

Information-based fitness and the emergence of criticality in living systems

Jorge Hidalgo^{a,1}, Jacopo Grilli^{b,1}, Samir Suweis^b, Miguel A. Muñoz^{a,2}, Jayanth R. Banavar^c, and Amos Maritan^{b,2}

^aDepartamento de Electromagnetismo y Física de la Materia and Instituto Carlos I de Física Teórica y Computacional, Universidad de Granada, 18071 Granada, Spain; ^bDipartimento di Fisica “G. Galilei”, Consorzio Nazionale Interuniversitario per le Scienze Fisiche della Materia and Istituto Nazionale di Fisica Nucleare, Università di Padova, 35131 Padua, Italy; and ^cDepartment of Physics, University of Maryland, College Park, MD 20742

Edited* by William Bialek, Princeton University, Princeton, NJ, and approved May 27, 2014 (received for review October 12, 2013)

Empirical evidence suggesting that living systems might operate in the vicinity of critical points, at the borderline between order and disorder, has proliferated in recent years, with examples ranging from spontaneous brain activity to flock dynamics. However, a well-founded theory for understanding how and why interacting living systems could dynamically tune themselves to be poised in the vicinity of a critical point is lacking. Here we use tools from statistical mechanics and information theory to show that complex adaptive or evolutionary systems can be much more efficient in coping with diverse heterogeneous environmental conditions when operating at criticality. Analytical as well as computational evolutionary and adaptive models vividly illustrate that a community of such systems dynamically self-tunes close to a critical state as the complexity of the environment increases while they remain noncritical for simple and predictable environments. A more robust convergence to criticality emerges in coevolutionary and coadaptive setups in which individuals aim to represent other agents in the community with fidelity, thereby creating a collective critical ensemble and providing the best possible tradeoff between accuracy and flexibility. Our approach provides a parsimonious and general mechanism for the emergence of critical-like behavior in living systems needing to cope with complex environments or trying to efficiently coordinate themselves as an ensemble.

evolution | adaptation | self-organization

Physical systems undergo phase transitions from ordered to disordered states on changing control parameters (1, 2). Critical points, with all their remarkable properties (1, 2), are only observed upon parameter fine tuning. This is in sharp contrast to the ubiquity of critical-like behavior in complex living matter. Indeed, empirical evidence has proliferated that living systems might operate at criticality (3)—i.e. at the borderline between order and disorder—with examples ranging from spontaneous brain behavior (4) to gene expression patterns (5), cell growth (6), morphogenesis (7), bacterial clustering (8), and flock dynamics (9). Even if none of these examples is fully conclusive and even if the meaning of “criticality” varies across these works, the criticality hypothesis—as a general strategy for the organization of living matter—is a tantalizing idea worthy of further investigation.

Here we present a framework for understanding how self-tuning to criticality can arise in living systems. Unlike models of self-organized criticality in which some inanimate systems are found to become critical in a mechanistic way (10), our focus here is on general adaptive or evolutionary mechanisms, specific to biological systems. We suggest that the drive to criticality arises from functional advantages of being poised in the vicinity of a critical point.

However, why is a living system fitter when it is critical? Living systems need to perceive and respond to environmental cues and to interact with other similar entities. Indeed, biological systems constantly try to encapsulate the essential features of the huge variety of detailed information from their surrounding complex

and changing environment into manageable internal representations, and they use these as a basis for their actions and responses. The successful construction of these representations, which extract, summarize, and integrate relevant information (11), provides a crucial competitive advantage, which can eventually make the difference between survival and extinction. We suggest here that criticality is an optimal strategy to effectively represent the intrinsically complex and variable external world in a parsimonious manner. This is in line with the hypothesis that living systems benefit from having attributes akin to criticality—either statistical or dynamical (3)—such as a large repertoire of dynamical responses, optimal transmission and storage of information, and exquisite sensitivity to environmental changes (2, 5, 12–16).

As conjectured long ago, the capability to perform complex computations, which turns out to be the fingerprint of living systems, is enhanced in “machines” operating near a critical point (17–19), i.e., at the border between two distinct phases: a disordered phase, in which perturbations and noise propagate unboundedly—thereby corrupting information transmission and storage—and an ordered phase where changes are rapidly erased, hindering flexibility and plasticity. The marginal, critical situation provides a delicate compromise between these two impractical tendencies, an excellent tradeoff between reproducibility and flexibility (12, 13, 16) and, on larger time scales, between robustness and evolvability (20). A specific example of this general framework is genetic regulatory networks (19, 21). Cells ranging from those in complex organisms to single-celled

Significance

Recently, evidence has been mounting that biological systems might operate at the borderline between order and disorder, i.e., near a critical point. A general mathematical framework for understanding this common pattern, explaining the possible origin and role of criticality in living adaptive and evolutionary systems, is still missing. We rationalize this apparently ubiquitous criticality in terms of adaptive and evolutionary functional advantages. We provide an analytical framework, which demonstrates that the optimal response to broadly different changing environments occurs in systems organizing spontaneously—through adaptation or evolution—to the vicinity of a critical point. Furthermore, criticality turns out to be the evolutionary stable outcome of a community of individuals aimed at communicating with each other to create a collective entity.

