

Information processing in living systems

Gašper Tkačik^{1*} and William Bialek^{2†}

¹*Institute of Science and Technology Austria, Am Campus 1, A-3400 Klosterneuburg, Austria*

²*Joseph Henry Laboratories of Physics and Lewis-Sigler Institute for Integrative Genomics,
Princeton University, NJ 08544 Princeton, USA*

(Dated: December 31, 2014)

Life depends as much on the flow of information as on the flow of energy. Here we review the many efforts to make this intuition precise. Starting with the building blocks of information theory, we explore examples where it has been possible to measure, directly, the flow of information in biological networks, or more generally where information theoretic ideas have been used to guide the analysis of experiments. Systems of interest range from single molecules (the sequence diversity in families of proteins) to groups of organisms (the distribution of velocities in flocks of birds), and all scales in between. Many of these analyses are motivated by the idea that biological systems may have evolved to optimize the gathering and representation of information, and we review the experimental evidence for this optimization, again across a wide range of scales.

I. INTRODUCTION

Cells and organisms sense, compute, and make decisions: to proliferate, to find food, to protect themselves against predators and unfavorable environmental changes, to act in unison with their neighbors within a collective, and—broadly speaking—to invest their limited resources to their maximal benefit. These processes span a wide range of temporal and spatial scales, and reflect dynamics in phase spaces of widely varying intrinsic dimensionality. In bacteria, for example, the chemotactic signaling network comprises just a handful of chemical species, responds to changes in the local nutrient concentration on the one second timescale, and modulates bacterium’s swimming by controlling the direction of flagellar rotation. At the other extreme of complexity, millions of neurons in the brain of a pianist are engaged in processing roughly one megabit per second of visual, auditory, and tactile information, combining these data with the score retrieved from memory, and computing the fine motor program controlling the pianist’s fingers. Other processes take place over much longer timescales: DNA sequences in a population of organisms are subject to the evolutionary forces of mutation, selection, and random drift, and changes in the environment slowly shape the distribution of genomic sequences. In all these cases, intuition tells us that information is “flowing” from the outside environment to a representation internal to the organism or a population, and that this information flow is essential for life. Can this intuition be formalized? If so, can it be connected, quantitatively, to experimental data? Most ambitiously, can a focus on information flow help us to unify our description of the complex and diverse phenomena that we see in the living world?

The purpose of this review is to explain how the intuition about “biological information flows” can be formal-

ized in the language of information theory, and what the resulting research program teaches us about the physics of living systems. We introduce the basic concepts of information theory (Section II), show how this mathematical structure leads to new views of experimental data (Section III), and explore the possibility that some of life’s mechanisms can be understood as solutions to the problem of optimizing information flow subject to physical constraints (Section IV). Throughout, we follow the physics tradition of emphasizing the unity of theoretical ideas at the expense of some of the biological complexity.

Before we begin, it seems worth discussing how ideas about information flow relate to the more familiar theoretical structures in physics. Intuitions about information flow are not confined to biological systems. When we look a ferromagnet, for example, we know that all the details of the electronic configurations on one side of the sample are not relevant if we try to predict those on the other side of the sample; what does carry information across long distances is the order parameter, in this case the magnetization. But precisely because we can identify the order parameter, we don’t need to speak abstractly about information. Rather than trying to turn “carry information” into a precise mathematical statement, we just compute correlation functions of the order parameter evaluated at distant points.

For most of the biological systems that we will be discussing, finding something like an order parameter remains a distant hope. Symmetries are absent, and even locality is not much of a guide to understanding, for example, a network of neurons in which each cell may be connected to many thousands of neighbors. Under these conditions, some more general approaches to formalizing our intuitions about information flow would be useful. Interestingly, as we write this, information theoretic ideas are becoming more important in the study of topological phases in correlated electron systems, where no local order parameter is possible.

*gtkacik@ist.ac.at

†wbialek@princeton.edu

II. BUILDING BLOCKS OF INFORMATION THEORY

Here we give a brief exposition of information theory, focusing on topics that are most relevant for our subsequent discussion; see also Ref [1]. The foundational papers by Shannon are wonderfully readable [2], there is an excellent textbook account aimed at the core of the subject [3], and a more recent text that emphasized connections between ideas information theory and statistical physics [4]. Ideas which have been most widely used in thinking about biological systems are not necessarily those that have been most useful in the (enormously successful) application of information theory to communication technology, so our perspective here is somewhat different from that in Ref [3].

When we ask a question, and hear the answer, we gain information. In 1948 Shannon asked whether we could attach a number to the amount of information we gain [2]. He imagined that, upon asking the question, we could envision a set of possible answers ($n = 1, 2, \dots, N$) and that we could assign probabilities $\mathbf{p} \equiv \{p_1, p_2, \dots, p_N\}$ to each of these answers. Any acceptable measure of information must obey some constraints: if all N possible answers are equally likely, then the information must grow with N ; if the question can be decomposed into independent parts, then the information gained on hearing the answer to each part should add to the total; if we can decompose the full question into a tree of choices (as we do when playing twenty questions), then the total information should be the weighted sum along the paths through the tree. Remarkably, these postulates allow only one possible measure of information, the entropy of the distribution of answers,¹ $S[\mathbf{p}] = -k \sum_n p_n \log p_n$.

A seemingly very different question concerns the amount of space required to write down, or represent the answer to our question. While each possible answer might require a different amount of space, if this is a question we can ask many times (what will the temperature be at noon tomorrow?) then it makes sense to ask about the minimum amount of space required, per answer, in the limit of a long sequence of question/answer pairs. The answer is again the entropy. There is some ambiguity about units—the number of characters we need to write the answer depends on our choice of alphabet, and the entropy itself has an arbitrary constant k —but if we choose, for example, a binary alphabet, we can set $k = 1$ and $\log \equiv \log_2$, then all of the ambiguity is resolved. When entropy or information is measured in these units, they are called ‘bits.’

Getting the precise answer to a question is a rare thing.

More common is that we are interested in x but can only measure y . Before a measurement we know only that x is drawn from the distribution $P_X(x)$; after the measurement we know that it is drawn from the conditional distribution $P(x|y)$. The same arguments that establish the uniqueness of the entropy now show that the only acceptable measure of the information about x gained by observing y is the reduction in entropy,

$$I(y \rightarrow x) = - \sum_x P_X(x) \log_2 P_X(x) - \left[- \sum_x P(x|y) \log_2 P(x|y) \right]. \quad (1)$$

If we ask not about the information gained from observing a particular y , but the average over all the y 's we can observe, this average information is actually symmetric in x and y , and is usually called the “mutual information.” It can be written in a manifestly symmetric form, and unlike the entropy itself it is well defined even when x and y are continuous variables,

$$I(x; y) = \iint dx dy P(x, y) \log_2 \left[\frac{P(x, y)}{P_X(x)P_Y(y)} \right], \quad (2)$$

where $P(x, y)$ is the joint distribution of x and y . This can be extended to the case where x and y are multi-dimensional, or even functions, in which case the sums over states become functional integrals as in quantum mechanics or field theory.

One way to think about the mutual information is as a measure of dependency or correlation. Clearly we have $I(x; y) = 0$ if the variables x and y are independent, so that $P(x, y) = P_X(x)P_Y(y)$, and Shannon’s arguments tell us that $I(x; y)$ is the only measure of correlation that satisfies a set of sensible requirements. If x and y are jointly Gaussian, then

$$I(x; y) = -\frac{1}{2} \log_2(1 - c^2), \quad (3)$$

where c is the usual linear correlation coefficient. But it is easy to imagine relations between x and y that are highly informative but not linear. Even if the underlying variables are linearly correlated, the things we measure might be nonlinearly transformed version of these variables, and in this case computing the mutual information becomes an especially powerful way of analyzing the data; see Section IIID below.

The more conventional way of thinking about the mutual information is in terms of a communication channel. There is an input message x , and it is sent along a channel that produces the output y . The fact that y is chosen from a probability distribution $[P(y|x) = P(x|y)P_Y(y)/P_X(x)]$ means that the communication is noisy, and this must limit the amount of information that can be transmitted. If we think of the mutual information as being a functional of the distributions $P_X(x)$ and $P(y|x)$, then the convexity of the logarithm means that

¹ Shannon’s theorem makes precise the vague statement made in many statistical mechanics classes, that entropy is related to our lack of information about the microscopic state of a system. The more positive version is that the entropy is the information we would gain upon learning the full microscopic state.

there is a maximum as functional of the distribution of inputs; this maximum is called the channel capacity,² and is the limit to information transmission set by noise in the channel.³

We note that actually finding the capacity of a channel, given a detailed model of the input/output relation and noise as represented by $P(y|x)$, can be challenging. Simple models, such as y being a linearly filtered version of x with added Gaussian noise, are tractable, and generalize in a straightforward fashion to cases where the signals are varying in time. For this reason, Gaussian channels have dominated the discussion of both biological and engineered systems, and the capacity-achieving strategies that depend on the system’s input and output noise levels are well-understood [1, 3, 6]. Beyond the Gaussian channel, progress has been made using various approximative schemes, such as by assuming that noise is small compared to the dynamic range of the signal (the small-noise approximation), but there is certainly room for further theoretical developments using biophysically-informed noise models.

Examples of information transmission through a noisy channel are everywhere in biological systems. Most obviously, our eyes and ears take inputs from the outside world and the neurons emerging from our sense organs generate streams of identical electrical pulses, termed action potentials or spikes. Even single celled organisms have sensors, as with the bacterial chemotaxis system, which responds to changing concentrations of molecules in the environment and generates at the output changes in the direction of swimming. There are also purely internal examples, as when a cell synthesizes “transcription factor” proteins that bind to DNA and influence the rate at which encoded information is read out to make other proteins; here the transcription factor concentration is the input and the resulting protein concentration is the output. In some of those cases variables are naturally continuous (sound pressure), in other cases they are naturally discrete (counting individual molecules), and in many cases it is useful to pass between descriptions that are either discrete or continuous (e.g. concentration instead of counting molecules); it is attractive that we can discuss information transmission in all these cases, in the same language and the same units.

² It is a remarkable fact that one can transmit information *without error* along a channel that has a non-zero noise level, provided that the rate of transmission does not exceed this capacity. Strictly speaking “capacity” usually is used to describe this maximum rate of errorless transmission, rather than the maximum of the mutual information between input and output [3].

³ In engineered systems, much emphasis is placed on devising coding schemes, i.e., algorithms that transform inputs x into messages to be sent through the channel (and likewise recover x from the channel output), so that information can be transmitted over noisy channels as close to capacity as possible using bounded processing resources. While interesting, coding questions have remained largely unexplored in biological systems.

Information is a difference in entropies, as in Eq (1), and because information is mutual we can write this in two ways,

$$\begin{aligned} I(x; y) &= S[P_Y(y)] - \langle S[P(y|x)] \rangle_{P_X(x)} & (4) \\ &= S[P_X(x)] - \langle S[P(x|y)] \rangle_{P_Y(y)}, & (5) \end{aligned}$$

where $\langle \cdot \rangle_{p(x)}$ denotes an average over $p(x)$. If either the input or the output is discrete, then both the entropy and the conditional entropy are positive semi-definite, and hence the entropy sets an upper bound on the information. Further, if we can assign a cost to every output state y —e.g., because generating more action potentials costs more energy, then there is a maximum entropy consistent with a given average cost (for more on maximum entropy see Section III C). Thus there is a rigorous path from limits on physical resources to limits on information transmission.

Another way of bounding information is through the data processing inequality. If the output y receives information about the input x only through some intermediate z , we can write

$$P(y|x) = \sum_z P_{\text{II}}(y|z)P_{\text{I}}(z|x). \quad (6)$$

Now there are two communication channels, first $x \rightarrow z$ (channel I) and then $z \rightarrow y$ (channel II). Each channel is associated with a mutual information, $I_{\text{I}}(x; z)$ and $I_{\text{II}}(z; y)$, and each of these is an upper bound on the information transmission through the combined channel, $I(x; y)$. Another way of saying this is that any processing from z to y can only result in a loss of information, so that $I(x; y) \leq I_{\text{I}}(x; z)$. In this sense, information can be destroyed, but not created. This result is conceptually important, and we will see that it is of practical use in analyzing experiments (Section III D). As for the mutual information, so too for the channel capacities: to support mutual information $I(x; y)$, between x and y , the capacities of both channels I and II must be larger than $I(x; y)$.

The universality of information-theoretic quantities is their strength, but it also leaves one with an uncomfortable feeling that there is more to the story of biological signal processing than bits. In particular, not all bits may be of equal value or utility to the organism. Information processing usually has a goal beyond communication, and this seems especially true in biological systems. Perhaps because Shannon called his work a mathematical theory of communication, it is not as widely appreciated that information theory has anything to say about such goal-defined tasks. The key idea is rate-distortion theory, which again traces back to Shannon [5, 6].

