal derivative that needs to be predicted is correlated with the current signal derivative, it is useful to include in the prediction strategy the current signal derivative, leading to a non-zero value of  $b^{\text{opt}}$ . Perhaps more surprisingly, if the future signal derivative is also correlated with the current signal *value*, then the system can enhance the prediction accuracy by also including the current signal value, yielding a non-zero  $a^{\text{opt}}$ . Clearly, in general, to optimally predict the future signal change, the system should base its prediction on both the current signal value and its derivative.

The degree to which the system bases its prediction on the current value versus the current derivative depends on the relative magnitudes of  $a^{\text{opt}}$  and  $b^{\text{opt}}$ , respectively. In Appendix B2, we show that when the concentration change over the timescale  $\tau_v$ ,  $\sigma_v \tau_v$ , is much smaller than the range of possible concentrations  $\sigma_\ell$  that the bacterium can experience, i.e. when  $\sigma_v \tau_v \ll \sigma_\ell$  such that

$\omega_0 \ll \tau_v^{-1}$ , the cross-correlation coefficient  $\rho_{\ell_0 v_\tau}$  vanishes, such that  $a^{\text{opt}}$  becomes zero (see Eq. 11). The optimal kernel has become a perfectly adaptive, derivative-taking kernel. We emphasize that while we have derived this result for the class of signals defined by Eqs. 8 and 9, the idea is far more generic. In particular, while we do not know the temporal structure of the ligand statistics that *E. coli* experiences, we do know that it can detect concentration changes over a range of background concentrations that is much wider than the typical concentration change over a run, such that the correlation between the concentration value and its future change is likely to be very small. As our analysis shows, a perfectly adaptive kernel then emerges naturally from the requirement to predict the future concentration change.

While the class of signals specified by Eqs. 8 and 9 is arguably limited, it does describe the biologically important regime of chemotaxis in shallow gradients. In the limit that  $\omega_0 \ll \tau_v^{-1}$ , Eq. 9 reduces to  $\dot{v} = -v/\tau_v + \eta_v$ . In shallow gradients, the stimulus only weakly affects the swimming behavior, such that the perceived signal is mostly determined by the intrinsic orientational dynamics of the bacterium in the absence of a gradient. In this regime, the temporal statistics of the concentration derivative  $v$  is completely determined by the steepness of the concentration gradient  $g$  and the swimming statistics of the bacterium in the absence of a gradient:

$$\langle \delta v(0) \delta v(\tau) \rangle = g^2 \bar{\ell}^2 \langle \delta v_x(0) \delta v_x(\tau) \rangle \simeq \sigma_{v_x}^2 e^{-\tau/\tau_{v_x}}, \quad (13)$$

where the latter is the autocorrelation function of the (positional) velocity of the bacterium in the absence of a gradient. It is a characteristic of the bacterium, not of the environment, and has been measured to decay exponentially with a correlation time  $\tau_{v_x}$  [18], precisely as our model, with  $\tau_v = \tau_{v_x}$ , predicts. This correlation time is on the order of the typical run time of the bacterium in the absence of a gradient,  $\tau_v \sim 0.9$  s [18].

#### Finite resources prevent the chemotaxis system from taking an instantaneous derivative and reaching the information bound

The above analysis indicates that the chemotaxis system seems ideally designed to predict the future concentration change, because its integration kernel is nearly perfectly adaptive [15, 25–28]. But how close can this system come to the information bound for the non-Markovian signals specified by Eqs. 8 and 9?

To address this, we consider a molecular model that can accurately describe the response of the chemotaxis system to a wide range of time-varying signals [29–32]. In this model, the receptors are partitioned into clusters. Each cluster is described via a Monod-Wyman-Changeux model [33]. While each receptor can switch between an active and an inactive conformational state, the energetic cost of having different conformations in the same cluster is prohibitively large. Each cluster is thus either active or

inactive. Ligand binding favors the inactive state while methylation does the opposite. Lastly, active receptor clusters can via the associated kinase CheA phosphorylate the downstream messenger protein CheY.

