Information: currency of life?

Daniel Polani¹

¹Adaptive Systems Research Group, School of Computer Science, University of Hertfordshire,
1 College Lane, Hatfield, Hertfordshire AL10 9AB, UK

(Received 10 April 2009; accepted 15 June 2009; published online 8 September 2009)

In biology, the exception is mostly the rule, and the rule is mostly the exception. However, recent results indicate that known universal concepts in biology such as the genetic code or the utilization of ATP as a source of energy may be complemented by a large class of principles based on Shannon's concept of information. The present position paper discusses various promising pathways toward the formulation of such generic informational principles and their relevance for the realm of biology. [DOI: 10.2976/1.3171566]

CORRESPONDENCE

Daniel Polani: d.polani@herts.ac.uk

In considering biological systems, one is confronted with an enormous amount of complexity and diversity, governed—as it appears—by exceptions rather than the rule. Under such circumstances, finding generalizing principles that can be applied across a wide range of phenomena to be described is a daunting task. One of the success stories is, of course, the high degree of universality found in genetic code; on the other hand, the fact that it is not fully universal suggests that it may have resulted from a selection principle (Wong, 1976; Vetsigian *et al.*, 2006) and thus that the resulting code could, in turn, reflect deeper principles about the function and origin of the proteins to be coded.

This is one prominent example that biological processes need not be entirely without generalizable structure which not only organizes what the biological organism can or should do, but which also can be “read” and understood by scientific investigation.

MOTIVATION

Compared, e.g., to physics, established universality principles in biology are far and few in between. If one compares biology and physics, the success of physics in obtaining expressive general statements are, among others, due to its following properties: (1) its op-

eration in a space with simple mathematical rules, and a significant degree of symmetries; (2) the presence of simple bookkeeping principles, such as conservation of momentum and of energy/mass. In physics one finds that such principles, whenever they are seen to fail, can be successfully extended to save the bookkeeping principles by including suitable quantities that were hitherto excluded (e.g., the discovery that heat is also a form of energy or that energy can be carried by hitherto unobserved particles). Such a bookkeeping principle provides powerful predictive power. (3) Formalisms to factor out unknowns or unknowables in a controlled fashion (such as exhibited in different ways in statistical physics or in quantum physics).

We shall now consider the possible relevance of these points for attempting a similar approach in biology. Concerning point (1): for the space of biological operations one is far from being able to formulate a simple mathematical structure. Symmetries, while emerging in particular special cases (such as the bilateral symmetry of many organisms or the regular patterns of snail shells (Hoyle, 2006), are not—as far as is known today—inherent to the fundamental levels of biological operation, although there are attempts to under-

stand the structure of the genetic code from symmetry group considerations (Hornos and Hornos, 1993).

As for point 2, bookkeeping principles, of course, biological organisms have to obey the laws of momentum and of energy conservation; in addition, there is a separate conservation of mass and many other separate conservation laws that organisms obey (e.g., under the constraints of organic life, one can assume that there is no conversion between chemical elements, so one has stoichiometric conservation). Furthermore, a larger number of reactions may be excluded for energetic or other reasons (consider, e.g., the processing of dextro- or levorotatory sugars). One can imagine a larger number of hitherto unknown bookkeeping laws that govern organisms. Due to the multitude of such conserved quantities in biology, as compared to physics, relevant studies are again dominated by special cases. This, in turn, strongly dilutes the potential for universal predictions based on these conservation principles.

Finally, for point (3), the relative homogeneity of physical systems enables in many cases a description that “hides away” the details of a system into a statistical description. Only some dominant parameters are extracted, which provide a general high-level formulation of the system (Haken, 1983; Shalizi, 2001). One of the defining properties of biological systems is, however, their inhomogeneity, which goes along with their complexity and makes the identification of dominant descriptive parameters difficult.

The genetic code was already mentioned earlier as a concept largely universal to biology. Another one is represented by adenosine triphosphate (ATP) as the universal carrier of energy. More precisely, ATP is a carrier of free energy, allowing the organism to carry out useful amount of (physical) work consuming a given amount of ATP molecules (Avery, 2003). Biological processes obtain their free energy in the form of ATP. It turns out that this observation implies another universal bookkeeping quantity of central relevance for biology.

TOWARD PRINCIPLES

The concept of free energy distinguishes between different “qualities” of energy. According to the laws of physics, only “high-quality” energy can be exploited to carry out useful work. This distinction arises from the observation that certain types of energy (specifically heat energy), while obeying the basic energy conservation law (2) mentioned in the first section, cannot be exploited to create (organized) work. Physics provides a systematic treatment for distinguishing energy that is useable from energy that is unusable for such a purpose. The usable component is what is commonly denoted by (Gibbs) free energy.

The origin of this phenomenon arises from the fact that the usable (free) energy in a thermodynamic system requires a certain macroscopic coordination of the system to allow its exploitation. If the system is uncoordinated, heat energy in

the system cannot be used to create macroscopic work. This lack of coordination is captured in classical thermodynamics by the concept of entropy, which models the lack of coordination (“disorder”) of a physical system (Reichl, 1980). Moreover, it turns out that this concept of disorder is complementary to the concept of Shannon information that quantifies how much data can be processed by a given communication channel (Shannon, 1949; Jaynes, 1957a, 1957b; Adami, 1998).