Author contributions: J.H., J.G., S.S., M.A.M., J.R.B., and A.M. designed research; J.H., J.G., and S.S. performed research; and J.H., J.G., S.S., M.A.M., J.R.B., and A.M. wrote the paper.

The authors declare no conflict of interest.

*This Direct Submission article had a prearranged editor.

¹J.H. and J.G. contributed equally to this work.

²To whom correspondence may be addressed. Email: mamunoz@onsager.ugr.es or amos.maritan@pd.infn.it.

This article contains supporting information online at www.pnas.org/lookup/suppl/doi:10.1073/pnas.1319166111/-DCSupplemental.

microbes such as bacteria respond to signals in the environment by modifying the expression of their genes. Any given genetic regulatory network, formed by the genes (nodes) and their interactions (edges) (22), can be tightly controlled to robustly converge to a fixed almost-deterministic attractor—i.e. a fixed “phenotype”—or it can be configured to be highly sensitive to tiny fluctuations in input signals, leading to many different attractors, i.e., to large phenotypic variability (23). These two situations correspond to the ordered and disordered phases, respectively. The optimal way for genetic regulatory networks to reconcile controllability and sensitivity to environmental cues is to operate somewhere in between the two limiting and impractical limits alluded to above (19) as has been confirmed in different experimental setups (5, 7, 24). Still, it is not clear how such tuning to criticality comes about.

Our goal here is to exploit general ideas from statistical mechanics and information theory to construct a quantitative framework showing that self-tuning to criticality is a convenient strategy adopted by living systems to effectively cope with the intrinsically complex external world in an efficient manner, thereby providing an excellent compromise between accuracy and flexibility. To provide some further intuition, we use genetic regulatory networks as a convenient guiding example, but one could equally well consider neural networks, models for the immune response, groups of animals exhibiting collective behavior, etc., with each specific realization requiring a more detailed modeling of its special attributes.

We uncover coevolutionary and coadaptive mechanisms by which communities of living systems, even in the absence of other forms of environmental complexity, converge to be almost critical in the process of understanding each other and creating a “collective entity.” The main result is that criticality is an evolutionary/adaptive stable solution reached by living systems in their striving

to cope with complex heterogeneous environments or when trying to efficiently coordinate themselves as an ensemble.

Results

Mathematical Framework. The external environment in which living systems operate is highly variable, largely unpredictable, and describable in terms of probability distribution functions. Living systems need to modify their internal state to cope with external conditions, and they do so in a probabilistic manner. To be specific, but without loss of generality, we represent an environmental cue “perceived” and processed by a living system as a string of N (binary) variables, $\mathbf{s} = (s_1, s_2, \dots, s_N)$. A specific environmental source is modeled by the probability distribution P_{src} with which it produces each of the 2^N possible states. For concreteness, this distribution is assumed to depend on a set of parameters, $\alpha = (\alpha_1, \alpha_2, \dots)$, accounting for environmental variability. We turn now to an individual living system or “agent,” which seeks to adapt itself to cope with the perceived stimuli/signals emanating from a given environmental source. This is accomplished by changing its internal state, encapsulated in a second probability distribution function, P_{int} , specified by a different—smaller in principle—parameter set $\beta = (\beta_1, \beta_2, \dots)$ aimed at capturing the essential features of P_{src} in the most efficient—although in general imperfect—way (see Fig. 1). Henceforth we will denote the external source and its internal representation by $P_{\text{src}}(\mathbf{s}|\alpha)$ and $P_{\text{int}}(\mathbf{s}|\beta)$ respectively.

In our guiding example, the external cues could be, for instance, the environmental (temperature, pH, ...) conditions, which are variable and can only be probabilistically gauged by a cell/bacterium. The binary vector $\mathbf{s} = (s_1, s_2, \dots, s_N)$ can be thought of as the on/off state of the different N genes in its (Boolean) genetic regulatory network (19, 21, 22). In this way, $P_{\text{src}}(\mathbf{s}|\alpha)$ can be interpreted as the probability that the most convenient state aimed at by the system to cope with a given