As an example, consider a single celled organism in an environment where various nutrients are present at concentrations $\mathbf{c} = \{c_1, c_2, \dots, c_N\}$, and over time this environmental state varies. To metabolize these molecules, the cell synthesizes e_1 molecules of enzyme 1, e_2 copies of enzyme 2, and so on up to enzyme K ; to fully convert the inputs into useful molecules, there may need to

be more enzymes than nutrients. The growth rate of the cell depends on both the environment and the state of the cell, $\lambda(\mathbf{c}, \mathbf{e})$. Under any particular environmental condition \mathbf{c} , there is an optimal setting of the internal state \mathbf{e} that maximizes the growth rate. If the cell finds this optimum exactly, then by looking at \mathbf{e} we would know exactly the environmental variable \mathbf{c} , and this exact knowledge of real variables corresponds to an infinite amount of information. At the opposite extreme, the cell could choose the copy numbers of various enzymes at random, and hope that, averaged over the distribution of environments, the average growth rate is not so bad; in this case the internal state of the cell carries zero information about the environment. But neither infinite nor zero information seems plausible; what happens in between? Rate–distortion theory tells us that if the cell needs to achieve a given average growth rate, then there is a minimum required mutual information between the state of the environment and the internal state of the cell, $I(\mathbf{c}, \mathbf{e}) = I_{\min}(\langle \lambda \rangle)$.

More generally, if we have inputs x and outputs y , then given some measure of performance $\mathcal{U}(x, y)$, we can plot the average performance vs. the mutual information $I(x; y)$. The rate–distortion curve $I_{\min}(\langle \mathcal{U} \rangle)$ divides this plane into accessible and inaccessible regions. Given a certain number of bits, there is a maximal achievable performance, and conversely achieving any given level of performance requires a minimum number of bits. Thus, although collecting and transmitting bits cannot be the goal of an organism, any real goal requires a minimum number of bits. If bits are cheap, then this is a true but not very useful observation. If, on the other hand, the limited resources available to cells and organisms mean that more bits are hard to get, then selection for better performance will translate into selection for mechanisms that gather and process more bits with the same resources.

Rate distortion is a framework that formalizes lossy data compression: signals x can be sent through the channel to yield outputs y which are then used to find the best achievable reconstruction $\hat{x}(y)$ of the input signal. The metric for the reconstruction quality is the distortion function, for example, the RMS error between the true signal x and its reconstruction \hat{x} . In this sense, rate distortion tells us the minimum number of bits to recover the input signal to within a desired level of quality. When thinking about biological systems, it is often convenient to turn this reasoning around and explicitly construct various decoding or readout schemes that reconstruct the input signals from the activity of the biological network. Classic examples involve decoding of motion trajectories from the spiking activity of neurons in the visual system [7, 8], or the decoding of hand trajectories from the activity of neurons in motor cortex during drawing [9]; for more on decoding of neural activity see Ref [10]. Decoding schemes can be used to put a lower bound on information transmission (Section III) and might further suggest how much of the total information is encoded ex-

PLICITLY, in a way that can be extracted using biologically plausible mechanisms. Indeed, the relationship between the total information and the component which is accessible in simple form is an issue in neural coding that needs to be explored more fully.

III. INFORMATION AND DATA ANALYSIS

In complex systems, we may be unsure which features of the system dynamics are the most relevant. Even if we can identify the correct features, we might not know the biologically meaningful metric that measures the similarity or difference along these relevant dimensions. Information theory gives us a way out of these difficulties, because it singles out measures that are completely general: while some particular correlation function of X and Y might vanish, if these variables are related in any way then there must be a nonzero mutual information $I(X; Y)$; further, there is an obvious sense in which the stronger the relationship, the greater the mutual information. Against this generality is the difficulty of actually estimating the entropy of a probability distribution from a limited set of samples.

A. Quantifying real information flows

In thermodynamics, entropy changes are related to heat flows, and this can be used to measure the entropy in equilibrium systems. Away from equilibrium statistical mechanics, we have no such general method for measuring entropy. Indeed, the entropy depends on the entire probability distribution, and so “measuring” the entropy is quite unlike measuring the usual observables.

We can estimate the entropy of a probability distribution by counting how often each state occurs, using this frequency as an estimate of the probability, and plugging into the definition $S = -\sum_n p_n \log_2 p_n$. While the errors in estimating the individual probabilities p_n are random, the convexity of the logarithm means that these random errors contribute to a systematic error that declines as $1/\sqrt{N_s}$, where N_s is the number of samples. The difficulty is that the coefficient of this systematic error is proportional to the number of possible states of the system, and hence accurate estimates require many more samples than states. The problem becomes worse when we try to estimate the mutual information between two variables, since the number of possible states is the product of the number of states for each variable taken separately.

The problem of systematic errors in estimating entropy and information was appreciated very early, in the 1950s, as information theoretic ideas began to be applied to biological problems [11]. Motivated in large part by efforts to measure the information that neural responses carry about sensory inputs, there has been a substantial effort to calculate and correct for these systematic errors [12–14]. A different direction is to note that the single num-

ber, entropy, might be well determined even if the full distribution is not. This idea can be made concrete in a Bayesian approach, placing a prior on the space of probability distributions; priors that are uniform in entropy, as opposed to uniform in probability, seem especially effective [15–18]. In many cases, the variables of interest are not naturally discrete, and so the problem of entropy estimation is tangled with that of density estimation [19]. Although much progress has been made, the difficulty of estimating entropy and information from finite samples leads to periodic inventions of other measures that might serve similar purposes (e.g., quantifying dependency beyond linear correlation); for a recent reminder of why mutual information is not just one measure among many, see Ref [20].

An important ingredient in estimating entropy and information from real data has been the use of bounds. It may be difficult to estimate the entire probability distribution, but it often is feasible to measure, with small error bars, several average quantities (e.g., moments), and construct the probability distribution that matches these averages exactly but has the maximum possible entropy (Section III C). The data processing inequality means, for example, that if we want to measure the mutual information between x and y , and y is of high dimensionality, any projection into a lower dimensional space, $y \rightarrow z$, can only result in a loss of information, $I(z; x) \leq I(y; x)$. Reducing dimensionality makes the problem of estimating distributions much easier, and thereby eases the sampling problems of entropy estimation; since the lower dimensional variable provides a lower bound to the true information, it makes sense to search for optimal projections (Section III D). Finally, as emphasized long ago by Ma [21], the probability that two randomly chosen states are the same, $P_c = \sum_n p_n^2$, provides a lower bound to the entropy, $S \geq -\log_2 P_c$, that can be estimated reliably from relatively few samples. For more details, see Appendix A.8 of Ref [1].

A modest though widespread use of information theory in the analysis of biological data concerns the description of regulatory sites along DNA. The readout of information encoded in the genome—the synthesis of proteins—is regulated in part by the binding of “transcription factor” proteins (TFs) to these sites, and the strength with which one TF binds depends on the DNA sequence at the binding site. In this sense, the local sequences can be thought of as “addresses” that distinguish certain genes from the rest, but it is known that transcription factors do not do exact string matching; instead, they bind to an ensemble of sites that permit some variation. Early work used simple arguments to compute how much “information” is contained in these binding site ensembles, i.e., as the difference between the entropy of random sequences and the difference of the ensemble that the TF will bind [22], and this has given rise to a popular graphical way of depicting TF binding site specificity in the form of “sequence logos” [23]. Subsequent work has refined these arguments and asked whether the information in the binding sites

is sufficient to address specific genes in genomes of pro- and eukaryotes [24].

Meaningful quantitative statements about information flow obviously depend on having high quality data. In the nervous system, the fact that the relevant output signals are electrical means that quantitative measurements have been the norm since the invention of the vacuum tube [25]. Steady improvements in technology now make it possible to collect sufficient data for information theoretic analyses in almost any neural system, even in freely behaving animals; in practice, much of the motivation for these analyses has come from ideas about the optimization of the neural code, and so these measurements are reviewed below (Section IV.B). In biochemical and genetic networks, the fact that outputs are the concentrations of molecules means that there is no universal measurement tool.

Classical methods for measuring the concentrations of signaling and control molecules inside cells include raising antibodies against particular molecules of interest, and then tagging antibodies against these antibodies with fluorescent labels. Recent work in the fly embryo shows that, with care, this classical technique can be pushed to give measurements of concentration that are accurate to $\sim 3\%$ of the relevant dynamic range [26]. More modern methods involve genetically engineering the organism to produce fusion of the protein of interest with a fluorescent protein [27], and here too there has been progress in demonstrating that the fusion proteins can replace the function of the native protein, quantitatively [28, 29]. Finally, rather than monitoring protein concentrations, one can count the messenger RNA molecules (mRNA) that are an intermediate step in protein synthesis; recent work has pushed to the point of single molecule sensitivity, so that counting is literal [30, 31]. Similarly, progress has been made in quantifying the activity of intracellular signaling networks, and subsequently using information theoretic approaches to compare how signals propagate in different network architectures [32], or how they are encoded in the dynamical properties of the response [33].

B. Quantifying correlations

In the last decade it has become possible to perform simultaneous recordings from tens or hundreds of nodes in interacting biological networks. Gene expression arrays and sequencing technology have enabled such readouts for gene expression patterns; cheap sequencing can map the distributions of gene variants in populations; multi-electrode arrays and lately optical imaging have pushed the boundaries in neuroscience to the point where full temporal activity traces for every neuron in small brains—a zebrafish larva, the worm *C. elegans*, or the fruit fly—don’t seem out of reach; and in the field of collective behavior, simultaneous tracing of positions of hundreds or thousands of individuals can be performed. Because the barriers between subfields of biology are large, each

of these experimental developments has led to an independent discovery of the need for new and more powerful analysis tools. Even the simple question of when the activity at two nodes in the network is correlated is complicated, because it is surely not true that linear correlations capture the full relationships among pairs of variables in these complex systems.

As microarray methods provided the first global pictures of gene expression in cells, several groups tried to analyze the covariation in expression levels simple by computing correlations. With enough data, it became possible to estimate the mutual information between expression levels [34]. These initial analyses revealed that, as expected, linear correlations did not exhaust the relationships between genes; while the most correlated pairs had the highest mutual information, at intermediate or even minimal levels of correlation one could find pairs of genes whose expression levels shared nearly half the mutual information of the most correlated pairs. In addition, it became clear that pairs of expression levels could share more than one bit of mutual information, which contributed to the realization that the control of gene expression involved more than just on/off switches.

Measurements of mutual information also provide a new basis for clustering, in which we try to maximize the average mutual information between the expression levels of pairs of genes in the same cluster [35]. This approach is interesting mathematically because it is invariant to any invertible transformations of the expression levels (and hence to systematic distortions in the measurement), and because the clusters have no “centers” or prototypical members. These information theoretic methods were able to reveal biologically significant structures even inside tightly correlated “modules” where other methods failed. Further developments place the estimates of mutual information among many pairs of gene expression levels at the core of algorithms for reconstructing the underlying network of interactions, where the data processing inequality provides a guide to disentangling direct interactions from more widespread correlations [36].

One can look at interactions between the genes in a very different way, analyzing joint presence or absence of particular genes in the sequences of many organisms. Here, each “sample,” i.e., one organism’s genome, is represented as a long binary vector, where 1 represents a presence of a particular gene and 0 its absence.⁴

⁴ Some care is needed here. Two different organisms almost never have genes with exactly the same sequences, but if we are comparing organisms that are not too far separated in evolutionary distance, then finding the corresponding or “homologous” genes is relatively straightforward. It is more difficult to be certain of the historical path that generated these homologs: descent from a common ancestor (orthologs’), convergent evolution, gene duplication (paralogs), For more about sequence similarity, see Section III C.

Estimating the mutual information between the presence/absence of pairs of genes is technically straightforward, and reveals the organization of genes into clusters [37]; this approach can be extended beyond pairs to probe higher-order interactions among genes. Information theoretic measures can also be used to quantify the correlations between the presence/absence of particular genes and various phenotypes, and these analyses point toward functional modularity in the organization of genomes [38].

C. Systems with many degrees of freedom

One of the great triumphs of the twentieth century is the identification and characterization of the molecular building blocks of life. But life is more than the sum of its parts. From the spectacular aerial displays by flocks of birds down to the coordinated movements of cells in a gastrulating embryo, many of the most striking phenomena in the living world are the result of interactions among hundreds, thousands, or even millions of individual elements. The enormous success of statistical physics in describing emergent phenomena in equilibrium systems has led many people to hope that it could provide a useful language for describing emergence in biological systems as well. Perhaps the best developed set of ideas in this direction concerns the dynamics of neural networks [39–41], but there have also been important developments in the description of collective animal behavior [42–44] and in models for genetic networks. While each case has its own special features, it seems fair to say that, in all these cases, the search for simple, analyzable models often takes us away from contact with the details of measurements on the real biological systems. In extreme cases, it is not clear whether these models are just metaphors (the alignment of birds in a flock is like the alignment of spins in a magnet), or if they are to be taken seriously as theories that make quantitative predictions. Over the past decade, progress in experimental techniques, as noted in the previous section, have held out the promise of much more detailed confrontations between theory and experiment.