In turn, this implies a subtle relation between free energy and the amount of information that a system can process. If the free energy difference of a reaction is given by ΔG , taking place at temperature T , then a simple calculation allows one to express $-\Delta G/T$ in units of bits per molecule; as example, for the metabolization of glucose into carbon dioxide and water at room temperature one obtains an information-processing capability of 1.7×10^3 bits/molecule (Avery, 2003).

This calculation, of course, represents the ideally achievable case. In a realistic situation only part of this information-processing capability can be actualized. Nevertheless, this makes clear that the information-processing capacity is tightly integrated with the available metabolic energy resources. As it turns out, information acquisition and processing indeed constitute a considerable component of an organism’s metabolism. Examples include the fly eye, which consumes about 10% of a fly’s metabolic energy (Laughlin *et al.*, 1998), as well as the human brain which can consume 20% of the human’s resting energy (Kandel *et al.*, 1991). These amounts are non-negligible and emphasize that information acquisition and processing need to be worth the metabolic investment. Already early in the heyday of cybernetic thought, the fundamental role of information has been suspected (Attneave, 1954; Barlow, 1959).

Now, information is not a conserved quantity such as energy. Nevertheless, a number of useful bookkeeping laws can be formulated for it. One is the data processing inequality, which essentially states that the input information processed over a chain of information-processing units can progressively degrade, as it is handed down the chain, but never again recover (unless it is reinforced or fed in again later on from another source). As a special case of this, any information that an organism has about its environment must have reached it through one or more of its sensors.

A second important statement is that, for given information-processing mechanisms, it is possible to quantify in absolute numbers the maximum amount of information that can be transmitted through this mechanism: this quantity is the channel capacity, measured in bits per time unit. The ability to formulate such a quantity in a consistent and well-defined way is a central result of Shannon’s original work (Shannon, 1949). In particular, for the physical processes underlying sensoric detection, it is possible to obtain estimates for the value of the channel capacity. It is further-

more possible to bound the control capability of an organism by the available sensoric information, which is expressed by the so-called “Law of Requisite Variety” and its variations (Ashby, 1956; Touchette and Lloyd, 2000, 2004). In one of its forms, the law of requisite variety states that (under certain conditions on the dynamics of the environment) to reduce entropy in the environment by a certain amount, an organism has to first acquire that amount of information from the environment.

Armed with this insight, one can now proceed to study how well biological sensors perform in comparison to the theoretical optimum dictated by physics and by information theory. Strikingly, it turns out that biological sensors often indeed operate close to the possible limit: the human ear can come close to the channel capacity constraint imposed by thermal noise (Denk and Webb, 1989). An adapted human eye can react to only small photon clusters (Hecht *et al.*, 1942), and, as a more extreme example, toad receptors can even identify individual photons (Baylor *et al.*, 1979).

These selected examples are quite typical and allow two important conclusions: first, the ability of organisms to acquire and process information is indeed limited by the fundamental laws of physics and information rather than by biological constraints on possible mechanisms. Second, they strongly indicate that information is a resource of prime importance for organisms, and that it trades off with the available metabolic energy (Laughlin, 2001). The above examples highlight sensoric information acquisition, but the principle behind this may well extend beyond it to the whole biological information-processing cascade (Bialek *et al.*, 2007). This suggests a parsimony principle, which proposes that there is little unused information transmission capacity in organisms (Polani *et al.*, 2007).

The rationale of the parsimony principle is that if organisms would develop a suboptimal information-processing strategy, this would waste metabolic energy. Such a disadvantage would then be selected against by evolution. On the other hand, when operating close to the information capacity limits, minor variations have the potential to explore possible evolutionary advantages of increased capacity, in which case the species would exhibit a slow climb on the information/energy cost trade-off curve over time.

Taking this hypothesis as a basis turns information into a quantity that fulfills two of the three criteria listed in the beginning of the paper, suggested as instrumental for the success of physics in providing universal statements: criterion (2), the formulation of bookkeeping principles and criterion (3), the ability to factorize out unknown components of the system while still retaining the capability to make general statements about the rest.

MODELING WITH INFORMATION

“Bookkeeping” principles

The bookkeeping principle applies in a straightforward manner: information acquisition by an organism increases the or-

ganization of its knowledge about the environment, and this has, as discussed above, an immediate correspondence in the free energy necessary for its achievement. Furthermore, the information acquisition process is tied to the organism’s sensors. Together with the parsimony principle, the bookkeeping principle then puts a constraint on possible information-processing architectures of an organism.

Depending on the statistics of the signals impinging on the sensors, the requirement of parsimony constitutes a significant organizing principle. An example is the biologically plausible form of the receptive fields emerging from Linsker’s mathematical infomax principle (Linsker, 1988), or the conceptual probabilistic mappings emerging purely through information flow maximization (Klyubin *et al.*, 2007).

According to the parsimony principle, an organism that realizes an evolutionarily successful behavior will at the same time attempt to minimize the required sensoric information to achieve this behavior to minimize unnecessary processing cost. Conversely, an increase in information acquisition capacity may permit behaviors which confer an advantage to the organism (e.g., in form of a quantitative utility or fitness); in that case, the organism (or species) will be able to slowly increase its utility at the cost of increased information processing. Finally, if one externally imposes a constraint on the amount of information that the organism can acquire or process, the achieved utility will necessarily decrease until it balances with the best utility that can be achieved under this informational constraint.

