Maximal, Enforced and Potential Entropy Production: towards an Understanding of Mechanisms for the Generation of Complexity

Daniel Polani, Philippe Capdepuy and Chrystopher Nehaniv
Algorithms and Adaptive Systems Groups
School of Computer Science
University of Hertfordshire
d.polani@herts.ac.uk

1 Introduction

Recent research has been increasingly emphasizing the central role that entropy plays in the understanding of complex systems. From just providing a "balance sheet", entropy has gone a long way to be used for the characterization of complex systems.

Apart from the fact the entropy provides a fundamental characterization of the "natural" state of a system when it achieves the maximum entropy under given constraints, Jaynes' view of entropy extended this by understanding it as a universal way to determine the "most unbiased" distribution (17; 18).

In physics, the macroscopic variables of a state in thermodynamic equilibrium define constraints under which by maximizing entropy then one determines the microscopic state distribution. However, with the increased interest in nonequilibrium dynamics to understand complex systems (25) and the introduction of the *maximum entropy production principle* (MEPP) (15; 13; 14; 23), one has attained a new tool to treat systems that are not in thermodynamic equilibrium, but that have a (still fixed) in- and outflow of matter, energy or other. The MEPP is, similar to the maximum entropy principle, a statistical principle. Instead of finding the most probable state (or, in the the language of the theory of types, the most *typical* state class (9)), in the MEPP one looks for the most typical trajectory.

This allows to address a number of nonequilibrium scenarios, whenever fixed in- and outflows are given as constraints. This is a generalization of the usual maximum entropy view of thermal equilibrium, and its use is becoming increasingly important also in the fields of Artificial Life and ecosystem analysis (33).

However, there is an alternative view of the MEPP. The special case of systems in dynamic equilibrium (i.e. with constant flows), while thermodynamically not in equilibrium, still embrace a purely statistical approach. The MEPP is a statement about probable trajectories, and thus an ultimately purely statistical phenomenon.

2 Enforced Entropy Production

However, the prevalence of this model suggests another, slightly different question: what if the system actually would *aim* for the maximization of entropy production?

The question arises from turning around the original premise, namely not considering the maximization of entropy (or entropy production) as the result of a statistical process, but the result of a process *aimed* at maximizing the entropy.

To understand why this might be a relevant viewpoint, we have to discuss the conditions under which the maximum entropy and maximum entropy production principles hold. For these principles to be valid, one requires either (for maximum entropy) a system in thermal equilibrium (in particular, no net flows) or (for MEPP) a system in dynamic equilibrium with fixed net in- and outflows (but still temperature needs to be locally well-defined).

The thermal equilibrium condition is microscopically essentially the condition of detailed balance — the net transition rates from a state A to a state B are the same as those from B to A. In naive Newtonian physics, this is enabled by microreversibility which ensures that the concept of temperature is well defined in many relevant cases. However, once one moves away from systems with microreversibility (e.g. in lasers, or in systems with complex dynamics, such as living matter), then the concept of detailed balance and thus of temperature does not have to be well-defined anymore.

In that moment also the validity of the conditions ensuring maximum entropy or MEPP breaks down. The systems in question do not anymore have to fulfil any of these conditions by default. Our question is now whether there could be situations where maximum entropy principles become favoured by the systems, but not as a result of statistics, but rather driven by the systems.

Indeed, if we cross the gap from the microscopic view of physics towards living agents, there are scenarios where the agents will strive to generate entropy in their environment. Consider as example, e.g. a von Neumann game matrix. For a typical zero-sum game the optimal strategy is indeed a probabilistic one. The optimal strategy in these cases is that the opponent has to be left in a state of most uncertainty (minimum information) about the action of the present player. This small example shows that there are indeed situations where the generation of entropy (removal of information) is a favourable property of an adaptive system.

3 Potential Entropy Production

Of course, the example in Sec. 2 is not the only one to be considered. In (20; 21), *empowerment*, the *potential* information flow that an agent can send through the environment and perceive through its own sensors, has been introduced as a universal utility for agents, i.e. as a quantity which could be plausibly optimized by intelligent agents[1]. Although this involves the potential to generate entropy (via the actuators), it is different from Sec. 2 in that the entropy needs not be actually created, only states with the potential of generating entropy are favoured.

This concept is closely related to the concept of mobility in games, in that it provides the player with the most options in a situation, and thus with a higher (unbiased) chance of responding to an unexpected situation. Whereas, the enforced entropy production from Sec. 2 does not maximize options, but actually exploits them.

The difference between potential entropy production (empowerment) and the actual entropy generation from Sec. 2 is also the context in which we hypothesize an agent will apply each of these: the first will be applied if there is no clear utility/fitness profile in the state of the agent. That is, the environment is either unbiased on the short term, or there is not enough feedback (reinforcement) information about which states are useful and which are not. If, thus, not much is known as to which states should be desirable to an agent, empowerment is a criterium that can be used to drive an agent's behaviour.