Although there are many interesting questions about emergent behavior in biological networks, we focus here on one question where information theoretic ideas have been crucial: if there a variable y_i that lives at each node of the network, can we make a model for the joint distribution of these variables, $P_y(\mathbf{y}) \equiv P_y(y_1, y_2, \dots, y_N)$? Progress on this problem has been made using the idea of maximum entropy [45].

Let us assume that experiments provide T observations on the state of the system, \mathbf{y} . Often, the dimension of the system is so large that T samples are completely insufficient to empirically sample $P_y(\mathbf{y})$. However, we might be able to estimate certain statistics on the data, which we denote by $\langle \hat{O}_\mu \rangle_{\text{expt}}$, where $\mu = 1, \dots, M$. These could be, for example, expectation values of individual network

nodes, $\langle y_i \rangle$, the covariances, $\langle y_i y_j \rangle$, or any other function for which the samples provide a reliable estimate. We look for a model distribution $\hat{P}_y(\mathbf{y})$ that will exactly reproduce the measured expectation values, but will otherwise be as unstructured, or random, as possible—hence maximum entropy. This amounts to solving a variational problem for \hat{P} , for which the solution is:

$$\hat{P}^{(\hat{O}_1, \dots, \hat{O}_M)}(\mathbf{y}) = \frac{1}{Z(\{g_\mu\})} \exp \left[\sum_{\mu=1}^M g_\mu \hat{O}_\mu(\mathbf{y}) \right]. \quad (7)$$

Thus, the maximum entropy distribution has an exponential form, where the coupling constants g_μ have to be chosen so that the constraints $\langle \hat{O}_\mu \rangle_{\text{expt}} = \langle \hat{O}_\mu \rangle_{\hat{P}}$ are exactly satisfied. Equation (7) has the form of a Boltzmann distribution in which every state of the system is assigned an “energy” $E(\mathbf{y}) = \sum_{\mu=1}^M g_\mu \hat{O}_\mu(\mathbf{y})$, and this equivalence to equilibrium statistical mechanics is an important source of intuition and calculational tools. But Eq (7) is *not* a description of a system in thermal equilibrium. If the only thing we measure about a system is its energy, then asking for the maximum entropy distribution consistent with the average energy does indeed correspond to a description of thermal equilibrium, but in most biological systems the quantities that we measure have no relation to the energy, and we are measuring many of them.

The maximum entropy method does not correspond to a single model for the system, but rather provides a systematic strategy for building a hierarchy of models that provide increasingly good approximations to the true distribution. At each step in the hierarchy, the models are parsimonious, having the minimal structure required to reproduce the expectation values that we are trying to match. For every new expectation value that we can measure, the new maximum entropy model provides a tighter upper bound on the entropy of the true distribution, and can uniquely decompose the total amount of correlation among the N nodes in the network into contributions from pairwise, triplet-, \dots , K -fold interactions [46]; the hope is that we arrive at a good approximation while $K \ll N$.

Equation (7) tells us the form of the maximum entropy distribution, but to complete the construction of the model we need to find the coupling constants $\{g_\mu\}$. Usually in statistical mechanics we are given the coupling and try to compute the expectation values or correlation functions ($\{g_\mu\} \rightarrow \{\langle \hat{O}_\mu(\mathbf{y}) \rangle\}$). Here we need to do inverse statistical mechanics, mapping the experimentally measured expectation values back to the coupling constants ($\{\langle \hat{O}_\mu(\mathbf{y}) \rangle\} \rightarrow \{g_\mu\}$). Such inverse problems have a history dating back (at least) to the demonstration by Keller and Zumino [47] that the temperature dependence of the second virial coefficient in a classical gas determines the interaction potential between molecules uniquely, provided that this potential is monotonic; the search for rigorous statements about inverse statistical mechanics has attracted the attention of mathematical

physicists [48, 49], although many questions are open. At a practical level, estimating the coupling constants from measured expectation values can be seen as part of the broader problem of learning probabilistic models from data, and hence part of machine learning in computer science. For a small sampling of the interplay between ideas in statistical physics and machine learning, see Refs [4, 50–53].

Interest in maximum entropy approaches to biological networks was stimulated by results from an analysis of activity patterns in small groups of neurons from the vertebrate retina [54]. As the retina responds to natural movies, pairwise correlations are weak, of mixed signs, and widespread. Models that ignore these weak correlations make predictions about the frequency of patterns in groups of ten neurons that can be wrong by many orders of magnitude, but these errors are largely cured in maximum entropy models that match the pairwise correlations. Even in these small networks, there are hints that what one sees in the network are genuinely collective states. By now, these ideas have been extended and applied to several different neural systems [55–65]. Most recently it has been possible to build models for the joint activity of more than 100 neurons in the retina, and these models are so accurate that they can predict distribution of the effective energy deep into the tail of patterns that occur only once in a two hour long experiment [64]. The states that dominate the distribution are indeed collective, so one can reconstruct the time dependent response of single cells from the behavior of the population, without reference to the visual stimulus. Most provocatively, that distribution that we reconstruct from the measured expectation values exhibits many of the signature of an equilibrium statistical mechanics problem near its critical point [66].

An independent stream of work has used maximum entropy ideas to think about the ensemble of amino acid sequences that form a family of proteins, folding into essentially the same three-dimensional structures. This work was inspired by experiments that actively explored sequence space, trying to respect pairwise correlations among amino acids at different sites [67, 68], using an algorithm that is equivalent in some limits to maximum entropy [69]. Although these correlations extend over long distances, the effective interactions among amino acids in these models were found to be spatially local; this raises the possibility that we can infer the spatial neighbor relations among amino acids from data on sequence variation, in effect folding the protein via sequence comparisons [70–73]. This new approach to the protein folding problem has generated considerable excitement, but there are many open questions about the sufficiency of the available data and the approximations that have been used to solve the inverse problem. Note that this effort to describe the distribution of sequences consistent with a given protein structure is the inverse of the usual protein folding problem, where we are given the sequence and asked to predict the structure, and there is a big con-

ceptual question about the relation between these two problems.

In the immune system we can look at special cases of sequences ensembles, the distribution of antibody molecules that organisms use to combat infection and the distribution of viral sequences in patients that have long term infections such as HIV. In the problem of antibody diversity, maximum entropy models built from pairwise correlations provide a surprisingly accurate description of the distribution as a whole [74], while for HIV the effective energy in maximum entropy models predicts viral fitness and vulnerability to immune attack [75]. This is a small window into a much larger literature on statistical physics approaches to the immune system.⁵

Birds in a flock are thought to follow only their near neighbors, yet the entire flock can decide to move in a single direction and at nearly uniform speed. Order spreads through the flock much as through spins in a magnet, and this is more than a metaphor: maximum entropy models consistent with measured local correlations in a flock provide successful quantitative, parameter-free predictions of correlations in flight direction and speed throughout the flock [76, 77]. It is also possible to do this analysis in more detail, building maximum entropy models that match the correlations between a bird and its nearest neighbor, its second neighbor, and so on; one finds that the resulting couplings which link each bird to its k^{th} neighbor decline exponentially with k , providing some of the most direct support yet for the hypothesis that interactions are local [78]. Real starling flocks exhibit long-ranged correlations of the fluctuations in flight direction and speed [79]; analysis of the maximum entropy models shows that the directional correlations arise from Goldstone modes associated with the spontaneous polarization of the flock [76], while the speed correlations are long-ranged only because the parameters of the model are tuned close to a critical point [77]; signatures of criticality have also been found in the behavior of insect swarms [80].

D. Information-theoretic parameter inference

A standard tool for fitting parametric models to data is maximum likelihood inference. This procedure, however, depends on our ability to write down the likelihood function: the probability of making specific experimental observations given the parameters of the model. In studying biological systems we often are faced with situations where this is impossible, for two very distinct reasons. First, while we might have a good idea for how a part of the system works (e.g., that transcription factors recognize and bind DNA sequences in a way that is gov-

erned by equilibrium statistical mechanics), what we can measure is far “downstream” of these events (the activation level of a gene), and we may not have a quantitative model for the intervening steps. Second, and less fundamentally, many experiments still are not as calibrated as one might hope, so we don’t have a model for the noise in the measurement itself.

Without the ability to write down the likelihood function, is there any way to do unbiased parametric inference of the process of interest? There isn’t much we can do about the general version of this problem, but in many biological systems one crucial step in the process that we are trying to describe involves a substantial reduction of dimensionality. Thus, in the visual system, we expect that neurons respond not to every possible detail of the image or movie falling onto the retina, but only to some limited set of features; in the simplest case these features might be described simply as linear projections onto a template or “receptive field.” If this picture is right, then the transformation from input movie x to output spikes z can be written as $x \rightarrow y \rightarrow z$, where x is a vector in a space of very large dimensionality (the space of movies), while both y and z are low dimensional, in the simplest case just scalars. The mapping $y \rightarrow z$ then is easy to characterize; the problem is how to choose the projection that describes $x \rightarrow y$. The data processing inequality tells us that the best possible predictions of z will be generated by a model in which the mapping $x \rightarrow y$ captures as much information as possible about z . Importantly, the difficulty of searching for such “maximally informative projections” does not vary much with the structure of the distribution of the inputs x , which means that we can use this approach to characterize the encoding of complex, naturalistic input signals [81]. Applications and extensions of this idea have been explored in several different systems [82–86].

The idea of searching for maximally informative projections has also been used to analyze experiments on the interactions of transcription factors with their target sites along DNA; here the projection is from the local DNA sequence to the binding energy of the TF [87]. An important success of this initial effort was to show that two independent experiments which characterized the same transcription factor in fact led to consistent models of TF–DNA binding, despite earlier claims to the contrary. This work also demonstrated the formal equivalence between information-theoretic parameter inference and “error-model-averaged” maximum likelihood inference. Subsequent experiments exploited these idea to give a much more detailed view of how a transcription factor and the RNA polymerase bind and interact at the lac operon [88], and related methods were developed to allow analysis of a wider range of input data [89].

Because the issues of model inference for processes embedded into unknown larger systems or subject to unknown noise processes are prevalent in biology, we expect to see further applications for information theoretic inference in biology. We also note that these applica-

⁵ See, for example, the program at the Kavli Institute for Theoretical Physics, <http://immuno-m12.wikispaces.com>.

tions can be seen as specific examples of a more general framework that relates learning (including parameter inference) to information flow between the data and its compressed representations. There are many interesting links between the learning problem, statistical mechanics, and field theory, which are beyond the scope of this review.

IV. OPTIMIZING INFORMATION FLOW

Information is essential for life, but bits are not free. Is it then possible that organisms are driven to extract and represent the maximum possible information given the physical constraints? This is an idea that arose almost immediately after Shannon’s original work, and has been a productive source for thinking about many specific biological systems, across many scales.

A. The genetic code

Five years after Shannon’s papers, Watson and Crick proposed the double helical structure for DNA, and after another five years Crick could articulate “the sequence hypothesis”: the sequence of amino acids in proteins is determined by the sequence of bases along the DNA [90]. Once one has this idea, the mapping between two major classes of biopolymers becomes a problem of coding, and it seemed natural to bring the new ideas of information theory to bear on this central problem of life [91]. In particular, it was hoped that the code might be efficient in the sense defined by Shannon, using as few bases as possible to encode the amino acid sequence of a protein. As we now know, the genetic code maps triplets of bases (codons) into one of twenty amino acids (plus “stop”), so that there is substantial redundancy.

The failure of the real genetic code to instantiate any of the early ideas about efficiency or optimality must have been a disappointment, and it is hard to find much serious interaction between information theory and molecular biology for some time after the code was worked out. This split may have been premature. In a triplet code, there are three possible “reading frames,” and the first organism to have its entire genome sequenced—the bacteriophage Φ X174—makes use of the same DNA in multiple reading frames [92], an unexpected multiplexing that enhances efficiency in the original Shannon sense. Another observation is that the redundancy of the code is structured, so that random changes of a single base along the DNA are likely to map one amino acid either into itself (‘silent’ mutations) or into amino acids with very similar chemical properties. This would not be true in a random assignment of codons to amino acids, and it has been suggested that the real code may even be nearly optimal in this respect [93, 94]; if correct this means that the genetic code is not an optimal noiseless compression, as in the earliest (incorrect) proposals, but may be near

the bounds set by rate–distortion theory [95]. A full development of this idea will require understanding the cost of amino acid substitutions, which relates back to the analysis of sequence ensembles in Section III C.