Bearing in mind that the principle is proposed to be universal, we will illustrate it with the help of a highly simplified abstract model example. Consider the grid world in Fig. 1. Starting out from any spot in this simple grid world, an organism needs to find the marked spot, e.g., a food source. The world is limited by a wall at the boundaries of the grid. We now assume that each step the organism has to take until it reaches the food spot entails a fixed cost (in the following,

Figure 1. Model grid with target spot.

we realize this as a negative utility value of -1). Ideally, the organism would like to minimize this cost (maximize its utility). Assume now that the organism can select as an action a single move in one of the four main directions (north, east, south, and west only—in the example, diagonal moves and turns are not permitted), one can formulate optimal behaviors for this model organism, namely, “zigzagging” from the starting location of the organism toward the food spot, using only moves closing in toward the food spot.

It turns out that, at each step, the organism needs to take in a certain minimum amount of information at each step about its state, i.e., location to find the food in the optimal way described above. If one limits the information acquisition capacity of the organism, e.g., if the organism acquires only blurred signals and does not have a clear idea of whether it is above or below the food, the optimal behavior cannot be realized, and any behavior using this impaired sensory information will be suboptimal. Still, one can ask for the best possible strategy under this constraint. In the most extreme case, the organism is entirely blind and cannot take in any information when deciding its appetitive move, as long as it is not at the food spot. Even in this case it is possible for the organism to find the food spot, namely, by random walk, but at the price of a significantly longer average search length (i.e., significantly more negative utility).

Quantitative treatment

For this model we now study quantitatively the trade-off between the average information to be acquired by the organism per time step and the achieved utility. More precisely we quantify the minimum amount of information that the organism needs to acquire per step to achieve a desired utility for its strategy. The reader less interested in the details of the quantitative methodology is invited to proceed directly to the results for this model in the next subsection.

Assume that the state, i.e., in our case the location of the organism, is described by a random variable S . In a given state, the organism then selects (not necessarily deterministically) an action, described by a random variable A . The information about the state S that is reflected in the choice of action A is then measured with Shannon’s mutual information given by

$$I(S;A) = \sum_{s,a} p(s,a) \log \frac{p(s,a)}{p(s)p(a)}, \quad (1)$$

where the sum runs over all possible concrete states s , and action a , $p(s)$ is the probability that a particular state s will be assumed, $p(a)$ is the probability that the action a will be selected, and $p(s,a)$ is the probability that a particular state s will be realized jointly with action a . If the logarithm is chosen for base 2, the mutual information is measured in bits. The action selection now in fact depends on the current state. This is captured by concept of a strategy. We denote by π the strategy of the (memoryless) organism, i.e., the condi-

tional probability $p(a|s)$ that the action a being taken if the organism’s state (i.e., location) is s . If the organism happens to be in a state s with probability $p(s)$, it follows for the joint distribution of states and actions that $p(s,a) = p(a|s)p(s)$. With this, a given strategy π implies that action A carries the amount of information $I(S;A)$ about the state S given by Eq. (1).

In the next step, we need to connect the selection of actions in a given state to their desirability. For this, we utilize the formalism from Polani *et al.* (2006), which adapts the relevant information concept from Tishby *et al.* (1999) to qualify the selection of actions. Consider a function $U(s,a)$ that denotes the utility of choosing an action a in a state s . If an organism happens to be in a state s with probability $p(s)$ and follows a strategy $\pi \equiv p(a|s)$, then the mean utility of the organism is given by $\mathbf{E}^\pi[U(s,a)] = \sum_{s,a} p(a|s)p(s)U(s,a)$.

To implement the principle of parsimony, one can now ask for the strategy π that minimizes $I(S;A)$ for a given mean utility level. If one has a fixed utility U , this is achieved by solving the minimization problem (details in Tishby *et al.*, 1999 and Polani *et al.*, 2006)

$$\min_{\pi} (I(S;A) - \beta \cdot \mathbf{E}^\pi[U(S,A)]). \quad (2)$$

As the values for β are swept from 0 toward infinity, the minimization singles out strategies π which minimize $I(S;A)$ for various given mean utility levels. It is not possible to give a closed expression for how a general value of β relates to the mean utilities. However, as β tends toward 0, the information $I(S;A)$ tends to 0, i.e., actions are increasingly independent from the state, while still trying to maximize the utility achievable for the given value of the information $I(S;A)$. On the other hand, as β tends toward infinity, strategies π are sought which are optimal with respect to U , but still minimize the information $I(S;A)$ required to achieve that optimal utility.

In our model, the utility derives from the cost of -1 for each step taken before the goal is reached. The utility therefore is not fixed but depends on the path that the organism takes. This path, in turn, depends on the strategy π of the organism. Thus, the utility is no longer fixed, but it itself depends on the strategy π , i.e., we have to replace U with U^π . Although this utility computation now involves averaging over all possible trajectories of the organism through the future under policy π , an elegant reduction is possible using reinforcement learning models (Sutton and Barto, 1998). One can write

$$U^\pi(s,a) = \sum_{s'} p(s'|s,a) [r(s,a,s') + V^\pi(s')], \quad (3)$$

where s and a are the current state and action of the organism, $p(s'|s,a)$ is the probability that action a in state s will move the organism to a successive state s' , $r(s,a,s')$ is the reward for this step which in our case, as mentioned above, is always realized as a penalty of -1 , and finally $V^\pi(s')$ is the

value of the successor state (i.e., the cumulated reward when starting in state s'), if strictly following strategy π . For the latter term, it can now be shown that the recursion relation holds (Bellmann equation)

$$V^\pi(s) = \sum_a p(a|s) \sum_{s'} p(s'|s, a) [r(s, a, s') + V^\pi(s')]. \quad (4)$$

With the strategy-dependent utility given by Eq. (3), the optimal trade-off between utility and information needs to simultaneously fulfill Eqs. (2) and (4). These trade-off curves can be systematically computed (Polani *et al.*, 2006) and will be discussed in the “Scenario results” section.