Linearizing around the steady state, we obtain:

$$\delta a_i(t) = \alpha \delta m_i(t) - \beta \delta \ell(t), \quad (14)$$

$$\delta \dot{m}_i = -\delta a_i(t) / (\alpha \tau_m) + \eta_{m_i}(t), \quad (15)$$

$$\delta \dot{x}^* = \gamma \sum_{i=1}^{R_T} \delta a_i(t) - \delta x^*(t) / \tau_r + \eta_x(t). \quad (16)$$

Here,  $\delta a_i(t)$  and  $\delta m_i(t)$  are the deviations of the activity and methylation level of receptor cluster  $i$  from their steady-state values, and  $R_T$  is the total number of receptor clusters;  $\delta \ell(t)$  and  $\delta x^*(t)$  are, respectively, the deviations of the ligand and CheY<sub>p</sub> concentration from their steady-state values;  $\tau_m$  and  $\tau_r$  are the timescales of receptor methylation and CheY<sub>p</sub> dephosphorylation;  $\eta_{m_i}$  and  $\eta_x$  are independent Gaussian white noise sources. In Eq. 14, we have assumed that ligand binding is much faster than the other timescales in the system, so that it can be integrated out. There is therefore no need to time average receptor-ligand binding noise, which means that, in the absence of running costs, the optimal receptor integration time  $\tau_r$  is zero. In what follows, we set  $\tau_r$  to the value measured experimentally,  $\tau_r \approx 100\text{ms}$  [10, 34]. We consider the non-Markovian signals specified by Eqs. 8 and 9 in the physiologically relevant limit  $\omega_0 \rightarrow 0$ , such that the optimal kernel is perfectly adaptive, like that of *E. coli*. For these signals, we determine the accessible region of  $I_{\text{past}}$  and  $I_{\text{pred}}$  under a resource constraint  $C = R_T + X_T$  (see Fig. 4) by optimizing over the methylation time  $\tau_m$  and the ratio of readout over receptor molecules  $X_T/R_T$ . The forecast interval  $\tau$  is set to  $\tau_v$ , but we emphasize that the optimal design is independent of the value of  $\tau$  (see Appendix F4).

Fig. 4A shows that the chemotaxis system is, in general, not at the information bound that maximizes the predictive information  $I_{\text{pred}} = I(x_0; v_\tau)$  for a given past information  $I_{\text{past}} = I(x_0; \mathbf{L}_p)$ . The optimal systems that maximize  $I_{\text{pred}}$  under a resource constraint  $C$ , marked by the red dots, are indeed markedly away from the information bound. Yet, as the resource constraint is relaxed and  $C$  is increased, the optimal system moves towards the bound. Panel B shows that the methylation time  $\tau_m$  rises along the three respective isocost lines of panel A. It highlights that there exists an optimal methylation time  $\tau_m^{\text{opt}}$  that maximizes the predictive information  $I_{\text{pred}}$ . Moreover,  $\tau_m^{\text{opt}}$  decreases as the resource constraint is relaxed. Along the respective isocost lines,  $X_T/R_T$  varies only mildly (see Fig. 5 in the appendix).

These observations can be understood by noting that the system faces a trade-off between taking a derivative that is recent versus one that is robust. All the information on the future derivative, which the cell aims to predict, is contained in the current derivative of the signal; measuring the current derivative would allow the