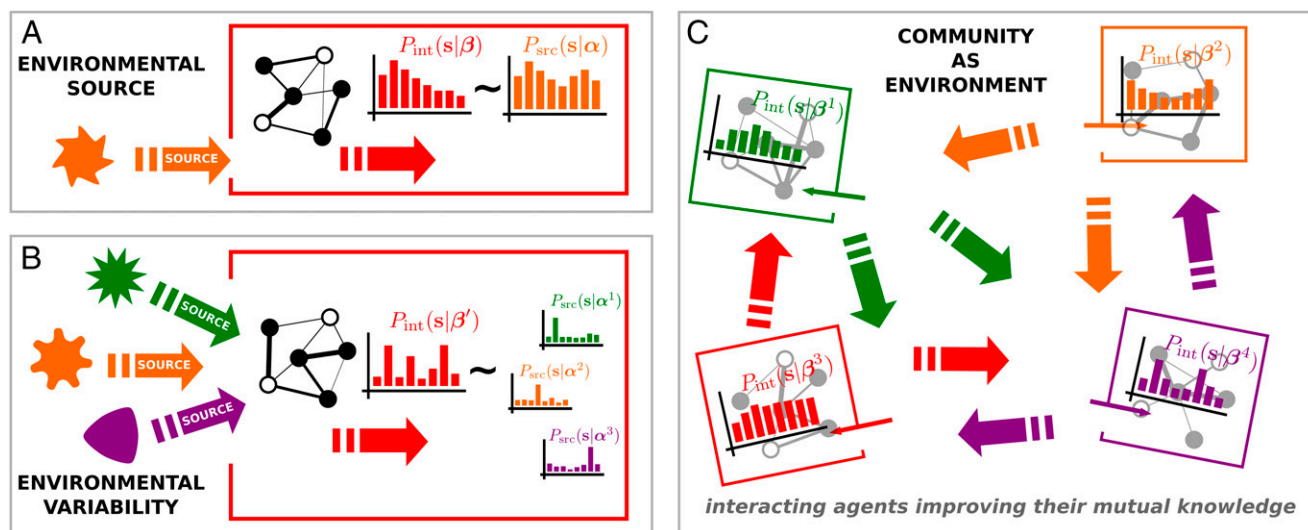


Fig. 1. Living systems coping with the environment. **A** illustrates a living system responding to an environmental source (e.g., a bacteria responding to some external conditions such as the presence/absence of some nutrients, pH concentration, or temperature). A given source, labeled by the set of parameters α , can only be probabilistically gauged by the system. $P_{\text{src}}(\mathbf{s}|\alpha)$ is the most accurate representation that the system can potentially generate in terms of the Boolean variables (or bits) \mathbf{s} . However, such a representation might not be accessible to the system by merely changing its internal state parameters, β , and the actual internal state, $P_{\text{int}}(\mathbf{s}|\beta)$ (e.g., the probability of a gene expression pattern), is usually an imperfect proxy for $P_{\text{src}}(\mathbf{s}|\alpha)$. The optimal choice of parameters β —aiming at capturing the most relevant features of the environment—is obtained by minimizing the KL divergence of $P_{\text{int}}(\mathbf{s}|\beta)$ from $P_{\text{src}}(\mathbf{s}|\alpha)$. In genetic networks, changing internal parameters is equivalent to changing the interactions between the different (Boolean) variables (nodes of the networks in the figure). **B** shows a more complex scenario, where the system has to cope with multiple and diverse sources. The internal state has to be able to accommodate each of them. In **C**, the environment is not imposed ad hoc but, instead, it is composed of other individuals, and every agent needs to cope with (“understand”) the states of the others. Each agent evolves similarly to the others in the community, trying to exhibit the same kind of state, generating in this way a self-organized environment. In the case of sufficiently heterogeneous externally imposed sources as well as in the self-organized case, we find that evolutionary/adaptive dynamics drive the systems to operate close to criticality.