Scenario results

Figure 2 shows the resulting optimal trade-off curve between utility and information; more precisely, it shows the utility achieved by the best strategy that requires no more than the given amount of information per time step. On the horizontal axis, the required information is plotted on the vertical axis, the best utility (i.e., the negative cost) that can be achieved at this given information intake level. Higher values of utility are more desirable. Every combination not above this curve is achievable.

There are two special cases: the rightmost spot on the trade-off curve corresponds to the information use for the optimal strategy. The leftmost spot on the trade-off curve denotes the best utility that can be achieved by an organism which is entirely blind, i.e. which does not utilize any information at all. In between these two extreme cases, the trade-off curve indicates the best utility achievable when a certain informational bandwidth is available, or, read the other way around, the minimum information required to achieve a certain utility.

For closer discussion, consider first the case where the organism uses the optimal strategy. In this case, if the organism starts randomly on one of the spots in Fig. 1 except for

Figure 2. Trade-off between information (horizontal axis, in bit/step) and utility (vertical axis, denoted as negative values—the more negative, the more costly the behavior; a value of -40 means an average path length of 40 steps until reaching the goal).

the target spot itself, the optimal strategy achieves an average utility of -3.5, i.e., it requires 3.5 time steps on average to reach the target spot. The rightmost point of the trade-off curve now specifies the minimal amount of information required by the agent to select actions realizing such an optimal strategy. This amount is per time step, i.e., for each individual action selection. The following simplifying assumptions are made in this calculation: the organism is memoryless and does not keep track of its actions selected in earlier time steps, i.e., the actions are independently selected at each time; furthermore, it is assumed that the distribution of states with which the organism is confronted is not changed by the actions, but that it remains an equilibrium distribution.

As one moves towards the left on the trade-off curve, one increasingly trades achievable utility for savings in required information until one reaches the leftmost spot on the curve, which marks the highest achievable utility without information intake. This corresponds to an entirely blind organism for which the best achievable utility drops off to approximately -70, indicating that under the best possible strategy it needs on average 70 steps to reach the goal. However, the slope of the trade-off curve at that spot is virtually vertical, so that already with only a very small amount of information a considerable improvement in performance can be achieved; for instance, for as little as 0.1 bits/step, the optimally achievable performance improves to an average of only about 12 time steps to the goal. Thus, under the hypothesis that biological systems indeed adapt toward the principle of parsimony, they already may profit from very small informational bandwidths.

More importantly, one should note that the information bookkeeping property ensures that it is completely irrelevant how the organism acquires its information about where the food spot is, whether per chemotaxis, per optical signals or a different sensoric modality. In fact, the information trade-off is so fundamental that it also would hold for artificial robot implementations attempting the same task. In all cases, given an average cost (path length), the minimal information it has to acquire about the food position per single step is universally determined by the curve in Fig. 2. The particular sensors that the organism possesses may or may not provide this information, but in no case can the organism possibly do better than the trade-off curve in Fig. 2. They could do worse, though, i.e., achieving a lower utility or requiring more information, i.e., the solutions below the curve in Fig. 2 are viable, whereas solutions above the curve are not.

The principle of parsimony predicts that, for utility measures that are central for survival, organisms will evolve as to operate close to the trade-off curve separating the viable from the nonviable region. Where exactly they will lie on this curve depends on the balance between the metabolic/fitness cost for maintaining a particular (sensory) information channel and the utility (i.e., evolutionary) advantage that this

channel confers. Indeed, recent experiments measuring growth rates of organisms show that such trade-offs between fitness and information acquisition can be seen (Taylor *et al.*, 2007). That principle has been suggested to extend to population level (Bergstrom and Lachmann, 2004).

THE CHALLENGE OF UNIVERSALITY

Universal utilities

All aforementioned results have in common that they emerge without requiring an intricate knowledge of the mechanisms behind them. The results do not depend on the detailed mechanisms that guide the dynamics of decision. The same limitations hold for biological systems as for artificially constructed ones. In analogy to the energy concept in physics, the bookkeeping property of information allows one to hide away the details of the inner workings of a particular system to be able to make nontrivial general statements. In spirit, this is similar to the idea behind thermodynamics in physics, which provides an explicit framework that separates everything that is known about a system from everything that is unknown about it and nevertheless allows for quantitative statements.

In doing so, there are several points that deserve particular notice. In thermodynamics, the key property that makes quantitative statements possible about statistical systems is the second law. In Jaynes' celebrated formulation, it is essentially equivalent to the statement that a system adopts a probability distribution over states for which entropy is maximized (Jaynes, 1957a,b). In the language of earlier sections, entropy maximization corresponds to information minimization. Stated informally, this corresponds to minimizing the assumptions about the given system beyond the known parameters and constraints.