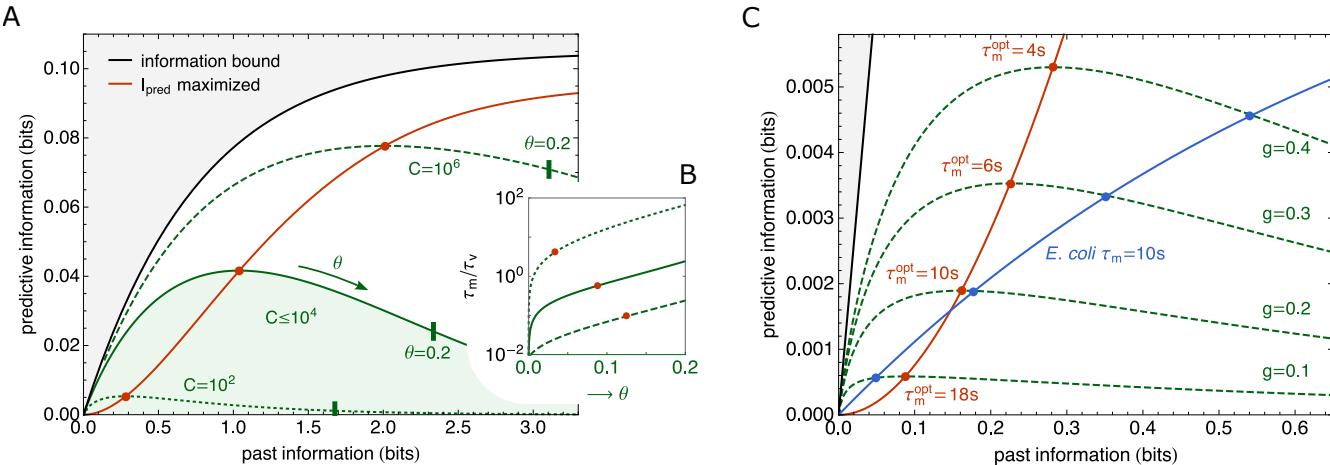
system to reach the information bound. However, computing the recent derivative is extremely costly. The cell takes the temporal derivative of the ligand concentration at the level of the receptor via two antagonistic reactions that occur on two distinct timescales: ligand binding rapidly deactivates the receptor, while methylation slowly reactivates it [30]. The receptor ligand-occupancy thus encodes the current concentration, the methylation level stores the average concentration over the past  $\tau_m$ , and the receptor activity reflects the difference between the two—the temporal derivative of the signal over the timescale  $\tau_m$ . To obtain an instantaneous derivative,  $\tau_m$  must go to zero. However, this dramatically reduces the gain; in fact, in this limit, the gain is zero, because the receptor activity instantly adapts to the change in the ligand concentration. Since the push-pull network downstream of the receptor is a device that samples the receptor stochastically [10, 36], the gain, i.e. the change in the receptor activity due to the signal, must be raised to lift the signal above the sampling noise. This requires a finite methylation time  $\tau_m$ : as we show in Appendix F3, the gain increases monotonically with  $\tau_m$ . The trade-off between a recent derivative and a reliable one gives rise to an optimal methylation time  $\tau_m^{\text{opt}}$  that maximizes the predictive information for a given resource cost.

The same analysis also explains why the optimal methylation time  $\tau_m^{\text{opt}}$  decreases and the predictive information increases when the resource constraint is relaxed. The sampling noise in estimating the average receptor activity decreases as the number of readout molecules increases [10, 36]. A smaller gain is thus required to lift the signal above the sampling noise. In addition, a larger number of receptors decreases the noise in the methylation level, which also allows for a smaller gain, and hence a smaller methylation time. These two effects together explain why  $\tau_m^{\text{opt}}$  decreases and  $I_{\text{pred}}$  increases with  $C$ .

Fig. 4A also shows that the past information  $I_{\text{past}} = I(x_0; \mathbf{L}_p)$  does not return to zero along the contourline of constant resource cost. Along the contourline, the methylation time  $\tau_m$  rises (Fig. 4B). While the predictive information  $I_{\text{pred}}$  exhibits an optimal methylation time  $\tau_m^{\text{opt}}$ , the past information  $I_{\text{past}}$  continues to rise with  $\tau_m$  because the system increasingly becomes a copying device, rather than one that takes a temporal derivative.

### Comparison with experiment

To test our theory, we study the predictive power of the *E. coli* chemotaxis system as a function of the steepness of the ligand concentration gradient, keeping the resource constraint at the biologically relevant value of  $C = R_T + X_T = 10^4$  [35]. Panel C of Fig. 4 shows  $I_{\text{pred}}$  and  $I_{\text{past}}$  for cells swimming in an exponential concentration gradient  $\ell(x) = \ell_0 e^{gx}$ , for different values of the gradient steepness  $g$ ; along the green iso-steepness lines  $\tau_m$  is varied and  $X_T/R_T$  is optimized to maximize  $I_{\text{pred}}$  and  $I_{\text{past}}$ , with the red dots marking  $\tau_m^{\text{opt}}$ , while along