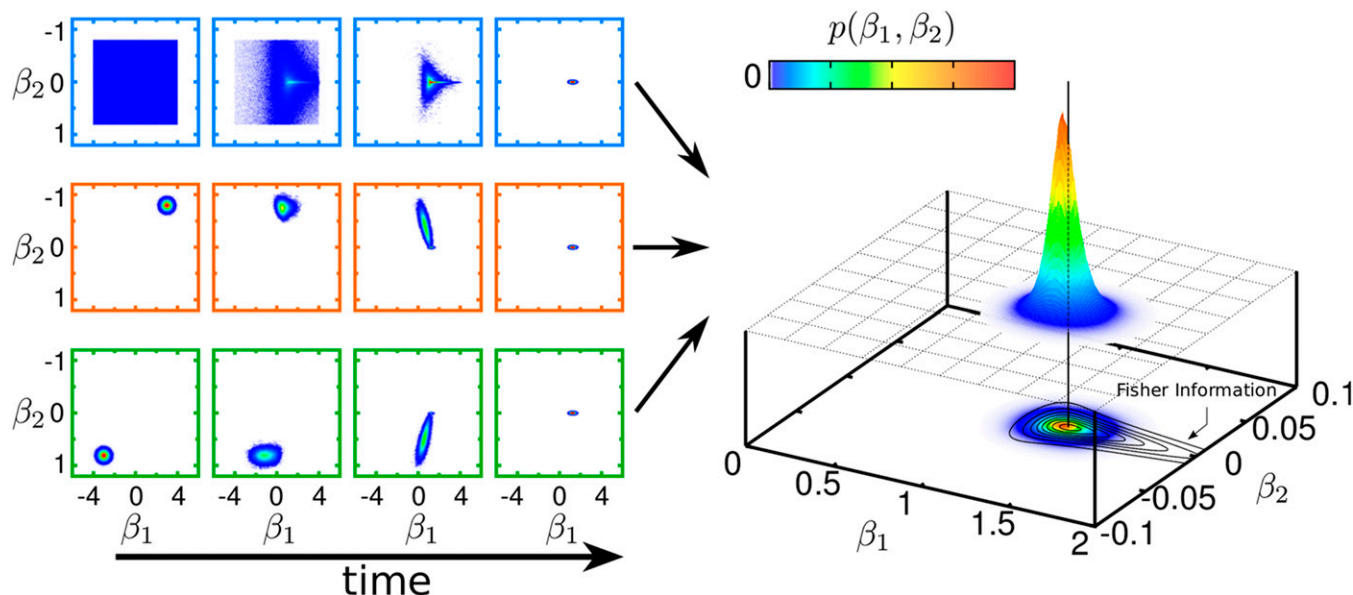


Fig. 2. Coevolutionary model leads self-consistently to criticality: A community of M living systems (or agents) evolves according to a genetic algorithm dynamics (27). Each agent i ($i = 1, \dots, M$) is characterized by a two-parameter (β_1, β_2) internal state distribution $P_{\text{int}}(\mathbf{s}|\beta_1, \beta_2)$, and the rest of the community acts as the external environment it has to cope with, i.e., the agents try to “understand” each other. At each simulation step, two individuals are randomly chosen and their respective relative fitnesses are computed in terms of the KL divergence from each other’s internal state probability distribution. One of the two agents is removed from the community with a probability that is smaller for the fitter agent; the winner produces an offspring, which (except for small variations/mutations) inherits its parameters. (Left) These coevolutionary rules drive the community very close to a unique localized steady state. As shown (Right), this is localized precisely at the critical point, i.e., where the generalized susceptibility or Fisher information of the internal state distribution exhibits a sharp peak (as shown by the contour plots and heat maps). The internal state distributions are parameterized as $P_{\text{int}}(\mathbf{s}|\beta_1, \beta_2) \propto \exp\left\{\beta_1 \frac{N}{2} \left(\sum_{k=1}^N \frac{s_k}{N}\right)^2 + \beta_2 \sum_{k=1}^N s_k\right\}$ representing a global (all-to-all) coupling of the internal nodes (see *Materials and Methods*). Much more complex probability distributions in which all units are not coupled to all other units—i.e., more complex networked topologies—are discussed in *SI Appendix, section S4*.

stimuli: $d(\rho_{\text{src}}|\beta) := \int d\alpha \rho_{\text{src}}(\alpha) D(\alpha|\beta)$. In the special case of just one internal state parameter, β_1 , we find that upon iterating the genetic algorithm (see *Materials and Methods*), the distribution evolves toward a steady-state stable distribution. Computer simulations show that in the case of very homogeneous environments, occurring when all sources in the pool are similar— $\rho_{\text{src}}(\alpha)$ sufficiently narrow—the optimal β strongly depends on the specific sources, resulting in detail-specific internal states (see Fig. 3 Bottom). On the other hand, if the external world is sufficiently heterogeneous (see *SI Appendix, section S5*), the optimal internal state becomes peaked near the critical point (see Fig. 3 and *Movie S3* illustrating the evolution of agents toward the vicinity of the critical point). We note that the approach to criticality is less precise in this model than in the co-evolutionary one, in which the environment changes progressively as the agents coevolve, allowing the system to systematically approach the critical point with precision. Similar conclusions hold for an analogous “adaptive model” (see *SI Appendix, section S6*). Finally, one might wonder whether the resulting closeness to criticality in these models is not just a byproduct of the environment itself being critical in some way. In fact, it has been recently shown that complex environments, when hidden variables are averaged out, can be effectively described by the Zipf’s law (28), a signature of criticality (3). This observation applies to some of the heterogeneous environments analyzed here which, indeed, turn out to be Zipfian; however, as shown in *SI Appendix, section S5*, there are simple yet heterogeneous environments, which are not Zipfian but nevertheless result in the same behavior.