What would correspond to Jaynes' maximum entropy assumption in our earlier informational picture of biology? If one remembers that Fig. 2 shows the best possible trade-off between an achieved utility and the information utilized, then any realizable system needs to lie anywhere on or below that separator line. There is—from an informational point of view—no limitation on where below this line the system could lie in principle. However, if we interpret the indicators discussed in the “Toward Principles” section, the parsimony principle hypothesizes that biological systems actually operate at the very limit of the viability region, and indeed close to the optimality line.

The quantities in question (entropy in physics, information in biology) are related. The parsimony principle, however, is clearly different from physics. The entropy maximization principle in physics makes a statement about a maximum lack of information by an external observer under given material and physical constraints; it is also at the core of more recent variants of that principle, such as the maximum entropy production principle (Dewar, 2003; Martyushev and Seleznev, 2006). On the other hand, how-

ever, the hypothesized information parsimony principle for biology rather says something about the best possible information acquisition by an organism (under the constraints of metabolism and drives in question).

Beyond the fundamental limits, the information parsimony principle does not impose any restrictions on the substrate implementing the loop. In particular, it does not make detailed statements about the mechanisms realizing the principle. However, it is implied that the adaptation processes on different time scales (this includes the long-term evolutionary process, but also lifetime and more short-term adaptation/learning dynamics) “conspire” as to result in parsimonious information processing.

The parsimony principle, in addressing the amount of information necessary to achieve a particular utility, is implicitly a measure of the cost per time to process the sensoric information for generating a desired behavior. Thus, it essentially measures the cost for the complexity of the particular task a behavior is addressing. Some utilities and the associated tasks or drives turn out to be informationally cheap, others may be more expensive to achieve.

The above approaches assume the existence of utilities that can be quantified. There are various situations when this is directly possible, such as for specific given tasks or (on the population level) in studying growth rates, as in Taylor *et al.* (2007). Is it possible to address this, however, in the case of generic behaviors when canonical utilities cannot be identified?

As a way of doing so, rather than considering particular tasks or goals of an organism, recent work has begun to the general properties and abilities of the perception-action loop of a given organism (Fig. 3). Here, the information picture can again provide a suitable framework.

In Bialek *et al.*, 2001 and subsequent work, the authors propose to consider predictive information, i.e., the information that the sensoric past of an organism carries about the sensoric future. They suggest this as a quantity that an organism aims to maximize. Not being tied to a particular utility, this measure is generic in the sense that it depends only on the dynamic of the organism in the environment. The proposed maximization of predictive information can be interpreted as poising the organisms in states in the environment

Figure 3. Structure of the perception-action loop of an organism.

with a high richness of successor states, which are, however, at the same time predictable from the organism's point of view. This principle has been used in artificial agents to produce a rich set of self-motivated agent behaviors without any further assumptions about the ecological niche of the agents (Ay *et al.*, 2008). The behavior dictated by predictive information maximization is determined purely by how the organism interacts with the environment, but not by any specific drives. In fact, it has been hypothesized as the underlying generic principle for the degeneration of drives, prior to any organism-specific needs.

An alternative generic approach to derive fundamental candidates for organismic drives from properties of the perception-action loop emphasizes the organism as an entity that is able to select its actions. It considers the informational channel capacity between the organism's actions at a given time and the sensoric inputs at a later time (Klyubin *et al.*, 2005a,b). Intuitively, this is a measure to which extent the organism's actions could potentially influence its environment (i.e., how much information the organism could potentially "inject" into the environment) in such a way that the intensity of this influence can later be detected again by the organism. This quantity, empowerment, measures the organism's power to change the environment and to be aware that it did so. This can be formalized by measuring the maximal mutual information that can possibly be introduced into the environment by a suitable distribution of actions.

Quantitative models of empowerment

To discuss empowerment quantitatively, we discuss its simplest form. Consider an agent that starts out at time $t'=0$ in the state $S_0=s$. If the agent can potentially perform a single action (denoted by a random variable A_0), and then observe the following state S_1 , then the empowerment is given by $\mathfrak{E}(s) = \max_{p(A_0)} I(A_0; S_1 | S_0=s)$, where the maximum is taken over all possible distributions $p(A_0)$ of the selected A_0 . This can be readily generalized to consider not single actions, but action sequences $A_0, A_1, \dots, A_{t-1}$ of length t and their influence on the state observed after this time, leading to t -step empowerment

$$\mathfrak{E}^t(s) = \max_{p(A_0, A_1, \dots, A_{t-1})} I(A_0, A_1, \dots, A_{t-1}; S_t | S_0=s). \quad (5)$$

Further generalizations are possible (Klyubin *et al.*, 2008).

Empowerment essentially measures the informational efficiency of the sensorimotor loop. In a similar vein to the parsimony principle or the predictive information maximization, the empowerment maximization hypothesis states that an organism behaves in a way that attempts to maximize its empowerment. Various scenarios demonstrate that using the empowerment maximization principle as a drive indeed provides intuitive and plausible behavior patterns (Klyubin *et al.*, 2008). Unlike the parsimony principle, which immediately trades off a given cost versus information, the empowerment maximization hypothesis is less direct; however,

apart from providing intuitively desirable behaviors, there are some reasons to suggest that either the principle itself (or a related principle) may indeed hold.