The genetic code really is a strategy for communicating from DNA bases to amino acids, and here too there are errors. Although one often pictures the specificity of life’s mechanisms as deriving from the specific binding between complementary molecular components (as in the pairing of A with T and C with G in DNA), a beautiful chapter in the physics of life is the appreciation that these equilibrium interactions are not sufficiently selective, and hence that cells must build Maxwell demons to sort molecules with higher precision; the broad class of mechanisms that implement these demons is called kinetic proofreading [96, 97] [see also Section 4.5 of Ref [1]]. This creates an interesting tradeoff for the cell, which can operate at high fidelity but high energetic cost (paying for the demon), or trade the cost of errors against the reduced energetic cost of correcting them [98]. Although the costs and benefits of information flow in reading the genetic code have been much discussed, this still has not been given a fully information theoretic formulation.

B. Efficient representations in sensory processing

One of the earliest uses of information theory to think about biological systems was an attempt to estimate the amount of information carried by neuronal action potentials [10, 99]. The answer depends crucially on the time resolution with which such spikes can be generated and detected by neurons. If the average rate of action potentials is \bar{r} , and the system can work with a time resolution $\Delta\tau \ll 1/\bar{r}$, then the stream of spikes becomes a binary sequence of $1/\Delta\tau$ symbols per second, and the probability of a 1 (spike) is $\bar{r}\Delta\tau$. The maximum entropy of this sequence, which sets the capacity for the neuron to carry information, is then $S \sim \log_2(e/\bar{r}\Delta\tau)$ bits/spike. In contrast, if the system is limited to counting spikes in windows much larger than the typical interspike interval, then the entropy becomes $S \sim \log_2(e\bar{r}\Delta\tau)/(\bar{r}\Delta\tau)$ bits/spike. Evidently higher time resolution allows for higher information rates, and dramatically so: if spike rates are a few per second (as often true, on average), then time resolutions of a few milliseconds could allow for up to ten bits per spike. In 1952, when MacKay and McCulloch first made such estimates, locating real neural codes on the continuum from timing to counting was a matter of speculation. But this work established an agenda: could we measure the information that spike trains carry, for example about sensory inputs, and does this information approach the limits set by the entropy of the spike sequences themselves?

MacKay and McCulloch were thinking about a single neuron. But signals often are conveyed by larger populations of neurons, for example in the retina where many neurons respond to overlapping regions of the vi-

sual world. As early as 1959, Barlow suggested that processing in the retina would have the function of removing redundancy among the outputs of neighboring ganglion cells, and argued that several major features of these cells' responses could be understood in this way: a reduction in the response to sustained stimuli (termed adaptation) serves to reduce redundancy in time, while the tendency of cells to respond to differences between light intensity in a small patch and the surrounding region ("center-surround" organization) serves to reduce redundancy in space [100, 101].

It would take many years until these ideas about efficient coding in sensory processing would be made quantitative. An important qualitative idea, however, is that the efficiency of a code depends on the distribution of inputs. Thus, we might expect that real organisms use codes that are matched to the distribution of signals that they encounter in their natural environment; put simply, the function of the code would make sense only in context, and the relevant context is quite complex. This realization stands in contrast to the long tradition of using highly simplified stimuli in laboratory studies of sensory function, and has led to the emergence of a subfield of neuroscience concerned specifically with the structure of natural sensory signals and the neural response to these signals [102, 103].

The first stages of retinal coding occur with neurons that produce analog voltages rather than discrete spikes. Laughlin focused on cells in the first stage of processing after the photoreceptors in the fly retina, and asked if the nonlinear input/output relation that transforms light intensity into voltage is matched to the environmental distribution of light intensities [104]. If the voltage noise is small and constant throughout the dynamic range, then optimal information transmission implies that a (normalized) input/output relation should be equal to the cumulative probability distribution of the light intensity. By sampling natural scenes with a custom-built photodetector to predict the optimal input/output relation and comparing it with the measured one, Laughlin found a remarkable agreement with the theory, especially considering that there are no free parameters. Although there are obvious open questions, this was a really beautiful result that inspired the community to take these ideas more seriously.

For spiking neurons, measurements of information transmission in response to complex, dynamic stimuli began to appear in the 1990s. Initial measurements were based on decoding the spike trains to recover estimates of the input sensory signals [7, 105, 106], and then analysis tools for more direct estimates were developed [13, 107]. In the primary sensory neurons of the frog inner ear, the mechanical sensors of the cockroach, ganglion cells in the vertebrate retina [108] and subsequent layers of visual processing in mammals [109, 110], and motion-sensitive neurons deeper in fly visual system, as well as other systems, it was found that spike trains could transmit more than one bit per spike, and typically reached 30–50% of

the limit set by the spike train entropy, even with time resolutions in the millisecond range. In accord with the idea of matching to natural signals, information rates and coding efficiencies are larger with distributions of stimuli that capture features of the natural ensemble [111], and in some cases truly natural stimuli reveal codes that maintain their efficiency down to sub-millisecond resolution [112].⁶ Despite much prejudice to the contrary, there are demonstrations of efficient coding, with time resolutions of a few milliseconds, by single neurons in the cerebral cortex of our primate cousins [113, 114], suggesting strongly that these principles extend to brains much like our own.

If neurons implement coding strategies that are matched to the distribution of sensory inputs, what happens when this distribution changes? When we walk from a dark room out into the bright sunlight, the mean light intensity changes, and it takes a moment for our visual system to adapt to this new level. This adaptation is visible in the response of individual neurons in the retina, and this is one of the classical observations about neural coding, having roots in Adrian's first experiments in the 1920s [25]. But the efficient coding hypothesis predicts that the code should adapt not just to the mean, but to the whole distribution. In particular, natural images have an intermittent structure, so that the variance in light intensity varies significantly from point to point in a scene, and hence the variance seen by neurons responding to a small patch of the world will vary in time, so adaptation to the variance would enhance efficiency [115], and similar arguments could be made about other sensory modalities. Early experiments in the vertebrate retina demonstrated this adaptation to variance explicitly [116]. Subsequent experiments have shown that retinal adaptation is an even richer phenomenon, so that properly chosen distributions of inputs can drastically change even very basic features of the code exhibited by single neurons, and this is part of a growing realization that the retina is capable of much more complex computations than originally believed [117].

Adaptation to the variance of input signals has now been seen in many systems [118]: motion-sensitive neurons of the fly visual system [119, 120], auditory neurons at several stages of processing in mammals and birds [121–126], whisker motion sensing neurons in rodents [127], and in the farthest reaches of primate visual cortex, where cells responsive to the shapes of objects adapt to the variance along different parameters describing these shapes [128]. In many of these systems, changing the variance of the input distribution simply produces a rescaling of the input/output relation, as first observed in the fly motion-sensitive cells [119]. If the typical sig-

⁶ These results depend on getting control over the problems of estimating entropy and information from limited samples, as discussed in Section III A.

nals are large compared to internal noise levels but not so large as to saturate the system, then the dynamic range of the inputs themselves provides the only scale in the problem, and hence optimization must lead to the observed scaling behavior; in the fly it was even possible to check that the scaling factor chosen by the adaptation mechanism served to optimize the transmitted information [119]. When the input distribution shifts suddenly, the input/output relation adjusts with remarkable speed, perhaps so rapidly that there is no sharp distinction between adaptation and the nonlinear transient response. The time scales involved are essentially the shortest ones allowed by the need to gather statistics and verify that the distribution has, in fact, changed [120, 129].

Barlow’s original idea about redundancy reduction in the retina was implemented mathematically by Atick and Redlich [130], who approximated retinal ganglion cells as analog devices that respond to a linearly filtered version of their visual inputs. They searched for optimal filters, and showed that in low background lights ganglion cells should integrate over small regions of space and time, while at high light levels the response should be differentiating; this corresponds to a crossover between noise reduction and redundancy reduction being the dominant factors in enhancing information transmission, and agrees with the data.

If we imagine that neurons filter their inputs to remove redundancy, then the outputs should become close to uncorrelated or “white.” But many natural input signals have power spectra in time that are approximately $1/\omega$, and whitening this spectrum requires a transfer function between input and output $\sim \sqrt{\omega}$, which is not so easy to realize. Nonetheless, the linear transfer function from the photoreceptor cells to the second order neurons studied by Laughlin has precisely this form [131]. The transfer of the signal from one cell to the next is mediated, here as elsewhere in the brain, by the neurotransmitter molecules that are packaged into discrete vesicles. Careful measurements of the signal and noise characteristics in both cells leads to very high estimates of the information capacity of this connection (synapse), more than one thousand bits per second [132], and this is close to the physical limit set by the entropy of counting the vesicles [1, 10].

In the actual retina, of course, not all cells are identical, and is interesting to ask if we can predict the distribution over cell types from optimization principles. There is an old discussion about the choice of three color sensitivities for the photoreceptors in the bee’s eye, arguing that the particular combination that is found in a whole family of related species serves to provide maximal information about the identity of the flowers from which bees collect nectar [133]. In the primate retina, more rigorous arguments have been used to predict the distribution of the three cone types, which is both highly non-uniform and varies with position on the retina [134]. Retinal processing involves a plethora of cell types, but all visual systems divide signals into ‘on’ and ‘off’ pathways, which

respond with more action potentials when the light level rises above or below its local average value; Ratliff et al [135] have argued that the preponderance of off cells is matched to an asymmetry in the distribution of light intensities found in natural scenes, and that other differences between on and off pathways also serve to optimize information transmission.

If we focus on a single class of retinal ganglion cells, we find that they form a lattice across the surface of the retina [see, for example, Ref [56]]. The geometrical properties of this lattice, specifically, the size over which each cell integrates the signal relative to the lattice spacing, can be predicted from information optimization [136]. These lattices are, however, not perfectly regular due to discreteness and imperfections in the spatial arrangement of the photoreceptors, and in principle these irregularities should lead to a loss of information. Remarkably, these variations turn to be correlated with the local variations in the connectivity from the photoreceptor array to the ganglion cell layer in such a way as to almost perfectly compensate for the information loss [137].

Attempts to analyze the optimization problem for spiking neurons were quite scattered until the work of Smith and Lewicki on the coding of complex sounds [138, 139]. Rather than trying to build models of the process by which continuous inputs are converted into spikes, and then optimizing the parameters of these models, they chose the spike times themselves as the variables describing the encoding process and optimized the code assuming that the goal was to reconstruct the inputs using a simple linear algorithm, following [7]. The result of the calculation is then a set of spike trains for a population of neurons responding to acoustic inputs, and these spike trains could be analyzed in the same way that experimentalists routinely analyze data from real neurons. The results were impressive: optimal spike encoders for speech-like sounds, but not for other sound ensembles, almost perfectly match the characteristics of real auditory neurons. Further, the rate-distortion functions for these neural codes reach or even surpass the performance of the best artificial systems.

The problem of optimal coding has also been studied from a theoretical angle. In a population of rate neurons that suffer from negligible intrinsic noise, the optimal code decorrelates the inputs: this should happen not only at pairwise order (by whitening), as suggested by early works of Barlow, van Hateren and others [100, 140], but at all orders of correlation; in other words, the optimal code will minimize the mutual information between different neurons at the output. This led to the idea of ‘independent component analysis’ (ICA) and sparse codes. Imagining that neurons in the visual pathway perform linear filtering on the incoming stimuli and optimizing the filters using ICA over natural scenes, two groups recovered the well-known Gabor-like filters found in the primary visual cortex [141, 142].

In the presence of nonlinearities and noise, optimizing information is much more difficult; early work in this di-

rection was done by Linsker [143]. If neurons are noisy, perfect decorrelation is no longer optimal. In this case, some amount of redundancy (in engineered codes in the form of error-correcting bits) can be beneficial [144]. It is interesting to note that real neuronal spike trains are usually weakly correlated at the pairwise level and redundant [145], and at the level of a large population this redundancy can suffice to predict well the behavior of one neuron from the state of the rest of the network [146]. If noise is small, information optimization can be carried out in a continuous nonlinear system, and the tradeoffs between decorrelation and redundancy reduction can be explored in detail; in this framework, Karklin and Simoncelli managed to derive many properties of the retina simultaneously from the statistical structure of natural scenes [147]. This optimization, as well as previous work [148], also incorporated an important constraint on metabolic efficiency (cost of spiking). By departing from analytically tractable cases, numerical exploration has been used to study optimal coding at arbitrary noise levels with high-dimensional stimuli [149]. Using a model for an interacting neural population that can be linked to maximum entropy models with an Ising form (Section III C), small networks experienced a transition between redundant coding and decorrelation based on the statistical structure of the inputs and the level of noise in single units. Surprisingly, when noise in single units is high enough, the redundant strategy with information-maximizing connectivity yields a network with strong attractor states to which different stimuli map uniquely, recovering Hopfield-like associative memory [39] from an optimization principle. Mapping out a complete ‘phase diagram’ of optimal coding strategies for high-dimensional noisy nonlinear input/output maps remains a challenge.