**FIG. 4. Finite resources prevent chemotaxis system from reaching the information bound.** (A) The region of accessible predictive information  $I_{\text{pred}} = I(x_0; v_\tau)$  and past information  $I_{\text{past}} = I(x; v_p)$  for the chemotaxis system under a resource constraint  $C = R_T + X_T$ , for the non-Markovian signals specified by Eqs. 8 and 9 (green). The black line shows the information bound at which  $I_{\text{pred}}$  is maximized for a given  $I_{\text{past}}$ . The chemotaxis system is not at the information bound, but it does move towards it as  $C$  is increased. The red line connects the red points where  $I_{\text{pred}}$  is maximized for a given resource cost  $C$ . The accessible region of  $I_{\text{pred}}$  and  $I_{\text{past}}$  under a given resource constraint  $C = R_T + X_T$  is obtained by optimizing over the methylation time  $\tau_m$  and the ratio of readout over receptor molecules  $X_T/R_T$ . The forecast interval is  $\tau = \tau_v$ . (B) The methylation time  $\tau_m$  over the input correlation time  $\tau_v$  as a function of the distance  $\theta$  along the three respective isocost lines shown in panel A. The methylation time  $\tau_m$  increases along the isocost line, but there exists an optimal  $\tau_m$  that maximizes the predictive information, marked by the red points;  $\theta \rightarrow 0$  corresponds to the origin of panel A,  $(I_{\text{pred}}, I_{\text{past}}) = (0, 0)$ ; the points where  $\theta = 0.2$  along the isocost lines of panel A are marked with a bar. As the resource constraint is relaxed (higher  $C$ ), the optimal  $\tau_m$  decreases: the system moves towards the information bound, where it takes an instantaneous derivative, corresponding to  $\tau_r, \tau_m \rightarrow 0$ . (C) The contourlines of  $I_{\text{pred}}$  and  $I_{\text{past}}$  for increasing values of the steepness  $g$  of an exponential ligand concentration gradient  $\ell(x) = \ell_0 e^{gx}$ , keeping the total resource cost fixed at  $C = R_T + X_T = 10^4$ ;  $\tau_m$  and  $X_T/R_T$  have been optimized. It is seen that the maximal predictive information  $I_{\text{pred}}$  under the resource constraint  $C$  (marked by the red points) increases with the gradient steepness. The blue line shows  $I_{\text{pred}}$  and  $I_{\text{past}}$  for the *E. coli* chemotaxis system with  $\tau_m = 10\text{s}$  and  $X_T = R_T = 5000$  fixed at their measured values [35]. Our analysis predicts that this system has been optimized to detect shallow gradients. Parameter values unless specified:  $\tau_r = 100\text{ms}$  [10, 34];  $\tau_v = 0.9\text{s}$  and  $\sigma_v^2 = g^2 \bar{\ell}^2 \sigma_{v_x}^2$ , with  $\bar{\ell} = 100\mu\text{M}$  and  $\sigma_{v_x}^2 = 157.1\mu\text{m}^2\text{s}^{-2}$  [18];  $\omega_0 \rightarrow 0$ ;  $g$  is given in units of  $\text{mm}^{-1}$ ; in A,  $g = 4/\text{mm}$ .

the blue line  $\tau_m$  and  $X_T$  and  $R_T$  are fixed at their experimentally measured values [29, 30, 35]. Clearly, both the predictive and the past information rise as the gradient steepness  $g$  increases—a steeper concentration gradient yields a larger change in the concentration, and thus a stronger signal.

More interestingly, in the optimal system  $I_{\text{pred}}$  rises much faster with  $I_{\text{past}}$  (red line) than in the *E. coli* system (blue line). A steeper gradient  $g$  yields a stronger input signal, which raises the signal above the sampling noise more. This allows the optimal system to take a more recent derivative, with a smaller  $\tau_m$ , which is more informative about the future. In contrast, the methylation time  $\tau_m$  of the *E. coli* chemotaxis system is fixed. As Fig. 4C shows, this value is beneficial for detecting shallow gradients,  $g \lesssim 0.2\text{mm}^{-1}$ . Moreover, in this regime, not only  $I_{\text{pred}}$  but also  $I_{\text{past}}$  are close to the respective values for the optimal system. For steeper gradients  $I_{\text{past}}$  becomes much higher in the *E. coli* system than in the optimal one, even though  $I_{\text{pred}}$  remains lower. The bacterium increasingly collects information that is less informative about the future. Taken together, these results strongly

suggest that the system has been optimized to predict future concentration changes in shallow gradients, which necessitate a relatively long methylation time.