To obtain an intuition how empowerment typically operates, we adapt an example from Klyubin *et al.* (2005a). Consider an (infinitely sized) grid world with an organism inside. An action of the organism is again one step to the north, east, south, or west, and we include also a "stop" action, which does not move the organism. We now consider empowerment for $t=5$ steps, i.e., the organism can select a sequence of actions freely for times $t'=0, \dots, 4$. The empowerment $\mathfrak{E}^t(s)$ then measures how much these various action sequences can at most affect the state of the organism. In the special case of a deterministic world $\mathfrak{E}^t(s)$ evaluates simply to $\log N$, where N is the number of all states that can be reached at the end of the action sequence. This reflects the achievable richness of system states that can be reached in the given time horizon. In the above example with five time steps, the empowerment is given by $\mathfrak{E}^5(s) = \log_2 61 \approx 5.9$ bit everywhere, since, in total, 61 states can be reached in five time steps. This is the amount of information that can potentially be injected by the organism into the environment through this action sequence.

We now modulate the scenario by placing a box at the center position (0,0) in the grid world, and the organism can sense not just its own location, but also the location of the box. We consider two cases: the box can be either immovable or movable. First, consider the immovable case. Figure 4, top, shows the empowerment on the vertical axis for various positions on the grid, with the box in the middle. Far away from the box, the empowerment is approximately 5.9 bits, the same as for the empty world mentioned above, as the box is not encountered during the travel of the organism. How-

Figure 4. Empowerment for immovable and movable box. The horizontal plane denotes the organism location, and the vertical axis gives the empowerment value. Note that for better visibility the diagrams use different scale; the empowerment differences for the immovable box are much smaller than for the movable.

ever, as the organism starts closer to the box, the box hems in the freedom of action of the organism. An increasing number of possible target states become unreachable and the empowerment drops, as the organism starts out increasingly closer to the box. On the box itself, the empowerment is high again, as the organism in its first move is free to stay on the box or fall off it in any direction, again allowing it to reach the same number of states as in the free-moving case. Thus, except for the central spot, there is a slight advantage to stay away from the box, which takes away some degrees of freedom (the empowerment differences are only small, about 0.2 bit).

Now, consider the scenario of the movable box. The scenario is changed in such a way that if the agent is beside the box and takes a move toward it, it thereby pushes it in the direction of its movement. The rest of the dynamics remains the same. The empowerment for this case is shown in Fig. 4, bottom. Here, the situation is entirely different. An agent far from the box achieves the usual 5.9 bit, however, close to the box, the organism cannot just influence its own location, but also manipulate the box. This increased influence on the world is immediately reflected in the empowerment value, and a considerable gain in empowerment can be achieved by the organism by placing itself close to the box (or on top of it, falling down at the first moving action). Thus, empowerment identifies directly the advantage of being close to an additional manipulable degree of freedom in the environment.

Empowerment is a fully organism-subjective measure, and an organism which cannot sense the location of the box will not identify the empowerment advantage it provides. If, in turn, the additional degrees of freedom provided by some feature in the environment (in above example, the box) would offer an evolutionary advantage to the organism, this would imply a selection pressure toward the evolution and subsequent adaptation of sensors able to detect that feature. The hypothesis is then that empowerment, as a measure for the efficiency of the perception-action loop, provides a virtually immediate feedback about its viability, prior to any further detailed assumptions about the necessary survival dynamics for the organism.

DISCUSSION

Consider an organism that does not fully exploit the sensorimotor capacities available to it by evolutionary adaptation, i.e., that either constrains the use of its actuators or whose sensors cannot detect the full range of possible actuator dynamics. This organism (or species) could then afford to degenerate either sensorics or actuators without loss, as its behavioral range does not exploit the available sensorimotor equipment or vice versa. Empowerment provides a direct measure for the quality of this adaptation, and as long as it is not affected by modifying actuators and sensors, adaptation will favor the least costly sensorimotor equipment among those with a given empowerment value.

Consider, in turn, an organism with an evolutionary his-

tory that led it to a particular phenotypic outcome and sensorimotor equipment. This implies that the particular make-up of the organism has proven to be viable with relation to the cost of sensorics and actuators. For this to come to full fruition, however, it is plausible to assume that during its lifetime the organism should attempt to seek a niche where its sensorics and actuators are indeed exploitable to the fullest extent. Note, however, that empowerment is a potential quantity. It quantifies the possibility to exploit the sensorimotor loop, not the actual richness of the behavior (unlike the predictive information mentioned earlier). Stated differently, according to the empowerment maximization principle, the organism will poise itself to maximize its options (in a way that is discernible by the organism itself), but not necessarily its actual actions.

Empowerment, measured by the information-theoretic quantity of sensorimotor channel capacity, is a candidate for a universal utility for the organism which does not require particular organism-specific knowledge for its formulation. The empowerment maximization principle can be interpreted in various ways for different organisms: attempting to maximize metabolic reserves to place oneself into a rich ecological niche, to place oneself in a situation of being able to select among various mating partners, and more. While the details depend on the individual level of consideration, the principle is universal. Furthermore, not only does the principle provide a generic way of deriving drives and generic behavior rules, but also a framework for the prediction of sensorimotor designs that would be advantageous to evolve in a given niche; simple examples for a model system are given in Klyubin *et al.* (2005b, 2008).