Finally, we note that many popular models for coding by large populations of neurons are inefficient in the extreme. It is widely assumed that neurons in a single region of the visual cortex, for example, are selective for a very limited number of features (e.g., the velocity of motion, or the orientation of an edge). If all the neurons in the population respond independently to these features, with each of the N cells preferring a slightly different value, and we assume that what matters is the number of spikes each neuron generates in a relatively large window of time (“rate coding”), then the response of the population as a whole will allow us to determine the value of the feature with a precision $\sim 1/\sqrt{N}$. But this means that the information conveyed by the population is $\sim \log(N)$, and thus the information per neuron vanishes as N becomes large. Are we missing something? One possibility is that the diversity of dynamical responses among neurons with redundant static feature selectivity can be sufficient to create a combinatorial code for stimulus timing, allowing the information to grow linearly with the number of neurons over a much larger dynamic range [150]. A more radical alternative is suggested by the discovery of “grid cells,” neurons that encode an animal’s position

relative to a lattice [151]. In this system, it is possible for a population of neurons with different lattice spacing to generate a truly efficient code, with more neurons allowing exponentially more accuracy in the representation of position [152, 153]. These ideas point to the need for direct measurement of the information carried by large populations of cells, under reasonably natural conditions.

C. Biochemical and genetic networks

Even the simplest of bacteria have genes that code for many hundreds of proteins, and our own genome codes for more than ten thousand proteins. No single cell needs to make all these proteins at once, and there are many layers of regulation to control the readout or expression of individual genes. One important class of these mechanisms is the binding of protein molecules, called transcription factors, to specific sequences along the DNA, close to the start of the coding region of a gene. Transcription factor binding can both repress or activate protein synthesis; in bacteria there are simple geometrical pictures for how this can happen that are probably good approximations to reality, but the situation in higher organisms is less clear. In all cases, transcription factors act at concentrations in the nanoMolar range, and $1\text{ nM} \sim 1\text{ molecule}/\mu\text{m}^3$. But a bacterium has a volume of only a few cubic microns, which means that the regulation of transcription is controlled by just handfuls of molecules.

The small numbers of molecules involved means that the regulation of gene expression is, inevitably, noisy, and this noise was measured directly in the early 2000s [154–156]. How does this noise limit the transmission of information through the system, from the input concentration of transcription factors to the output level of the protein encoded by the target gene? Here information provides a measure of regulatory power, counting the number of reliably distinguishable output states that can be accessed by adjusting the inputs [157, 158]. We recall that if the input/output relation and noise in the system are fixed, then we can optimize information transmission by adjusting the distribution of input signals. This matching condition, which is essentially the same as in Laughlin’s discussion of the fly retina, allows the system to get as much regulatory power as possible out of a limited number of molecules. An important difference from Laughlin’s discussion is that noise in the case of transcriptional regulation must be (and is measured to be) strongly dependent on the levels of input and output.

Early events in development of the fruit fly embryo provide a productive testing ground for ideas about the regulation of gene expression. During the construction of the egg, the mother places the mRNA for a transcription factor at the end that will become the head; the combination of synthesis, diffusion and degradation of the protein (Bicoid) translated from this mRNA results in a spatially varying concentration along the length of

embryo—a gradient in the concentration of a primary morphogen, to use the conventional language. Bicoid is the input to an interacting network of genes (called “gap genes” and “pair-rule genes”) that end up being expressed in distinct spatial patterns, which will define the final body plan of the larval fly. Importantly, these patterns are visible within a few hours after the egg is laid, before cells start to move and even before there are membranes marking the boundaries between the cells.

A first generation of quantitative experiments on this system mapped the relationship between the input transcription factor concentration and the output level of one of the gap genes, making measurements nucleus by nucleus in single embryos, and demonstrating that the results are consistent across embryos [28, 159]. If we think of x as the input concentration and y as the output concentration, then these experiments allow us to estimate the behavior of the communication channel, $P(y|x)$. We can solve numerically for the distribution of inputs that optimizes information transmission, and compare with experiment. Because the mean input/output relation has broad, flat segments, many different input distributions give essentially the same mutual information between x and y , but all these near-optimal input distributions have the same output distribution, so we make the comparison of the predicted $P_{\text{out}}(y)$ with experiment [160]. The agreement, which involves no adjustable parameters, is excellent. Importantly, we can also estimate directly the mutual information $I(x;y)$ from the data, and find that it is 0.88 ± 0.09 of the optimum, so that the system really is transmitting nearly the maximal information given the measured noise levels. This information is about a bit and a half, showing that the regulation of gene expression in this case is more than just an on/off switch.

Thinking of transcriptional regulation in the early embryo as having one input and one output is a drastic oversimplification, but dealing with multiple inputs and multiple outputs is very difficult. It is thus useful to note that the information of relevance for development is information about the position of the cell. If we can measure the expression levels of multiple genes simultaneously, we can ask how much information these level provide about position even if we can’t also measure all the input signals. Indeed, the qualitative concept of positional information is very old [161], and it is attractive to try and make this quantitative [162, 163]. Refinement of classical staining methods make it possible to measure, simultaneously and with high accuracy, the normalized concentrations of four gap genes that provide dominant signals along the central 80% of the embryo’s length [26]. Analysis of these experiments shows that it is possible to “decode” the expression levels and estimate the position of a nucleus to an accuracy of better than 1% of the length of egg, corresponding to a total of 4.14 ± 0.05 bits of information about position [162]. This precision is comparable to the spacing between neighboring cells, and to the precision of subsequent developmental events, suggesting that the embryo really does use this informa-

tion.

Each of the four gap genes has a complicated profile of concentration vs. position along the embryo, and the noise levels have an even more complicated structure. But since cells are almost uniformly distributed, matching of the distribution on inputs to the noise characteristics of the system would require that the error in estimating position be almost uniform as well. This is, in fact, what we see in the data [162]; indeed, if we imagine trying to redistribute the cells so as to transmit more information, we could gain less than 2%. This seems a striking signature of optimization.

Optimization of information flow should involve more than matching: we should be able to find the architecture and parameters of genetic networks that allow for maximal information transmission. One ambitious idea is to define a dynamics in the space of networks that is a schematic of the evolutionary process, and then use information as a surrogate for fitness [164]. The networks that emerge from such simulations have many qualitative features in common with the networks functioning in real embryos, including the fruit fly. One can find examples of multiple genes driven by a common maternal morphogen, each one activating itself and repressing other genes in the network, much as for the gap genes. A conceptually striking feature of this work is that it explicitly constructs paths through the space of possible networks that lead from simpler to more complex body plans, through gene duplication followed by mutation. Importantly, almost all the steps along the path are either neutral or increase the fitness. As the authors emphasize, this “... provides a quantitative demonstration that continuous positive selection can generate complex phenotypes from simple components by incremental evolution, as Darwin proposed.”

Rather than reconstructing entire evolutionary trajectories for networks, we can try to find locally optimal networks within some broad class of models [165–167]. Here again we can see features that remind us of real networks, but there are surprises. When the available concentration of transcription factors is severely limited, the optimal strategy for transmitting information is for a single transcription factor to broadcast to multiple, completely redundant targets, essentially using the redundancy to make multiple measurements of a weak signal. Observations of redundancy in genetic networks usually is taken as *prima facie* evidence against optimization and in favor of an interpretation in which network architecture is just an artifact of evolutionary history, but here we have a quantitative demonstration that such qualitative arguments can be wrong. More subtly, these calculations show that even qualitative features of the networks which maximize information transmission depend on some molecular details of the regulatory mechanisms. This is slightly worrisome, but also suggests a path for some of the intricacy of complex molecular interactions to influence macroscopic behavior. A related approach to these same problems is inspired not by embryos but by ar-

tificial genetic circuits that have more limited topologies and inputs, allowing for an exhaustive exploration [168]. In these cases, the optimization of information transmission leads to very broad maxima in the space of network parameters, so that network topology is the dominant consideration. These efforts at full optimization have been limited to static inputs and outputs; information transmission in dynamic contexts raises new questions, which are just starting to be explored [169, 170]. Finally, interesting questions that have been highlighted recently include the link between information transmission and cellular decision making [171], and using information theory to put strict bounds on the reliability of biochemical processes even when many of these processes are unknown or unobserved [172].

D. More ambitious goals

Success at gambling, and perhaps at life, plausibly depends on how much information we have. If we know more about the outcome of a coin flip (or a horse race, or the stock market, or ...), then we should be able to win our bets more often, and hence make more money. But it is not obvious that the knowledge measured by Shannon’s bits quantifies the useful knowledge in this context. Kelly analyzed a simple gambling game and showed that the maximum rate at which winnings can grow is precisely the information that bettors have about the outcome, and his proof is constructive, so we actually know how to achieve this maximum [173]. We can think of the choices that organisms make—e.g., to express different levels of proteins—as being bets on the state of the environment, and there have been several efforts to construct scenarios that match the structure of Kelly’s argument, so that the growth rate of the organisms becomes equal to the information that they have about environmental conditions [174–176]. More generally, from rate distortion theory we know that achieving a given average growth rate across a wide range of conditions will require some minimum amount of information about the environment [177]. These connections between growth rates and information provide a path for evolution to select for organisms that gather and represent more information about their environment.

Another interesting connection of information theoretic optimization to biological function is in the problem of search [178]. Organisms search for food, for shelter, for mates, and more; it may even be possible to think more generally about computations or decisions as the search for solutions. If we imagine that we are looking for something in one of N possible places, we can order these by our current estimate of the probability P_n that we will find what we are looking for in the n^{th} place, so that $P_1 > P_2 > \dots > P_n$. Then if we look first in the most likely place, and then down the list, the mean number of tries before we find what we are looking for is $\bar{n} = \sum_{n=1}^N nP_n$. Given \bar{n} , there is a probability distribu-

tion ($P_n \propto e^{-\lambda n}$) that has maximal entropy. Conversely, if our uncertainty corresponds to a probability distribution with entropy S , then the mean search time \bar{n} has a minimum, independent of all other details. The only way we can speed up our search is to reduce the entropy of the distribution, or to gain information in the Shannon sense; while there is no guarantee that our search will go as fast as allowed by this information, it cannot go faster.

The idea of “infotaxis” is to replace the task of finding something with the task of gaining information about its location [178]. In the context of a flying insect searching for the source of an odor, where new information arrives only intermittently when the flight path crosses odor plumes, optimizing the gathering of information by properly alternating upwind flight with cross-wind ‘casting’ in fact leads to the source, and the flight paths that are generated by this algorithm are reminiscent of the paths taken by insects in natural environments.⁷ For the purposes of this review, perhaps the most important result from the infotaxis idea is the demonstration that a biologically grounded, goal-directed behavior can be replaced with an abstract, information theoretic optimization principle. This certainly encourages us to think that some of the nitty-gritty of biology should be derivable from more general principles.

In their earliest work, Attneave and Barlow expressed the hope that the search for efficient representations of sensory data would drive discovery of features and objects that form our perceptual vocabulary, and Attneave even tried simple experiments to show that human subjects use efficient strategies in abstracting simple figures into sketches [100, 101, 180]. As an example, if our world consists of a finite set of rigid objects, each seen from many different perspectives, then the most efficient representation of the image data reaching our eyes would consist of a discrete label for object identity and a few parameters to describe the viewing geometry. In this case, optimization in an information theoretic sense would require our brains to recognize objects, and the labeling of these objects would constitute a ‘naming’ of the meaningful features of the visual world, without any additional instructions. At its most ambitious, the claim would be that these features acquire meaning because they are maximally informative, as opposed to having meaning imposed upon them.

Shannon frequently used English as an example of information theoretic ideas, and even played a “guessing

⁷ As an aside, we note that insects seem to “know” about the intermittent structure of turbulent flows [179]. In particular, moths exposed to smooth ribbon plumes, which provide a continuous signal and should in principle be easy to follow, fly along strongly zig-zag trajectories, spending most their time casting crosswind, while the same moths exposed to intermittent or turbulent plumes fly efficiently toward the source of the odor. This is one of the strongest, qualitative results demonstrating that neural processing is tuned to the statistical structure of physical signals that arise in the natural environment.

game” with human subjects to estimate the entropy of written texts [181]. On the other hand, there is a persistent suspicion that there must be more than just statistics involved in the structure and meaning of language, and one of the foundational papers of modern linguistics is essentially an attack on Shannon’s probabilistic models for language [182]. Some of these early arguments are related to the crude Markovian models that Shannon used as illustrations, and some are related to the confusion between learning a probabilistic model and simply counting events,⁸ but questions remain, and Shannon himself eschewed any connection to the problem of meaning. In what proved to be a more direct attack on this question, Pereira et al [184] developed an algorithm to cluster nouns based on the distribution of verbs that occur in the same sentence; their approach can be thought of as compressing the description of the noun while preserving information about the identity of the verb, and hence is yet another example of efficient representation [185]. The resulting clusters exhibit a striking degree of semantic (rather than just syntactic) coherence, and it hard to resist saying that this purely information theoretic approach is extracting meaning. These references are just a small sampling from the roots of a huge literature. By now, we all interact with probabilistic models of language every day (through internet search engines, automatic translation, ...), so perhaps the larger conceptual question of whether the search for efficient representation captures meaning has become less urgent. Nonetheless, we suspect that there are many open questions here, including whether we really understand how to write down models that capture the very long-ranged correlations that are present in texts.