Analytical Results for the Dynamical Models

A generic probability distribution can be rewritten to parallel the standard notation in statistical physics, $P(\mathbf{s}|\gamma) = \exp(-H(\mathbf{s}|\gamma))/Z(\gamma)$, where the factor $Z(\gamma)$ is fixed through normalization. The function H can be generically written as $H(\mathbf{s}|\gamma) = \sum_{\mu} \gamma_{\mu} \phi^{\mu}(\mathbf{s})$,

where $\phi^{\mu}(\mathbf{s})$ are suitable functions (“observables”) of the variables \mathbf{s} . For a specific set of parameters α characterizing an environmental source, the best possible internal state—minimizing the KL divergence—can be shown to obey $\langle \phi^{\mu}_{\text{int}} \rangle_{\alpha} = \langle \phi^{\mu}_{\text{int}} \rangle_{\beta}$, where the index μ runs over the whole set of parameters and $\langle \phi^{\mu}_{\text{int}} \rangle_{\alpha} := \sum_{\mathbf{s}} \phi^{\mu}_{\text{int}}(\mathbf{s}) P_{\text{int}}(\mathbf{s}|\alpha)$ and $\langle \phi^{\mu}_{\text{int}} \rangle_{\beta} := \sum_{\mathbf{s}} \phi^{\mu}_{\text{int}}(\mathbf{s}) P_{\text{int}}(\mathbf{s}|\beta)$. This result implies that the optimal internal state is the one which best reproduces the lowest moments of the original source distribution it seeks to cope with (the number of moments coinciding with—or being limited by—the number of free parameters). By evaluating the second derivatives (Hessian matrix), it is easy to verify that, if a solution exists, it actually corresponds to a minimum of the KL divergence (see *SI Appendix, section S3*).

To proceed further, we need to compute the internal state distribution in the presence of diverse sources distributed with $\rho_{\text{src}}(\alpha)$. In this case, we compute the value of β which minimizes the average KL divergence to the sources α as written above (an alternative possibility—which is discussed in *SI Appendix, section S3*—is to identify the optimal β for each specific source and then average over the source distribution), leading to the condition: $\langle \phi^{\mu}_{\text{int}} \rangle_{\beta} = \int d\alpha \rho_{\text{src}}(\alpha) \langle \phi^{\mu}_{\text{int}} \rangle_{\alpha}$. We consider the simple example in which both the sources and the system are characterized by a single parameter and, assuming that a phase transition occurs at some parameter value $\alpha = \alpha_c$, i.e., $\langle \phi \rangle_{\alpha}$ has a sigmoid shape (which becomes steeper as N increases) with an inflection point at $\alpha = \alpha_c$ (our analysis can be extended to more general cases where there is no built-in phase transition in the source distributions but they are merely sufficiently heterogeneous). The two plateaus of the sigmoid function correspond to the so-called disordered and ordered phases, respectively. When $\rho_{\text{src}}(\alpha)$ has support on both sides of the sigmoid function, i.e., when it is “heterogeneous,” by solving the equation for the optimal β , it is obvious that the moment to be reproduced lies somewhere in between the two asymptotic values of the sigmoid