The empowerment maximization principle can also be interpreted in another way: an organism should poise itself in such a way as to be able to react in a most effective way to possible perturbations of its favorite state. The higher an organism's empowerment, the better the "power" of the organism to control potential perturbations. Since the advent of the cybernetic picture, it has been hypothesized that homeostatic principles guide the stability of organismic dynamics on various levels (Ashby, 1956). The problem with this picture is that one needs to identify beforehand the homeostatic variables whose stabilization is necessary for the organism. These variables need to be rediscovered anew for each new type of organism. A principle, such as the empowerment maximization principle, however, provides a path for a first-principle derivation of candidates for homeostatic variables (Klyubin *et al.*, 2008), derived just from the properties of the given sensorimotor niche. In other words, essential drives (which may differ for different organisms) may be derived directly from a fundamental principle, which is essentially the same for all organisms.

CONCLUSION AND OUTLOOK

In the current paper we have discussed various information-theoretic principles as candidates for understanding the in-

formation ecology of organisms. Among these were the parsimony principle and the predictive information and empowerment maximization principles. There are other related principles—for instance, a free energy principle based on Bayesian models applied to the brain has been introduced in [Friston et al. \(2006\)](#). Possibly in the future, it will be possible to identify informational principles which will encompass and unify several of the above.

For this purpose, it is necessary to identify experimental settings to validate the various informational hypotheses about organismic sensorimotor dynamics and behavior. Developing appropriate experimental scenarios is ongoing work, but there are already developments suitable for such an undertaking. For instance, there are clear indications that human behavior can be explained under the assumption of Bayesian models underlying the decision process ([Körding and Wolpert, 2004](#)). Bayesian models are optimal in the sense that they make best use of the available information. Similarly, an “infotaxis” principle based on optimal information gain indicates the emergence of behaviors that correspond to the observed “coasting” behavior of moths in search of a mating partner ([Vergassola et al., 2007](#)).

To test for the validity of the parsimony principle, behavioral experiments will have to investigate goal-directed scenarios where a utility is defined. In experiments with humans, one would then investigate the deterioration of the effectiveness by which a goal is attained for varying degrees of information bandwidth (where the bandwidth is limited to various extent, e.g., by multiple simultaneous or distracting tasks) and compare the results with the theoretical trade-off curve. Another scenario is to consider taxis behavior observed in various biological scenarios and quantify the trade-off between utilized information and performance.

For the informational optimality principles of predictive information or empowerment, various tests can be conceived. The most direct ones could be behavioral experiments with humans. Here one could conceive tasks where a target is presented surprisingly and needs to be reached quickly, without and with distracting/perturbing subtasks or constraints. Such an experimental procedure would allow one to test the validity of the various principles. In particular, it would allow distinguishing predictive information and empowerment, or possibly a combination of these as possible principles obeyed by the observed behaviors. Here, one would expect that untrained test subjects’ behavior would come closer to behavior governed by predictive information optimality principles due to its exploratory character, while subjects more experienced in a scenario would move toward empowerment-informed behaviors, which optimize response to perturbations. Finally, a comparison of predictions from informational criteria with existing movement models, such as minimum jerk or the two-thirds power law ([Viviani and Flash, 1995](#)), could provide additional instrumental links to established results.

There is mounting evidence for the importance of information as a fundamental currency underlying the success of living organisms. The informational picture implies the existence of various quantitative constraints on an adapted organism’s possible sensoric make-ups, information-processing strategies, and behaviors. If it is indeed possible to firmly establish these constraints, this will, in turn, indicate a significant universality among the various information-processing mechanisms that abound in the realm of biology. This, finally, would introduce a new level of quantitative predictiveness into biology and lead to qualitatively novel insights on the principles that drive living organisms.

REFERENCES

- Adami, C (1998). *Introduction to Artificial Life*, Springer, New York.
- Ashby, WR (1956). *An Introduction to Cybernetics*, Chapman & Hall, London.
- Attneave, F (1954). “Informational aspects of visual perception.” *Psychol. Rev.* **61**, 183–193.
- Avery, J (2003). *Information Theory and Evolution*, World Scientific, Singapore.
- Ay, N, Bertschinger, N, Der, R, Güttler, F, and Olbrich, E (2008). “Predictive information and explorative behavior of autonomous robots.” *Eur. Phys. J. B* **63**, 329–339.
- Barlow, HB (1959). “Possible principles underlying the transformations of sensory messages.” In *Sensory Communication: Contributions to the Symposium on Principles of Sensory Communication*, Rosenblith, WA, ed., pp. 217–234, MIT, Cambridge, MA.
- Baylor, D, Lamb, T, and Yau, K (1979). “Response of retinal rods to single photons.” *J. Physiol. (London)* **288**, 613–634.
- Bergstrom, CT, and Lachmann, M (2004). “Shannon information and biological fitness.” In *Information Theory Workshop*, pp. 50–54, IEEE, Piscataway, NJ.
- Bialek, W, de Ruyter van Steveninck, RR, and Tishby, N (2007). “Efficient representation as a design principle for neural coding and computation.” arXiv.org: 0712.4381.
- Bialek, W, Nemenman, I, and Tishby, N (2001). “Predictability, complexity and learning.” *Neural Comput.* **13**, 2409–2463.
- Denk, W, and Webb, WW (1989). “Thermal-noise-limited transduction observed in mechanosensory receptors of the inner ear.” *Phys. Rev. Lett.* **63**(2), 207–210.
- Dewar, R (2003). “Information theory explanation of the fluctuation theorem, maximum entropy production and self-organized criticality in non-equilibrium stationary states.” *J. Phys. A* **36**(3), 631–641.
- Friston, K, Kilner, J, and Harrison, L (2006). “A free energy principle for the brain.” *J. Physiol. (Paris)* **100**, 70–87.
- Haken, H (1983). *Advanced Synergetics*, Springer, Berlin.
- Hecht, S, Schlaer, S, and Pirenne, M (1942). “Energy, quanta and vision.” *J. Opt. Soc. Am.* **38**, 196–208.
- Hornos, JEM, and Hornos, YMM (1993). “Algebraic model for the evolution of the genetic code.” *Phys. Rev. Lett.* **71**(26), 4401–4404.
- Hoyle, R (2006). *Pattern Formation*, Cambridge University Press, Cambridge, MA.
- Jaynes, ET (1957a). “Information theory and statistical mechanics. I.” *Phys. Rev.* **106**(4), 620–630.
- Jaynes, ET (1957b). “Information theory and statistical mechanics. II.” *Phys. Rev.* **108**(2), 171–190.
- Kandel, ER, Schwartz, JH, and Jessell, TM (1991). *Principles of Neural Science*, 3rd Ed., McGraw-Hill, New York.
- Klyubin, A, Polani, D, and Nehaniv, C (2007). “Representations of space