The entropy of a probability distribution is defined without reference to anything else, at least in the case of discrete variables, but information is always about something. Thus it makes no sense to ask about the information content of a signal; one has to ask how much information the signal provides about some quantity of interest.⁹ Arguably the most successful uses of information

theory in thinking about biological systems are in those cases where the question “information about what?” has a clear answer. But this makes us worry that more of the predictive power of our theories resides in this qualitative knowledge of the biologically relevant information than in the more mathematical principle of optimizing information transmission. Can we do something more general than just making a list of which information is relevant for each particular system?

Information cannot be useful to an organism unless it carries predictive power: if we use information to guide our actions, it takes time to act, and hence any measure of the success of our actions will be in relation to a future state of the world [186]. While not all predictive information is equally useful, non-predictive information is useless, and so separating predictive from non-predictive information is an important task. If we observe a system for time T , then the distribution from which trajectories are drawn has an entropy $S(T)$ that we expect to be extensive at large T . But the mutual information between, for example, observations on the past ($-T < t \leq 0$) and an equally long future ($0 < t \leq T$) is given by $I_{\text{pred}}(T) = 2S(T) - S(2T)$. Importantly, any extensive component of the entropy cancels, and the information that the past provides about the future must be sub-extensive, even if we allow for an unlimited future, so most of the bits we observe must be discarded if we want to isolate the predictive bits. This sub-extensive entropy is related to our intuitive notions of the complexity of the underlying processes [187, 188], and providing an efficient representation of predictive information is connected to problems ranging from signal processing to learning [186, 189]. Recent work provides preliminary evidence that populations of neurons in the retina indeed provide such an efficient representation, separating predictive information from the non-predictive background at least in simplified visual worlds [190], but much more remains to be done in this direction.

V. OUTLOOK

We can trace the use of information theoretic ideas to describe biological systems back roughly sixty years, almost to the origins of information theory itself. We have emphasized that these connections between information and the phenomena of life have two very different flavors, one grounded in data and the other aiming for a theory in which the behaviors of real biological systems are derivable from some sort of optimization principle. Both directions have been productive, and as the volume of quantitative experimental data on biological systems

⁸ Thus, Chomsky famously argued that statistical approaches could not distinguish between sentences that have never occurred but are instantly recognizable as grammatical (Colorless green ideas sleep furiously) and sentences that never occurred because they are forbidden by grammatical rules (Furiously sleep ideas green colorless) [182]. As a concrete rejoinder, Pereira constructed a simple Markov model in which successive words belong to clusters, so that $P(W_{n+1}, W_n) = \sum_C Q(W_{n+1}|C)Q(W_n|C)P(C)$, and learned the model on a modest body of newspaper texts using just sixteen clusters. He found that Chomsky’s two example sentences have probabilities that differ by a factor of more than 10^5 [183]. But between the dates of these references, a much more sophisticated view of the learning problem had emerged.

⁹ We could be interested in the value of the signal itself, in which case the information is equal to the entropy. But this is seldom the case. If the signal is the text of a newspaper, for example, we probably are not interested in the text itself, but what it tells

us about events in the world. In particular, the entropy of the text is increased if we introduce random spelling errors, but the information about events is decreased.

grows, it seems clear that the utility of the approaches to data analysis described in Section III will grow in parallel. The search for real theories is more ambitious.

It is important that we have concrete examples of biological systems that are operating near an optimum of information transmission or efficiency of representation, as reviewed in Section IV. What we don't know is whether these are isolated instances, or examples of a general principle. It is relatively easy to work out what optimization means in simplified settings, but much harder to address the theoretical issues that arise in more realistic contexts, e.g. with many interacting degrees of freedom and the complexities of fully natural signals. At the same time, it took decades to get to the point where even the most basic quantities—the actual information carried by biological signals—could be measured reliably. Now that many of these foundational steps have been solidified, we

believe that the coming years will see much more meaningful confrontations between theory and experiment, in a wide range of systems.

Acknowledgments

We thank our many collaborators and colleagues, whose names appear in the references, for teaching us so much and for making the exploration of these ideas so pleasurable. Our work was supported in part by the US National Science Foundation (PHY-1305525 and CCF-0939370), by the Austrian Science Foundation (FWF P25651), by the Human Frontiers Science Program, and by the Simons and Swartz Foundations.

-
- [1] W Bialek, *Biophysics: Searching for Principles*. W Bialek (Princeton University Press, Princeton NJ, 2012).
 - [2] CE Shannon, A mathematical theory of communication. *Bell Sys Tech J* **27**, 379–423 & 623–656 (1948).
 - [3] TM Cover and JA Thomas, *Elements of Information Theory* (Wiley, New York, 1991).
 - [4] M Mézard and A Montanari, *Information, Physics, and Computation* (Oxford University Press, Oxford and New York, 2009).
 - [5] CE Shannon, Coding theorems for a discrete source with a fidelity criterion. *IRE Nat Conv Rec* **4**, 142–163 (1959).
 - [6] T Berger, *Rate-Distortion Theory: A Mathematical Basis for Data Compression* (Prentice-Hall, Englewood Cliffs NJ, 1971).
 - [7] W Bialek, F Rieke, RR de Ruyter van Steveninck, and D Warland, Reading a neural code. *Science* **252**, 1854–1857 (1991).
 - [8] O Marre, VB Soler, KD Simmons, T Mora, G Tkačik, and MJ Berry II, High accuracy decoding of dynamical motion from a large retinal population. arXiv.org:1407.5946 [q-bio.NC] (2014).
 - [9] AB Schwartz, Direct cortical representation of drawing. *Science* **265**, 540–542 (1994).
 - [10] F Rieke, D Warland, R de Ruyter van Steveninck, and W Bialek, *Spikes: Exploring the Neural Code*. (MIT Press, Cambridge, 1997).
 - [11] GA Miller, Note on the bias of information estimates. In *Information Theory in Psychology II-B*, H Quastler, ed, pp 95–100 (Free Press, Glencoe IL, 1955).
 - [12] S Panzeri and A Treves, The upward bias in measures of information derived from limited data samples. *Neural Comp* **7**, 399–407 (1995).
 - [13] SP Strong, R Koberle, RR de Ruyter van Steveninck, and W Bialek, Entropy and information in neural spike trains. *Phys Rev Lett* **80**, 197–200 (1998).
 - [14] L Paninski, Estimation of entropy and mutual information. *Neural Comp* **15**, 1191–1253 (2003).
 - [15] I Nemenman, F Shafee, and W Bialek, Entropy and inference, revisited. In *Advances in Neural Information Processing 14*, TG Dietterich, S Becker, and Z Ghahramani, eds, pp 471–478 (MIT Press, Cambridge, 2002).
 - [16] I Nemenman, W Bialek, and R de Ruyter van Steveninck, Entropy and information in neural spike trains: Progress on the sampling problem. *Phys Rev E* **69**, 056111 (2004).
 - [17] E Archer, I Park, JW Pillow, Bayesian estimation of discrete entropy with mixtures of stick-breaking priors. *Adv Neural Info Proc Syst* **25**, 2024–2032 (2012).
 - [18] E Archer, I Park, JW Pillow, Bayesian entropy estimation for binary spike train data using parametric prior knowledge. *Adv Neural Info Proc Syst* **26**, 1700–1708 (2013).
 - [19] JD Victor, Binless strategies for estimation of information from neural data. *Phys Rev E* **66**, 51903 (2002).
 - [20] JB Kinney and GS Atwal, Equitability, mutual information, and the maximal information coefficient. *Proc Natl Acad Sci (USA)* **111**, 3354–3359 (2014).
 - [21] S-K Ma, Calculation of entropy from data of motion. *J Stat Phys* **26**, 221–240 (1981).
 - [22] OG Berg and PH von Hippel, Selection of DNA binding sites by regulatory proteins. *J Mol Biol* **193**: 723–50.
 - [23] TD Schneider and RM Stephens, Sequence logos: a new way to display consensus sequences. *Nucl Acids Res* **18**, 6097–6100.
 - [24] Z Wunderlich and LA Mirny, Different gene regulation strategies revealed by analysis of binding motifs. *Trends Genet* **25**, 434–40.
 - [25] ED Adrian, *The Basis of Sensation: The Action of the Sense Organs* (WW Norton, New York, 1928).
 - [26] JO Dubuis, R Samanta, and T Gregor, Accurate measurements of dynamics and reproducibility in small genetic networks. *Mol Sys Bio* **9**, 639 (2013).
 - [27] RY Tsien, The green fluorescent protein. *Annu Rev Biochem* **67**, 509–544 (1998).
 - [28] T Gregor, EF Wieschaus, AP McGregor, W Bialek, and DW Tank, Stability and nuclear dynamics of the Bicoid morphogen gradient. *Cell* **130**, 141–152 (2007).
 - [29] AH Morrison, M Scheeler, J Dubuis, and T Gregor, Quantifying the Bicoid morphogen gradient in living fly embryos. In *Imaging in Developmental Biology*. J

- Sharpe and RO Wang, eds, pp 398–406 (CSHL Press, 2012).
- [30] A Raj A, P van den Bogaard, SA Rifkin, A van Oudenaarden, and S Tyagi, Imaging individual mRNA molecules using multiple singly labeled probes. *Nat Methods* **5**, 877–879 (2008).
- [31] SC Little, M Tikhonov, and T Gregor, Precise developmental gene expression arises from globally stochastic transcriptional activity. *Cell* **154**, 789–800 (2013).
- [32] R Cheong, A Rhee, CJ Wang, I Nemenman, and A Levchenko, Information transduction capacity of noisy biochemical signaling networks. *Science* **334**, 354–8 (2010).
- [33] J Selimkhanov, B Taylor, J Yao, A Pilko, J Albeck, A Hoffmann, L Tsimring, and R Wollman, Accurate information transmission through dynamic biochemical signaling networks. *Science* **346**, 1370–1373 (2014).
- [34] N Slonim, GS Atwal, G Tkačik, and W Bialek, Estimating mutual information and multi-information in large networks. arXiv:cs.IT/0502017 (2005).
- [35] N Slonim, GS Atwal, G Tkačik, and W Bialek, Information-based clustering. *Proc Nat'l Acad Sci USA* **102**, 18297–302 (2005).
- [36] AA Margolin, I Nemenman, K Basso, C Wiggins, G Stolovitzky, R Dalla Favera, and A Califano, ARACNE: an algorithm for the reconstruction of gene regulatory networks in a mammalian cellular context. *BMC Bioinformatics* **S7** (2006).
- [37] PM Bowers, SJ Cokus, D Eisenberg, and TO Yeates, Use of logic relationships to decipher protein network organization. *Science* **306**, 2246–9 (2004).
- [38] N Slonim, O Elemento, and S Tavazoie, Ab initio genotype-phenotype association reveals intrinsic modularity in genetic networks. *Mol Syst Biol* **2**, 2006.0005 (2006).
- [39] JJ Hopfield, Neural networks and physical systems with emergent collective computational abilities. *Proc Natl Acad Sci USA* **79**, 2554–2558.
- [40] DJ Amit (1989) *Modeling Brain Function: The World of Attractor Neural Networks* (Cambridge University Press, Cambridge).
- [41] J Hertz, A Krogh, and RG Palmer (1991) *Introduction to the Theory of Neural Computation* (Addison Wesley, Redwood City).
- [42] J Toner and Y Tu, Long-range order in a two-dimensional XY model: How birds fly together. *Phys Rev Lett* **75**, 4326–4329 (1995).
- [43] J Toner and Y Tu, Flocks, herds, and schools: A quantitative theory of flocking. *Phys Rev E* **58**, 4828–4858 (1998).
- [44] S Ramaswamy, The mechanics and statistics of active matter. *Annu Rev Cond Matt Phys* **1**, 323–345 (2010).
- [45] ET Jaynes, Information theory and statistical mechanics. *Phys Rev* **106**, 620–630 (1957).
- [46] E Schneidman, S Still, MJ Berry II, and W Bialek, Network information and connected correlations. *Phys Rev Lett* **91**, 238701 (2003); arXiv:physics/0307072 (2003).
- [47] JB Keller and B Zumino, Determination of intermolecular potentials from thermodynamic data and the law of corresponding states. *J Chem Phys* **30**, 1351–1353 (1959).
- [48] JT Chayes, L Chayes, and E Lieb, The inverse problem in classical statistical mechanics. *Commun Math Phys* **93**, 57–121 (1984).
- [49] E Caglioti, T Kuna, J Lebowitz, and E Speer, Point processes with specified low order correlation. *J Markov Proc and Related Fields* **12**, 257–272 (2006).
- [50] N Tishby, E Levin, and SA Solla, A statistical approach to learning and generalization in layered networks. *Proc IEEE* **78**, 1568–1580 (1990).
- [51] JS Yedidia, WT Freeman, and Y Weiss, Understanding belief propagation and its generalizations. In *Exploring Artificial Intelligence in the New Millennium* G Lake-meyer and B Nebel, eds, pp 239–269 (Morgan Kaufman, San Francisco, 2003).
- [52] M Mézard, Message passing between disciplines. *Science* **301**, 1685–1686 (2003).
- [53] P Mehta and DJ Schwab, An exact mapping between the variational renormalization group and deep learning. arXiv:1410.3831 [stat.ML] (2014).
- [54] E Schneidman, MJ Berry II, R Segev, and W Bialek, Weak pairwise correlations imply strongly correlated network states in a neural population. *Nature* **440**, 1007–1012 (2006); arXiv:q-bio.NC/0512013 (2005).
- [55] G Tkačik, E Schneidman, MJ Berry II, and W Bialek, Ising models for networks of real neurons. arXiv.org:q-bio/0611072 (2006).
- [56] J Shlens, GD Field, JL Gauthier, MI Grivich, D Petrusca, A Sher, AM Litke, and EJ Chichilnisky, The structure of multi-neuron firing patterns in primate retina. *J Neurosci* **26**, 8254–8266 (2006).
- [57] A Tang, D Jackson, J Hobbs, W Chen, JL Smith, H Patel, A Prieto, D Petruscam MI Grivich, A Sher, P Hottowy, W Dabrowski, AM Litke, and JM Beggs, A maximum entropy model applied to spatial and temporal correlations from cortical networks in vitro. *J Neurosci* **28**, 505–518 (2008).
- [58] G Tkačik, E Schneidman, MJ Berry II, and W Bialek, Spin-glass models for a network of real neurons. arXiv.org:0912.5409 (2009).
- [59] J Shlens, GD Field, JL Gauthier, M Greschner, A Sher, AM Litke, and EJ Chichilnisky, The structure of large-scale synchronized firing in primate retina. *J Neurosci* **29**, 5022–5031 (2009).
- [60] IE Ohiorhenuan, F Mechler, KP Purpura, AM Schmid, Q Hu, and JD Victor, Sparse coding and higher-order correlations in fine-scale cortical networks. *Nature* **466**, 617–621 (2010).
- [61] E Ganmor, R Segev, and E Schneidman, Sparse low-order interaction network underlies a highly correlated and learnable neural population code. *Proc Natl Acad Sci (USA)* **108**, 9679–9684 (2011).
- [62] G Tkačik, O Marre, T Mora, D Amodei, MJ Berry II, and W Bialek, The simplest maximum entropy model for collective behavior in a neural network. *J Stat Mech* P03011 (2013).
- [63] E Granot-Atedgi, G Tkačik, R Segev, and E Schneidman Stimulus-dependent maximum entropy models of neural population codes. *PLoS Comput Biol* **9**, e1002922 (2013).
- [64] G Tkačik, O Marre, D Amodei, E Schneidman, W Bialek, and MJ Berry II, Searching for collective behavior in a large network of sensory neurons. *PLoS Comput Biol* **10**, e1003408 (2014).
- [65] T Mora, S Deny, and O Marre, Dynamical criticality in the collective activity of a population of retinal neurons. arXiv:1410.6769 [q-bio.NC] (2014).
- [66] G Tkačik, T Mora, O Marre, D Amodei, MJ Berry II, and W Bialek, Thermodynamics for a network of neu-