## DISCUSSION

Cellular systems need to predict the future signal by capitalizing on information that is contained in the past signal. To this end, they need to encode the past signal into the dynamics of the intracellular biochemical network from which the future input is inferred. To maximize the predictive information for a given amount of information that is extracted, the cell should store those signal characteristics that are most informative about the future signal. For a Markovian signal obeying an Ornstein-Uhlenbeck process this is the current signal value, while for the non-Markovian signal corresponding to an underdamped particle in a harmonic well, this is the current signal value and its derivative. As we have seen here, cellular systems are able to extract these sig-

nal characteristics: the push-pull network can copy the current input into the output, while the chemotaxis network can take an instantaneous derivative. We have thus demonstrated that at least for two classes of signals, cellular systems are in principle able to extract the most predictive information, allowing them to reach the information bound.

Yet, our analysis also shows that extracting the most relevant information can be exceedingly costly. To copy the most recent input signal into the output, the integration time of the push-pull network needs to go to zero, which means that the chemical power diverges. Moreover, taking an instantaneous derivative reduces the gain to zero, such that the signal is no longer lifted above the inevitable intrinsic biochemical noise of the signalling system. In fact, taking the chemical power cost to drive the adaptation cycle into account [27, 37] would push the system away from the information bound even more.

While information is a resource—the cell cannot predict the future without extracting information from the past signal—the principal resources that have a direct cost are time, building blocks and energy. The predictive information per protein and energy cost is therefore most likely a more relevant fitness measure than the predictive information per past information. Our analysis reveals that, in general, it is not optimal to operate at the information bound: cells can increase the predictive information for a given resource constraint by moving away from the bound. Increasing the integration time in the push-pull network reduces the chemical power and makes it possible to take more concentration measurements per protein copy. And increasing the methylation time in the chemotaxis system increases the gain. Both enable the system to extract more information from the past signal. Yet, increasing the integration time or the methylation time also means that the information that has been collected, is less informative about the future signal. This interplay gives rise to an optimal integration and methylation time, which maximize the predictive information for a given resource constraint. This argument also explains why the respective systems move towards the information bound when the resource constraint is relaxed: Increasing the number of receptor and readout molecules allows the system to take more instantaneous concentration measurements, which makes time averaging less important, thus reducing the integration time.

Increasing the number of readout molecules also reduces the error in sampling the receptor state. This makes it easier to detect a change in the receptor activity resulting from the signal, thus allowing for a smaller dynamical gain and a shorter methylation time.

Information theory shows that the amount of transmitted information depends not only on the characteristics of the information processing system, but also on the statistics of the input signal. While much progress has been made in characterizing cellular signalling systems, the statistics of the input signal is typically not known, with a few notable exceptions [38]. Here, we have focussed on two classes of input signals, but it seems likely that the signals encountered by natural systems are much more diverse. It will be interesting to extend our analysis to signals with a richer temporal structure [9], and see whether cellular systems exist that can optimally encode these signals for prediction.

Finally, while we have analyzed the design of cellular signaling networks to optimally predict future signals, we have not addressed the utility of information for function or behavior. It is clear that many functional or behavioral tasks, like chemotaxis [18], require information, but what the relevant bits of information are is poorly understood [7]. Moreover, cells ultimately employ their resources—protein copies, time, and energy—for function or behavior, not for processing information per se. Here, we have shown that maximizing predictive information under a resource constraint,  $C \rightarrow I_{\text{past}} \rightarrow I_{\text{pred}}$ , does not necessarily imply maximizing past information. This hints that optimizing a functional or behavioral task under a resource constraint,  $C \rightarrow I_{\text{pred}} \rightarrow \text{function}$ , may not imply maximizing the predictive information necessary to carry out this task.

## ACKNOWLEDGMENTS

We thank Jenny Poulton, Manuel Reinhardt, Michael Vennetilli and Daan de Groot for many useful discussions. This work is part of the Dutch Research Council (NWO) and was performed at the research institute AMOLF. This project has received funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation program (grant agreement No. 885065).