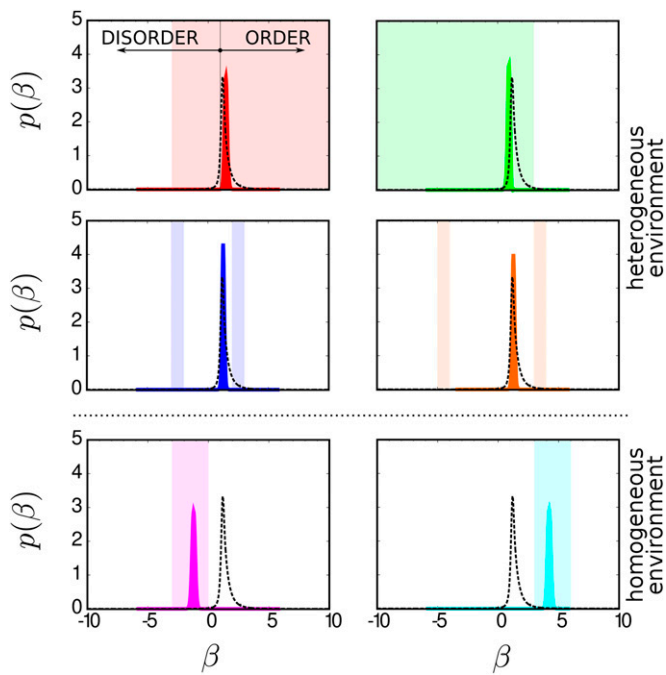


Fig. 3. Evolutionary model leading to near criticality in complex environments. A community of M agents undergoes a genetic algorithm dynamics (27). Each agent is simultaneously exposed to diverse stimuli s provided by S different sources, each one characterized by a probability $P_{src}(s|\alpha^u)$ with $u = 1, \dots, S$, fully specified by parameters α^u . At each time step, S sources are randomly drawn with probability $p_{src}(\alpha^u)$ (in this case, a uniform distribution with support in the colored region). Each agent i ($i = 1, \dots, M$) has an internal state $P_{int}(s|\beta)$ aimed at representing—or coping with—the environment. Agents' fitness increases as the mean KL divergence from the set of sources to which they are exposed decreases. The higher the fitness of an individual, the lower its probability of dying. An agent that is killed is replaced by a new individual with a parameter β inherited from one of the other agents (and, with some probability, a small variation/mutation). The community dynamically evolves and eventually reaches a steady state distribution of parameters, $p(\beta)$. The six panels in the figure correspond to different supports (colored regions) for uniform source distributions, $p_{src}(\alpha^u)$. The dashed line is the generalized susceptibility (Fisher information) of the internal probability distribution, which exhibits a peak at the critical point separating an ordered from a disordered phase. Heterogeneous source pools (Top and Middle) lead to distributions peaked at criticality, whereas for homogeneous sources (Bottom), the communities are not critical but specialized. Stimuli distributions are parametrized in a rather simplistic way as $P_{src}(s|\alpha^u) \propto \exp\left\{\alpha^u \frac{N}{2} \left(\sum_{k=1}^N \frac{s_k}{N}\right)^2\right\}$, while internal states are identical but replacing α^u by β^i (see Materials and Methods). In the guiding example of genetic regulatory networks, this example corresponds to an extremely simple fully connected network in which the state of each gene is equally determined by all of the other genes, and hence the probability of a given state depends only on the total number of on/off genes, controlled by a single parameter.

with the values of β for which intermediate moments are concentrated near the inflection or critical point, α_c . Indeed, as $\chi = -\frac{d\langle\phi\rangle}{d\beta}$, the critical region, where the generalized susceptibility χ has a peak, is the region of maximal variability in which different complex sources can be best accounted for through small parameter changes, in agreement with the finding that many more distinguishable outputs can be reproduced by models poised close to criticality (26).

Discussion and Conclusions

Under the mild assumption that living systems need to construct good although approximate internal representations of the outer

complex world and that such representations are encoded in terms of probability distributions, we have shown—by using concepts from statistical mechanics and information theory—that the encoding probability distributions do necessarily lie where the generalized susceptibility or Fisher information exhibits a peak (25), i.e., in the vicinity of a critical point, providing the best possible compromise to accommodate both regular and noisy signals.

In the presence of broadly different ever-changing heterogeneous environments, computational evolutionary and adaptive models vividly illustrate how a collection of living systems eventually clusters near the critical state. A more accurate convergence to criticality is found in a coevolutionary/coadaptive setup in which individuals evolve/adapt to represent with fidelity other agents in the community, thereby creating a collective “language,” which turns out to be critical.

These ideas apply straightforwardly to genetic and neural networks—where they could contribute to a better understanding of why neural activity seems to be tuned to criticality—but have a broader range of implications for general complex adaptive systems (21). For example, our framework could be applicable to some bacterial communities for which a huge phenotypic (internal state) variability has been empirically observed (29). Such a large phenotypic diversification can be seen as a form of “bet hedging,” an adaptive survival strategy analogous to stock market portfolio management (30), which turns out to be a straightforward consequence of individuals in the community being critical. Usually, from this point of view, generic networks diversify their “assets” among multiple phenotypes to minimize the long-term risk of extinction and maximize the long-term expected growth rate in the presence of environmental uncertainty (30). Similar bet-hedging strategies have been detected in viral populations and could be explained as a consequence of their respective communities having converged to a critical state, maximizing the hedging effect. Similarly, criticality has been recently shown to emerge through adaptive information processing in machine learning, where networks are trained to produce a desired output from a given input in a noisy environment; when tasks of very different complexity need to be simultaneously learned, networks adapt to a critical state to enhance their performance (31). In summary, criticality in some living systems could result from the interplay between their need for producing accurate representations of the world, their need to cope with many widely diverse environmental conditions, and their well-honed ability to react to external changes in an efficient way. Evolution and adaptation might drive living systems to criticality in response to this smart cartography.

Materials and Methods

Kullback–Leibler Divergence. Given two probability distributions $P(s)$ and $Q(s)$ for variables s , the KL divergence of $Q(s)$ from $P(s)$,

$$D(P|Q) := \sum_s P(s) \log \left(\frac{P(s)}{Q(s)} \right), \quad [2]$$

quantifies the loss of information when $Q(s)$ is used to approximate $P(s)$ (25). Indeed, in the large T limit, the probability \mathcal{L} that the model $Q(s)$ generates a sequence of T observations compatible with $P(s)$ can be computed as $\mathcal{L} \sim \exp(-TD(P|Q))$ up to leading order (see SI Appendix, section S2). Therefore, maximizing the likelihood of a trial probability distribution function Q is equivalent to minimizing its KL divergence from the original one, P . In Results we use the notation $D(\alpha|\beta)$ when $P(s) = P_{src}(s|\alpha)$ and $Q(s) = P_{int}(s|\beta)$.