- rons: Signatures of criticality. arXiv.org:1407.5946 [q-bio.NC] (2014).
- [67] M Socolich, SW Lockless, WP Russ, H Lee, KH Gardner, and R Ranganathan, Evolutionary information for specifying a protein fold. *Nature* **437**, 512–518 (2005).
- [68] WP Russ, DM Lowery, P Mishra, MB Yaffe, and R Ranganathan, Natural-like function in artificial WW domains. *Nature* **437**, 579–583 (2005).
- [69] W Bialek and R Ranganathan, Rediscovering the power of pairwise interactions. arXiv.org:0712.4397 [q-bio.QM] (2007).
- [70] M Weigt, RA White, H Szurmant, JA Hoch, and T Hwa, Identification of direct residue contacts in protein-protein interaction by message passing. *Proc Nat'l Acad Sci (USA)* **106**, 67–72 (2009).
- [71] DS Marks, LJ Colwell, R Sheridan, TA Hopf, A Pagnani, R Zecchina, and C Sander, Protein 3D structure computed from evolutionary sequence variation. *PLoS One* **6**, e28766 (2011).
- [72] TA Hopf, LJ Colwell, R Sheridan, B Rost, C Sander, and DS Marks, Three-dimensional structures of membrane proteins from genomic sequencing. *Cell* **149**, 1607–162 (2012).
- [73] JI Sulkowska, F Morocos, M Weigt, T Hwa, and JN Onuchic, Genomics-aided structure prediction. *Proc Natl Acad Sci (USA)* **109**, 10340–10345 (2012).
- [74] T Mora, AM Walczak, W Bialek, and CG Callan, Maximum entropy models for antibody diversity. *Proc Natl Acad Sci (USA)* **107**, 5405–5410 (2010); arXiv:0912.5175 [q-bio.GN] (2009).
- [75] AL Ferguson, JK Mann, S Omarjee, T Ndung'u, BD Walker, and AK Chakraborty, Translating HIV sequences into quantitative fitness landscapes predicts viral vulnerabilities for rational immunogen design. *Immunity* **38**, 606–617 (2013).
- [76] W Bialek, A Cavagna, I Giardina, T Mora, E Silvestri, M Viale, and A Walczak, Statistical mechanics for natural flocks of birds. *Proc Natl Acad Sci (USA)* **109**, 4786–4791 (2012); arXiv.org:1107.0604 [physics.bio-ph] (2011).
- [77] W Bialek, A Cavagna, I Giardina, T Mora, O Pohl, E Silvestri, M Viale, and A Walczak, Social interactions dominate speed control in poising natural flocks near criticality. *Proc Natl Acad Sci (USA)* **111**, 7212–7217 (2014); ; arXiv.org:1307.5563 [physics.bio-ph] (2013).
- [78] A Cavagna, L Del Castillo, S Dey, I Giardina, S Melillo, L Parisi, and M Viale, Short-range interaction vs long-range correlation in bird flocks. arXiv:1407.6887 (2014).
- [79] A Cavagna, A Cimorelli, I Giardina, G Parisi, R Santagati, F Stefanini, and M Viale, Scale-free correlations in starling flocks. *Proc Natl Acad Sci (USA)* **107**, 11865–11870 (2010).
- [80] A Attanasi, A Cavagna, L Del Castello, I Giardina, S Melillo, L Parisi, O Pohl, B Rossaro, E Shen, E Silvestri, and M Viale, Finite-size scaling as a way to probe near-criticality in natural swarms. *Phys Rev Lett* **113**, 238102 (2014).
- [81] TO Sharpee, NC Rust, and W Bialek, Analyzing neural responses to natural signals: maximally informative dimensions. *Neural Comput* **16**, 223–250 (2004).
- [82] TO Sharpee, H Sugihara, AV Kurgansky, SP Rebrik, MP Stryker, and KD Miller, Adaptive filtering enhances information transmission in visual cortex. *Nature* **439**, 936–942 (2006).
- [83] JD Fitzgerald, RJ Rowekamp, LC Sincich, and TO Sharpee, Second order dimensionality reduction using minimum and maximum mutual information models. *PLoS Comput Biol* **7**, e1002249 (2011).
- [84] K Rajan and W Bialek, Maximally informative “stimulus energies” in the analysis of neural responses to natural signals. *PLoS One* **8**, e71959 (2013); arXiv.org:1201.0321 [q-bio.NC] (2012).
- [85] M. Eickenberg, RJ Rowekamp, M Kouh, and TO Sharpee, Characterizing responses of translation-invariant neurons to natural stimuli: Maximally informative invariant dimensions. *Neural Comput* **24**, 2384–2421 (2012).
- [86] K Rajan, O Marre, and G Tkačik, Learning quadratic receptive fields from neural responses to natural stimuli. *Neural Comput* **25**, 1661–1692 (2013).
- [87] JB Kinney, G Tkačik, and CG Callan Jr, Precise physical models of protein-DNA interaction from high-throughput data. *Proc Nat'l Acad Sci (USA)* **104**, 501–506 (2007).
- [88] JB Kinney, A Murugan, CG Callan Jr, and EC Cox, Using deep sequencing to characterize the biophysical mechanism of a transcriptional regulatory sequence. *Proc Natl Acad Sci (USA)* **107**, 9158–9163 (2010).
- [89] O Elemento, N Slonim, and S Tavazoie, A universal framework for regulatory element discovery across all genomes and data types. *Molecular Cell* **28**, 337–350 (2007).
- [90] FHC Crick, On protein synthesis. *Symp Soc Exp Biol* **12**, 138–163 (1958).
- [91] FHC Crick, Recent excitement in the coding problem. *Prog Nucleic Acid Res Mol Biol* **1**, 163–217 (1963).
- [92] F Sanger, GM Air, BG Barrell, NL Brown, AR Coulson, JC Fiddes, CA Hutchinson III, PM Slocombe, and M Smith, Nucleotide sequence of bacteriophage Φ X174 DNA. *Nature* **265**, 687–695 (1977).
- [93] SJ Freeland and LD Hurst, The genetic code is one in a million. *J Mol Evol* **47**, 238–248 (1998).
- [94] SJ Freeland, RD Knight, LF Landweber, and LD Hurst, Early fixation of an optimal genetic code. *Mol Biol Evol* **17**, 511–518 (2000).
- [95] T Tlusty, A rate-distortion scenario for the emergence and evolution of noisy molecular codes. *Phys Rev Lett* **100**, 048101 (2008).
- [96] JJ Hopfield, Kinetic proofreading: A new mechanism for reducing errors in biosynthetic processes requiring high specificity. *Proc Natl Acad Sci (USA)* **71**, 4135–4139 (1974).
- [97] J Ninio, Kinetic amplification of enzyme discrimination. *Biochimie* **57**, 587–595 (1975).
- [98] M Ehrenberg and CG Kurland, Costs of accuracy determined by a maximal growth rate constraint. *Q Rev Biophys* **17**, 45–82 (1984).
- [99] D MacKay and WS McCulloch, The limiting information capacity of a neuronal link. *Bull Math Biophys* **14**, 127–135 (1952).
- [100] HB Barlow, Sensory mechanisms, the reduction of redundancy, and intelligence. In in *Proceedings of the Symposium on the Mechanization of Thought Processes, volume 2*, DV Blake and AM Utley, eds, pp 537–574 (HM Stationery Office, London, 1959).
- [101] HB Barlow, Possible principles underlying the transformation of sensory messages. In *Sensory Communication*, W Rosenblith, ed, pp 217–234 (MIT Press, Cam-