- and time in the maximization of information flow in the perception-action loop." *Neural Comput.* **19**(9), 2387–2432.
- Klyubin, AS, Polani, D, and Nehaniv, CL (2005a). "All else being equal be empowered." *Advances in Artificial Life, European Conference on Artificial Life (ECAL 2005)*, LNAI, Springer, Vol. 3630, 744–753.
- Klyubin, AS, Polani, D, and Nehaniv, CL (2005b). "Empowerment: A universal agent-centric measure of control." *Proc. IEEE Congress on Evolutionary Computation (CEC 2005)*, Edinburgh, Scotland, 128–135.
- Klyubin, AS, Polani, D, and Nehaniv, CL (2008). "Keep your options open: an information-based driving principle for sensorimotor systems." *PLoS ONE* **3**(12), e4018.
- Körding, KP, and Wolpert, DM (2004). "Bayesian integration in sensorimotor learning." *Nature (London)* **427**, 244–247.
- Laughlin, SB (2001). "Energy as a constraint on the coding and processing of sensory information." *Curr. Opin. Neurobiol.* **11**, 475–480.
- Laughlin, SB, de Ruyter van Steveninck, RR, and Anderson, JC (1998). "The metabolic cost of neural information." *Nat. Neurosci.* **1**(1), 36–41.
- Linsker, R (1988). "Self-organization in a perceptual network." *Computer* **21**(3), 105–117.
- Martyushev, LM, and Seleznev, VD (2006). "Maximum entropy production principle in physics, chemistry and biology." *Phys. Rep.* **426**, 1–45.
- Polani, D, Nehaniv, C, Martinetz, T, and Kim, JT (2006). "Relevant information in optimized persistence vs. progeny strategies." *Proc. Artificial Life X*, Rocha, LM, Bedau, M, Floreano, D, Goldstone, R, Vespignani, A, and Jaeger, L, eds., 337–343.
- Polani, D, Sporns, O, and Lungarella, M (2007). "How information and embodiment shape intelligent information processing." *Proc., 50th Anniversary Summit of Artificial Intelligence*, Lungarella, M, Iida, F, Bongard, J, and Pfeifer, R, eds., Springer, Berlin, 99–111.
- Reichl, L (1980). *A Modern Course in Statistical Physics*, University of Texas Press, Austin, TX.
- Shalizi, CR (2001). "Causal architecture, complexity and self-organization in time series and cellular automata." Ph.D. thesis, University of Wisconsin-Madison, Madison, WI.
- Shannon, CE (1949). "The mathematical theory of communication." In *The Mathematical Theory of Communication*, Shannon, CE, and Weaver, W, eds., The University of Illinois Press, Urbana, IL.
- Sutton, RS, and Barto, AG (1998). *Reinforcement Learning*, MIT Press, Cambridge, MA.
- Taylor, SF, Tishby, N, and Bialek, W (2007). "Information and fitness." arXiv.org: 0712.4382.
- Tishby, N, Pereira, FC, and Bialek, W (1999). "The information bottleneck method." *Proc., 37th Annual Allerton Conference on Communication, Control and Computing*, Urbana-Champaign, IL.
- Touchette, H, and Lloyd, S (2000). "Information-theoretic limits of control." *Phys. Rev. Lett.* **84**, 1156–1159.
- Touchette, H, and Lloyd, S (2004). "Information-theoretic approach to the study of control systems." *Physica A* **331**, 140–172.
- Vergassola, M, Villermanx, E, and Shraiman, BI (2007). "'Infotaxis' as a strategy for searching without gradients." *Nature (London)* **445**, 406–409.
- Vetsigian, K, Woese, C, and Goldenfeld, N (2006). "Collective evolution and the genetic code." *Proc. Natl. Acad. Sci. U.S.A.* **103**(28), 10696–10701.
- Viviani, P, and Flash, T (1995). "Minimum-jerk, two-thirds power law, and isochrony: converging approaches to movement planning." *J. Exp. Psychol. Hum. Percept. Perform.* **21**(1), 32–53.
- Wong, JT-F (1976). "The evolution of a universal genetic code." *Proc. Natl. Acad. Sci. U.S.A.* **73**(7), 2336–2340.