Fisher Information and Criticality. Given a probability distribution $P(s|\gamma)$ —where γ can stand either for α or β —the Fisher Information is defined as

$$\chi^{\mu\nu}(\gamma) := \left\langle \frac{\partial \log P(\cdot|\gamma)}{\partial \gamma_\mu} \frac{\partial \log P(\cdot|\gamma)}{\partial \gamma_\nu} \right\rangle_\gamma, \quad [3]$$

where μ and ν are parameter labels and the average $\langle \cdot \rangle_\gamma$ is performed with respect to $P(\cdot|\gamma)$. It measures the amount of information encoded in the states s about the parameters γ (25). This follows from the Cramér–Rao

inequality, which states that the error made when we estimate γ from one state \mathbf{s} is, on average, greater than (or at least equal to) the inverse of the Fisher information (25). In particular, if χ happens to diverge at some point, it is possible to specify the associated parameters with maximal precision (26). With the parametrization used in the main text, the Fisher information is the generalized susceptibility in the statistical mechanics terminology and measures the response of the system to parameter variations: $\chi^{\mu\nu}(\gamma) = -\frac{\partial \langle \phi^\mu \phi^\nu \rangle_\gamma}{\partial \gamma_\nu} = \langle \phi^\mu \phi^\nu \rangle_\gamma - \langle \phi^\mu \rangle_\gamma \langle \phi^\nu \rangle_\gamma$, and is well known to peak at critical points (1, 2).

Coevolutionary Model. The k th agent of the community is described by a probability distribution $P_{\text{int}}(\mathbf{s}|\beta^k) \propto \exp\{-H_{\text{int}}(\mathbf{s}|\beta^k)\}$, with $H_{\text{int}}(\mathbf{s}|\beta^k) = \sum_{\mu} \beta_{\mu}^k \phi_{\text{int}}^{\mu}(\mathbf{s})$, depending on parameters β^k . Starting with an ensemble of M agents whose internal parameters are extracted from an arbitrary distribution, $p(\beta)$, two individuals, i and j , are randomly selected at each time step. Their relative fitnesses $f_i^{(j)}$ and $f_j^{(i)}$ are computed as $f_i^{(j)} = 1 - D(\beta^i|\beta^j)/[D(\beta^i|\beta^j) + D(\beta^j|\beta^i)]$, and similarly for $f_j^{(i)}$ (as the KL divergence is not symmetric, $f_i^{(j)} \neq f_j^{(i)}$ unless $\beta^i = \beta^j$). One of the two individuals—selected with probability equal to its relative fitness—creates an offspring, while the other one is removed from the community. The offspring inherits its parameters from its ancestor (with probability $1 - \nu$) or mutates with a probability ν , modifying its parameters from β to $\beta \rightarrow \beta + \xi$, where ξ is a multivariate Gaussian random vector, with uncorrelated components, zero mean, and deviation σ . Time is updated to $t \rightarrow t + 1/M$, another couple of individuals i' and j' is picked, and the process is iterated. Variants of this model are described in *SI Appendix, section S4*.

Evolutionary Model. A community of agents receiving external stimuli from an outer and heterogeneous environment is modeled as follows. Every specific environmental source corresponds to a probability distribution $P_{\text{src}}(\mathbf{s}|\alpha) \propto \exp(-H_{\text{src}}(\mathbf{s}|\alpha))$, with $H_{\text{src}}(\mathbf{s}|\alpha) = \sum_{\mu} \alpha_{\mu} \phi_{\text{src}}^{\mu}(\mathbf{s})$, where the parameters α are drawn from the distribution $\rho_{\text{src}}(\alpha)$. The k th agent in the community constructs an internal representation of the observed source described by $P_{\text{int}}(\mathbf{s}|\beta^k) \propto \exp(-H_{\text{int}}(\mathbf{s}|\beta^k))$ with $H_{\text{int}}(\mathbf{s}|\beta^k) = \sum_{\mu} \beta_{\mu}^k \phi_{\text{int}}^{\mu}(\mathbf{s})$, with parameters β^k . We start with M individuals, each one equipped with some initial parameter set extracted from some arbitrary distribution $p(\beta)$. At every time step, we generate S external sources, $\{\alpha^u\}_{u=1, \dots, S}$, from the source pool $\rho_{\text{src}}(\alpha)$. Then we compute the average KL divergence of every individual's internal state distribution from the external sources $d(\{\alpha^u\}|\beta^k) := \sum_{u=1}^S D(\alpha^u|\beta^k)/S$. The k th individual of the community is removed with a probability proportional to its average KL divergence (or any increasing function of it) $P_{\text{kill}}(k) = d(\{\alpha^u\}|\beta^k)/\sum_l d(\{\alpha^u\}|\beta^l)$, and it is replaced by an offspring of another individual randomly selected from the rest of the community. The offspring inherits its parameters from the parent, and time is updated as in the coevolutionary model.