---

- [1] M. Monti, D. K. Lubensky, and P. R. ten Wolde, Robustness of Clocks to Input Noise, *Physical Review Letters* **121**, 078101 (2018), publisher: American Physical Society.
- [2] W. Pittayakanchit, Z. Lu, J. Chew, M. J. Rust, and A. Murugan, Biophysical clocks face a trade-off between internal and external noise resistance., *eLife* **7** (2018).
- [3] E. Kussell and S. Leibler, Phenotypic Diversity, Population Growth, and Information in Fluctuating Environments, *Science* **10.1126/science.1114383** (2005), publisher: American Association for the Advancement of Science.
- [4] I. Tagkopoulos, Y.-C. Liu, and S. Tavazoie, Predictive Behavior Within Microbial Genetic Networks, *Science* **10.1126/science.1154456** (2008), publisher: American Association for the Advancement of Science.
- [5] A. Mitchell, G. H. Romano, B. Groisman, A. Yona, E. Dekel, M. Kupiec, O. Dahan, and Y. Pilpel, Adaptive

prediction of environmental changes by microorganisms, *Nature* **460**, 220 (2009), number: 7252 Publisher: Nature Publishing Group.

- [6] W. Bialek, *Biophysics: searching for principles*, edited by Princeton Editorial Associates Inc. (Princeton University Press, Woodstock, Oxfordshire, 2012) pages: 563.
- [7] N. B. Becker, A. Mugler, and P. R. Ten Wolde, Optimal Prediction by Cellular Signaling Networks, *Physical Review Letters* **115**, 1 (2015).
- [8] N. Tishby, F. C. Pereira, and W. Bialek, The information bottleneck method, *Proceedings of 37th Allerton Conference on communication and computation* (1999).
- [9] V. Sachdeva, T. Mora, A. M. Walczak, and S. E. Palmer, Optimal prediction with resource constraints using the information bottleneck, *PLOS Computational Biology* **17**, e1008743 (2021), publisher: Public Library of Science.
- [10] C. C. Govern and P. R. ten Wolde, Optimal resource allocation in cellular sensing systems, *Proceedings of the National Academy of Sciences* **111**, 17486 LP (2014).
- [11] G. Malaguti and P. R. ten Wolde, Theory for the optimal detection of time-varying signals in cellular sensing systems, *eLife* **10**, e62574 (2021).
- [12] G. Chechik, A. Globerson, N. Tishby, and Y. Weiss, Information Bottleneck for Gaussian Variables, *Journal of Machine Learning Research* **6**, 165 (2005).
- [13] A. Goldbeter and D. E. Koshland, An amplified sensitivity arising from covalent modification in biological systems, *Proceedings of the National Academy of Sciences* **78**, 6840 (1981), publisher: National Academy of Sciences Section: Research Article.
- [14] U. Alon, *An Introduction to Systems Biology: Design Principles of Biological Circuits* (Chapman and Hall/CRC, New York, 2006).
- [15] J. E. Segall, S. M. Block, and H. C. Berg, Temporal comparisons in bacterial chemotaxis., *Proceedings of the National Academy of Sciences of the United States of America* **83**, 8987 (1986), publisher: Proc Natl Acad Sci U S A.
- [16] C. C. Govern and P. R. ten Wolde, Optimal resource allocation in cellular sensing systems, *Proceedings of the National Academy of Sciences of the United States of America* **111**, 17486 (2014).
- [17] M. Hinczewski and D. Thirumalai, Cellular Signaling Networks Function as Generalized Wiener-Kolmogorov Filters to Suppress Noise, *Physical Review X* **4**, 3 (2014).
- [18] H. H. Mattingly, K. Kamino, B. B. Machta, and T. Emonet, Escherichia coli chemotaxis is information limited, *Nature Physics* 2021 17:12 **17**, 1426 (2021), publisher: Nature Publishing Group.
- [19] T.-L. Wang, B. Kuznets-Speck, J. Broderick, and M. Hinczewski, The price of a bit: energetic costs and the evolution of cellular signaling, *bioRxiv* , 2020.10.06.322770 (2022).