- bridge, 1961).
- [102] EP Simoncelli and BA Olshausen, Natural image statistics and neural representation. *Annu Rev Neurosci* **24**, 1193–1216 (2001).
- [103] WS Geisler, Visual perception and the statistical properties of natural scenes. *Annu Rev Psych* **59**, 167–192 (2008).
- [104] SB Laughlin, A simple coding procedure enhances a neuron’s information capacity. *Z Naturforsch* **36c**, 910–912 (1981).
- [105] F Rieke, D Warland, and W Bialek, Coding efficiency and information rates in sensory neurons. *Europhys Lett* **22**, 151–156 (1993).
- [106] A Borst and FE Theunissen, Information theory and neural coding. *Nat Neurosci* **2**, 947–957 (1999).
- [107] RR de Ruyter van Steveninck, GD Lewen, SP Strong, R Koberle, and W Bialek, Reproducibility and variability in neural spike trains. *Science* **275**, 1805–1808 (1997).
- [108] K Koch, J McLean, R Segev, MA Freed, MJ Berry II, V Balasubramanian, and P Sterling, How much the eye tells the brain. *Curr Biol* **16**, 1428–1434 (2006).
- [109] P Reinagel and RC Reid, Temporal coding of visual information in the thalamus. *J Neurosci* **20**, 5392–5400 (2000).
- [110] RC Liu, S Tzonev, S Rebrük, and KD Miller, Variability and information in a neural code of the cat lateral geniculate nucleus. *J Neurophysiol* **86**, 2789–2806 (2001).
- [111] F Rieke, DA Bodnar, and W Bialek, Naturalistic stimuli increase the rate and efficiency of information transmission by primary auditory neurons. *Proc R Soc Lond Ser B* **262**, 259–265 (1995).
- [112] I Nemenman, GD Lewen, W Bialek, and RR de Ruyter van Steveninck, Neural coding of a natural stimulus ensemble: Information at sub-millisecond resolution. *PLoS Comput Biol* **4**, e1000025 (2008).
- [113] GT Buračas, AM Zador, MR DeWeese, and TD Albright, Efficient discrimination of temporal patterns by motion-sensitive neurons in primate visual cortex. *Neuron* **20**, 959–969 (1998).
- [114] P Kara, P Reinagel, and RC Reid, Low response variability in simultaneously recorded retinal, thalamic, and cortical neurons. *Neuron* **27**, 635–646 (2000).
- [115] DL Ruderman and W Bialek, Statistics of natural images: Scaling in the woods. *Phys Rev Lett* **73**, 814–817 (1994).
- [116] S Smirnakis, MJ Berry II, DK Warland, W Bialek, and M Meister, Adaptation of retinal processing to image contrast and spatial scale. *Nature* **386**, 69–73 (1997).
- [117] T Gollisch and M Meister, Eye smarter than scientists believed: Neural computations in circuits of the retina. *Neuron* **65**, 15–164 (2010).
- [118] B Wark, BN Lundstrom, and A Fairhall, Sensory adaptation. *Curr Opin Neurobiol* **17**, 423–429 (2007).
- [119] N Brenner, W Bialek, and R de Ruyter van Steveninck, Adaptive rescaling optimizes information transmission. *Neuron* **26**, 695–702 (2000).
- [120] AL Fairhall, GD Lewen, W Bialek, and RR de Ruyter van Steveninck, Efficiency and ambiguity in an adaptive neural code. *Nature* **412**, 787–792 (2001).
- [121] MN Kvale and CE Schreiner, Short-term adaptation of auditory receptive fields to dynamic stimuli. *J Neurophysiol* **91**, 604–612 (2004).
- [122] I Dean, NS Harper, and D McAlpine, Neural population coding of sound level adapts to stimulus statistics. *Nature Neurosci* **8**, 1684–1689 (2005).
- [123] KI Nagel and AJ Doupe, Temporal processing and adaptation in the songbird auditory forebrain. *Neuron* **21**, 845–859 (2006).
- [124] B Wen, GI Wang, I Dean, and B Delgutte, Dynamic range adaptation to sound level statistics in the auditory nerve. *J Neurosci* **29**, 13797–13808 (2009).
- [125] JC Rahmen, P Keating, FR Nodal, AL Schulz, and AJ King, Adaptation to stimulus statistics in the perception and neural representation of auditory space. *Neuron* **66**, 937–948 (2010).
- [126] NC Rabinowitz, BDB Willmore, JWH Schnup, and AJ King, Contrast gain control in auditory cortex. *Neuron* **70**, 1178–1192 (2011).
- [127] M Maravall, RS Petersen, AL Fairhall, E Arabzadeh, and ME Diamond, Shifts in coding properties and maintenance of information transmission during adaptation in barrel cortex. *PLoS Biology* **5**, e19 (2007).
- [128] W De Baene, E Premereur, and R Vogels, Properties of shape tuning of macaque inferior temporal neurons examined using rapid serial visual presentation. *J Neurophysiol* **97**, 2900–2916 (2007).
- [129] B Wark, A Fairhall, and F Rieke, Timescales of inference in visual adaptation. *Neuron* **61**, 750–761 (2009).
- [130] JJ Atick and AN Redlich, Toward a theory of early visual processing. *Neural Comp* **2**, 308–320 (1990).
- [131] SB Laughlin and RR de Ruyter van Steveninck, Measurements of signal transfer and noise suggest a new model for graded transmission at an adapting retinal synapse. *J Physiol (Lond)* **494**, P19 (1996).
- [132] RR de Ruyter van Steveninck and SB Laughlin, The rate of information transfer at graded potential synapses. *Nature* **379**, 642–645 (1996).
- [133] L Chittka and R Menzel, The evolutionary adaptation of flower colours and the insect pollinators’ colour vision. *J Comp Physiol A* **171**, 171–181 (1992).
- [134] P Garrigan, CP Ratliff, JM Klein, P Sterling, DH Brainard, and V Balasubramanian, Design of a trichromatic cone array. *PLoS Comput Biol* **6**, e1000677 (2010).
- [135] CP Ratliff, BG Borghuis, Y-H Kao, P Sterling, and V Balasubramanian, Retina is structured to process an excess of darkness in natural scenes. *Proc Natl Acad Sci (USA)* **107**, 17368–17373 (2010).
- [136] BG Borghuis, CP Ratliff, RG Smith, P Sterling, and V Balasubramanian, Design of a neuronal array. *J Neurosci* **28**, 3178–3189 (2008).
- [137] YS Liu, CF Stevens, and TO Sharpee, Predictable irregularities in retinal receptive fields. *Proc Natl Acad Sci (USA)* **106**, 16499–16514 (2009).
- [138] EC Smith and MS Lewicki, Efficient coding of time-relative structure using spikes. *Neural Comp* **17**, 19–45 (2005).
- [139] EC Smith and MS Lewicki, Efficient auditory coding. *Nature* **439**, 978–982 (2006).
- [140] JH van Hateren, A theory of maximizing sensory information. *Biol Cybern* **68**, 23–29 (1992).
- [141] BA Olshausen and DJ Field, Emergence of simple-cell receptive field properties by learning a sparse code for natural images. *Nature* **381**, 607 (1996).
- [142] AJ Bell and TJ Sejnowski, The “independent components” of natural scenes are edge filters. *Vision Res* **23**, 3327–3338.
- [143] R Linsker, How to generate ordered maps by maximiz-

- ing the mutual information between input and output signals. *Neural Comput* **1**, 402–411 (1989).
- [144] H Barlow, Redundancy reduction revisited. *Network* **12**, 241–253 (2001).
- [145] JL Puchalla, E Schneidman, RA Harris, MJ Berry II, Redundancy in the population code of the retina. *Neuron* **46**, 493–504.
- [146] G Tkačik, O Marre, D Amodei, E Schneidman, W Bialek, and MJ Berry II, Searching for collective behavior in a large network of sensory neurons. *PLoS Comput Biol* **10**, e1003408 (2014).
- [147] Y Karklin and E Simoncelli, Efficient coding of natural images with a population of noisy linear-nonlinear neurons. *Adv Neural Info Proc Sys* 999–1007 (2011).
- [148] V Balasubramanian, D Kimber, and MJ Berry II, Metabolically efficient information processing. *Neural Comput* **13**, 799–815.
- [149] G Tkačik, JS Prentice, V Balasubramanian, and E Schneidman, Optimal population coding by noisy spiking neurons. *Proc Natl Acad Sci (USA)* **107**, 14419–14424 (2010).
- [150] LC Osborne, SE Palmer, SG Lisberger, and W Bialek, The neural basis for combinatorial coding in a cortical population response. *J Neurosci* **28**, 13522–13531 (2008).
- [151] T Hafting, M Fyhn, S Molden, MB Moser, and EI Moser, Microstructure of a spatial map in the entorhinal cortex. *Nature* **436**, 801–806 (2005).
- [152] IR Fiete, Y Burak, and T Brookings, What grid cells encode about rat position. *J Neurosci* **28**, 6856–6871 (2008).
- [153] S Sreenivasan and IR Fiete, Grid cells generate an analog error-correcting code for singularly precise neural computation. *Nature Neurosci* **14**, 1330–1337 (2011).
- [154] MB Elowitz, AJ Levine, ED Siggia, and PD Swain, Stochastic gene expression in a single cell. *Science* **297**, 1183–1186 (2002).
- [155] E Ozbudak, M Thattai, I Kurtser, AD Grossman, and A van Oudenaarden, Regulation of noise in the expression of a single gene. *Nature Gen* **31**, 69–73 (2002).
- [156] WJ Blake, M Kaern, CR Cantor, and JJ Collins, Noise in eukaryotic gene expression. *Nature* **422**, 633–637 (2003).
- [157] G Tkačik, CG Callan Jr, and W Bialek, Information capacity of genetic regulatory elements. *Phys Rev E* **78**, 011910 (2008).
- [158] G Tkačik and AM Walczak, Information transmission in genetic regulatory networks: a review. *J Phys Condens Matt* **23**, 153102 (2011).
- [159] T Gregor, DW Tank, EF Wieschaus, and W Bialek, Probing the limits to positional information. *Cell* **130**, 153–164 (2007).
- [160] G Tkačik, CG Callan Jr, and W Bialek, Information flow and optimization in transcriptional regulation. *Proc Natl Acad Sci (USA)* **105**, 12265–12270 (2008).
- [161] L Wolpert, Positional information and the spatial pattern of cellular differentiation. *J Theor Biol* **25**, 1–47 (1969).
- [162] JO Dubuis, G Tkačik, EF Wieschaus, T Gregor, and W Bialek, Positional information, in bits. *Proc Natl Acad Sci (USA)* **110**, 16301–16308 (2013).
- [163] G Tkačik, JO Dubuis, MD Petkova, and T Gregor, Positional information, positional error, and read-out precision in morphogenesis: a mathematical framework. arXiv:1404.5599 [q-bio.MN] (2014).
- [164] P François and ED Siggia, Predicting embryonic patterning using mutual entropy fitness and in silico evolution. *Development* **137**, 2385–2395 (2010).
- [165] G Tkačik, AM Walczak, and W Bialek, Optimizing information flow in small genetic networks. *Phys Rev E* **80**, 031920 (2009).
- [166] AM Walczak, G Tkačik, and W Bialek, Optimizing information flow in small genetic networks. II: Feed-forward interaction. *Phys Rev E* **81**, 041905 (2010).
- [167] G Tkačik, AM Walczak, and W Bialek, Optimizing information flow in small genetic networks. III. A self-interacting gene. *Phys Rev E* **85**, 041903 (2012).
- [168] E Ziv, I Nemenman, and CH Wiggins, Optimal signal processing in small stochastic biochemical networks. *PLoS One* **2**, e1007 (2007).
- [169] F Tostevin and PR ten Wolde, Mutual information between input and output trajectories of biochemical networks. *Phys Rev Lett* **102**, 21801 (2009).
- [170] WH de Ronde, F Tostevin, PR ten Wolde, Effect of feedback on the fidelity of information transmission of time-varying signals. *Phys Rev E* **82**, 031914 (2010).
- [171] CG Bowsher and PS Swain, Environmental sensing, information transfer, and cellular decision-making. *Curr Opin Biotech* **28**, 149–155 (2014).
- [172] I Lestas, G Vinnicombe, and J Paulsson, Fundamental limits on the suppression of molecular fluctuations. *Nature* **467**, 174–178 (2010).
- [173] JL Kelly, Jr, A new interpretation of information rate. *Bell Sys Tech J* **35**, 917–926 (1956).
- [174] CT Bergstrom and M Lachmann, The fitness value of information. arXiv:q-bio.PE/0510007 (2005).
- [175] EL Kussell and S Leibler, Phenotypic diversity, population growth, and information in fluctuating environments. *Science* **309**, 2075–2078 (2005).
- [176] O Rivoire and S Leibler, The value of information for populations in varying environments. *J Stat Phys* **142**, 1124–1166 (2011).
- [177] SF Taylor, N Tishby, and W Bialek, Information and fitness. arXiv:0712.4382 [q-bio.PE] (2007).
- [178] M Vergassola, E Villermaux, and BI Shraiman, “Infotaxis” as a strategy for searching without gradients. *Nature* **445**, 406–409 (2007).
- [179] A Mafra-Neto and RT Cardé, Fine-scale structure of pheromone plumes modulates upwind orientation of flying moths. *Nature* **369**, 142–144 (1994).
- [180] F Attneave, Some informational aspects of visual perception. *Psych Rev* **61**, 183–193 (1954).
- [181] CE Shannon, Prediction and entropy of written English. *Bell Sys Tech J* **30**, 50–64 (1951).
- [182] N Chomsky, Three models for the description of language. *IRE Trans Inf Theory* **IT-2**, 113–124 (1956).
- [183] F Pereira, Formal grammar and information theory: Together again? *Phil Trans R Soc Lond A* **358**, 1239–1253 (2000).
- [184] FC Pereira, N Tishby, and L Lee, Distributional clustering of English words. In *31st Annual Meeting of the Association for Computational Linguistics*, LK Schubert, ed, pp 183–190 (1993).
- [185] N Tishby, FC Pereira, and W Bialek, The information bottleneck method. In *Proceedings of the 37th Annual Allerton Conference on Communication, Control and Computing*, B Hajek and RS Sreenivas, eds, pp 368–377 (University of Illinois, 1999); arXiv:physics/0004057

- (2000).
- [186] W Bialek, RR de Ruyter van Steveninck, and N Tishby, Efficient representation as a design principle for neural coding and computation. arXiv:0712.4381 [q-bio.NC] (2007).
- [187] W Bialek, I Nemenman, and N Tishby, Predictability, complexity and learning. *Neural Comp* **13**, 2409–2463 (2001).
- [188] W Bialek, I Nemenman, and N Tishby, Complexity through nonextensivity. *Physica A* **302**, 89–99 (2001).
- [189] F Creutzig, A Globerson, and N Tishby, Past–future information bottleneck in dynamical systems. *Phys Rev E* **79**, 041925 (2009).
- [190] SE Palmer, O Marre, MJ Berry II, and W Bialek, Predictive information in a sensory population. arXiv.org:1307.0225 [q-bio.NC] (2013).