ACKNOWLEDGMENTS. We are indebted to T. Hoang, D. Pfaff, J. Uriagereka, S. Vassanelli, and M. Zamparo for useful discussions and to W. Bialek and two anonymous referees for many insightful suggestions. A.M., J.G., and S.S. acknowledge Cariparo Foundation for financial support. M.A.M. and J.H. acknowledge support from J. de Andalucia P09-FQM-4682 and the Spanish MINECO FIS2009-08451.

- Stanley HE (1987) *Introduction to Phase Transitions and Critical phenomena* (Oxford Univ Press, London).
- Binney J, Dowrick N, Fisher A, Newman M (1993) *The Theory of Critical Phenomena* (Oxford Univ Press, Oxford).
- Mora T, Bialek W (2011) Are biological systems poised at criticality? *J Stat Phys* 144(2): 268–302.
- Beggs JM, Plenz D (2003) Neuronal avalanches in neocortical circuits. *J Neurosci* 23(35):11167–11177.
- Nykter M, et al. (2008) Gene expression dynamics in the macrophage exhibit criticality. *Proc Natl Acad Sci USA* 105(6):1897–1900.
- Furusawa C, Kaneko K (2012) Adaptation to optimal cell growth through self-organized criticality. *Phys Rev Lett* 108(20):208103.
- Krotov D, Dubuis JO, Gregor T, Bialek W (2014) Morphogenesis at criticality. *Proc Natl Acad Sci USA* 111(10):3683–3688.
- Chen X, Dong X, Be'er A, Swinney HL, Zhang HP (2012) Scale-invariant correlations in dynamic bacterial clusters. *Phys Rev Lett* 108(14):148101.
- Bialek W, et al. (2012) Statistical mechanics for natural flocks of birds. *Proc Natl Acad Sci USA* 109(13):4786–4791.
- Jensen HJ (1998) *Self-Organized Criticality: Emergent Complex Behavior in Physical and Biological Systems* (Cambridge Univ Press, Cambridge, UK).
- Edlund JA, et al. (2011) Integrated information increases with fitness in the evolution of animals. *PLOS Comput Biol* 7(10):e1002236.
- Chialvo DR (2010) Emergent complex neural dynamics. *Nat Phys* 6(10):744–750.
- Beggs JM (2008) The criticality hypothesis: How local cortical networks might optimize information processing. *Philos Trans A Math Phys Eng Sci* 366(1864):329–343.
- Kinouchi O, Copelli M (2006) Optimal dynamical range of excitable networks at criticality. *Nat Phys* 2(5):348–351.
- Mora T, Walczak AM, Bialek W, Callan CG, Jr. (2010) Maximum entropy models for antibody diversity. *Proc Natl Acad Sci USA* 107(12):5405–5410.
- Shew WL, Plenz D (2013) The functional benefits of criticality in the cortex. *Neuroscientist* 19(1):88–100.
- Langton C (1990) Computation at the edge of chaos: Phase transitions and emergent computation. *Physica D* 42(1–3):12–37.
- Bertschinger N, Natschläger T (2004) Real-time computation at the edge of chaos in recurrent neural networks. *Neural Comput* 16(7):1413–1436.
- Kauffman S (1993) *The Origins of Order: Self-Organization and Selection in Evolution* (Oxford Univ Press, New York).
- Wagner A (2005) *Robustness and Evolvability in Living Systems* (Princeton Univ Press, Princeton, NJ).
- Gros C (2008) *Complex and Adaptive Dynamical Systems: A Primer* (Springer, New York).
- de Jong H (2002) Modeling and simulation of genetic regulatory systems: A literature review. *J Comput Biol* 9(1):67–103.
- Huang S (2010) Cell lineage determination in state space: A systems view brings flexibility to dogmatic canonical rules. *PLoS Biol* 8(5):e1000380.
- Balleza E, et al. (2008) Critical dynamics in genetic regulatory networks: Examples from four kingdoms. *PLoS ONE* 3(6):e2456.
- Cover TM, Thomas J (1991) *Elements of Information Theory* (Wiley, New York).
- Mastromatteo I, Marsili M (2011) On the criticality of inferred models. *J Stat Mech* 2011:P10012.
- Goldberg DE (1989) *Genetic Algorithms in Search, Optimization, and Machine Learning* (Addison-Wesley Professional, Reading, MA).
- Schwab DJ, Nemenman I, Mehta P (2013) Zipf's law and criticality in multivariate data without fine-tuning. arXiv:1310.0448.
- Kussell E, Leibler S (2005) Phenotypic diversity, population growth, and information in fluctuating environments. *Science* 309(5743):2075–2078.
- Wolf DM, Vazirani VV, Arkin AP (2005) Diversity in times of adversity: Probabilistic strategies in microbial survival games. *J Theor Biol* 234(2):227–253.
- Goudarzi A, Teuscher C, Gulbahce N, Rohlf T (2012) Emergent criticality through adaptive information processing in Boolean networks. *Phys Rev Lett* 108(12): 128702.