- [20] S. Tanase-Nicola, P. B. Warren, and P. R. ten Wolde, Signal detection, modularity, and the correlation between extrinsic and intrinsic noise in biochemical networks, *Phys. Rev. Lett.* **97**, 068102 (2006).
- [21] E. Ziv, I. Nemenman, and C. H. Wiggins, Optimal signal processing in small stochastic biochemical networks, *PLoS One* **2**, 10.1371 (2007).
- [22] W. De Ronde, F. Tostevin, and P. R. ten Wolde, Effect of feedback on the fidelity of information transmission of time-varying signals, *Phys. Rev. E* **82**, 031914 (2010).
- [23] T. E. Ouldridge, C. C. Govern, and P. R. ten Wolde, Thermodynamics of computational copying in biochemical systems, *Physical Review X* **7**, 10.1103/PhysRevX.7.021004 (2017).
- [24] J. E. Keymer, R. G. Endres, M. Skoge, Y. Meir, and N. S. Wingreen, Chemosensing in Escherichia coli: Two regimes of two-state receptors, *Proceedings of the National Academy of Sciences* **103**, 1786 (2006), publisher: National Academy of Sciences.
- [25] N. Barkai and S. Leibler, Robustness in simple biochemical networks, *Nature* **387**, 913 (1997).
- [26] R. G. Endres and N. S. Wingreen, Precise adaptation in bacterial chemotaxis through “assistance neighborhoods”, *Proceedings of the National Academy of Sciences* **103**, 13040 (2006).
- [27] G. Lan, P. Sartori, S. Neumann, V. Sourjik, and Y. Tu, The energy-speed-accuracy trade-off in sensory adaptation, *Nature Physics* **8**, 422 (2012).
- [28] P. Sartori and Y. Tu, Free Energy Cost of Reducing Noise while Maintaining a High Sensitivity, *Physical Review Letters* **115**, 118102 (2015), 1505.07413.
- [29] T. S. Shimizu, Y. Tu, and H. C. Berg, A modular gradient-sensing network for chemotaxis in Escherichia coli revealed by responses to time-varying stimuli, *Molecular Systems Biology* **6**, 1 (2010), publisher: Nature Publishing Group.
- [30] Y. Tu, T. S. Shimizu, and H. C. Berg, Modeling the chemotactic response of Escherichia coli to time-varying stimuli, *Proceedings of the National Academy of Sciences of the United States of America* **105**, 14855 (2008).
- [31] F. Tostevin and P. R. Ten Wolde, Mutual information between input and output trajectories of biochemical networks, *Physical Review Letters* **102**, 1 (2009).
- [32] M. Reinhardt, G. Tkačik, and P. R. t. Wolde, Path Weight Sampling: Exact Monte Carlo Computation of the Mutual Information between Stochastic Trajectories, *arXiv* 10.48550/arxiv.2203.03461 (2022), 2203.03461.
- [33] J. Monod, J. Wyman, and J.-P. Changeux, On the nature of allosteric transitions: A plausible model, *Journal of Molecular Biology* **12**, 88 (1965).
- [34] V. Sourjik and H. C. Berg, Binding of the Escherichia coli response regulator CheY to its target measured in vivo by fluorescence resonance energy transfer, *Proceedings of the National Academy of Sciences of the United States of America* **99**, 12669 (2002), publisher: National Academy of Sciences.
- [35] M. Li and G. L. Hazelbauer, Cellular stoichiometry of the components of the chemotaxis signaling complex, *Journal of Bacteriology* **186**, 3687 (2004), publisher: J Bacteriol.
- [36] G. Malaguti and P. R. T. Wolde, Theory for the optimal detection of time-varying signals in cellular sensing systems, *eLife* **10**, 1 (2021), publisher: eLife Sciences Publications Ltd.
- [37] P. Sartori and Y. Tu, Free Energy Cost of Reducing Noise while Maintaining a High Sensitivity, *Physical Review Letters* **115**, 118102 (2015), publisher: American Physical Society.
- [38] J. O. Dubuis, G. Tkacik, E. F. Wieschaus, T. Gregor, and W. Bialek, Positional information, in bits, *Proceedings of the National Academy of Sciences of the United States of America* **110**, 16301 (2013).