Once, however, clear criteria for utility, fitness or success exist, the actual generation of entropy can become central, such as in above game-theoretical example. At present, there is no clear picture how to model the transition between "unbiased" models aiming for maximum potential entropy production and the "utility biased" models aiming for (actual) entropy production.

[1] There are indications that for minimally cognitive agents (4) empowerment or related measures could provide an intrinsic drive for "interesting" or "relevant" states in the world(11; 19; 30; 29; 26).

4 Transforming Entropies

Between the micro-level maximum entropy and maximum entropy production and the agent (i.e. meso-level) enforced entropy production and potential entropy production, there is still a conceptual gap. In the current section, we discuss several issues that have to be resolved in order to reduce or close this gap.

First, we should discuss what it implies if one considers an "entropy producing" agent. As noted in (1), the unitarity of the equations of physical motion (both in classical and in quantum physics) implies that, in a faithful microscopic model of the system, the entropy remains unchanged over time. In this picture, the second law of thermodynamics of entropy growing with time emerges through the fact that an observer aware of the inital state of the system will quickly lose track of it (1). The growth of the entropy reflects this loss of information about the system throughout observation. In other words, the coarse-graining (in Adami's view, in the form of an observer) is necessary to actually obtain an increase of entropy with time. Without that, entropy remains fixed throughout time.

A problem that is associated with requiring the presence of an observer to be able to define coarse-grained entropy is that there could be the danger of arbitrariness (16). Shalizi shows, however, a system can exhibit special "observers" (28), via their *predictive efficency*. This is the ratio between how much information the system's past has about the future (excess entropy) in ratio to its statistical complexity (i.e. the complexity of an ϵ-machine performing this prediction) (10). It turns out that for a simple thermodynamic system, this efficiency is close to 0 for a microscopic observer (in fact, the original system description), see (28). Even endowed with full information concerning the initial state of a system, a microscopic observer has to employ enormous effort to predict the progress of the system as time progresses. However, a judiciously selected macroscopic observer (e.g. in the simple thermodynamic system at hand, Shalizi uses energy, particle number and volume) can achieve a predictive efficiency of almost 1. It is this phenomenon that makes the selection of coarse-grained observers a natural choice in thermodynamics.

For our mesoscopic agents, this conservation of entropy, however, poses now an interesting question: if randomness needs to be produced on the one side, order has to be produced on the other. This means that the randomness produced by the agent on the mesoscopic scale (in which the agent acts) has to be extracted from the microscopic scale — see also the related discussion from (6). This suggests some interesting implications.

5 Growth of Complexity

In (5), Bennett hypothesizes that there is a possible low-level drive that pushes physical systems towards increasingly higher complexity. In particular, he argues, there may be a natural limitation into how much complexity can be accumulated per time. While entropy (for an incomplete observer) can grow very quickly — e.g. (28) —, the accumulation of complexity is limited and grows only slowly. The definition of complexity as "frozen alleles" as in (2) supports Bennett's hypothesis. Also here, structural information about sufficiently useful dynamics is slowly accumulated in the frozen bits of an artificial genome over an evolutionary simulation.

This complexity growth is a generation of order. In view of the conservation of entropy, it means that randomness has to be slowly "squeezed out" of the system to enable the accumulation of this organization. The open question remains whether this relatively limited randomness is the one that can be used to provide an advantageous strategy of entropy generation for a game-theoretical scenario, or whether one rather should employ thermal baths as the primary source of randomness for this purpose.

It should be noted that the growth of complexity is itself a question of significant importance and that the forces which drive this growth are still under research. In this respect, the question whether enforced entropy generation is linked with the long-term accumulation of complexity is of high interest. Since, however, significant layers of mechanisms and dynamics lie between the micro- and the mesoscale of the systems at hand, it is not clear how (and whether) the microscopic balance of entropy propagates to the meso-level in a visible way.

Consider now the potential entropy production (i.e. empowerment). This essentially marks the *option* of controlling external disturbances, as a generalization of the *law of requisite variety* (3; 31; 32). Here, once the potential entropy production is maximized, the actual entropy generation will be typically driven by external disturbances (which hence provide the randomness), not by an internal or "clandestine" reservoir of randomness. The hypothesis put forward in (20; 21) is that driving an agent along an empowerment-maximizing route provides, similar to other related approaches (12; 11; 29; 22; 26) a pathway towards increasingly complex organization of agent dynamics. For instance, empowerment-based behaviour optimization has been used to create organized collective behaviour in agent colonies (8).

As opposed to the entropy production, in empowerment it is the structure of the agent/environment interaction that drives the emergence of behaviour patterns. If complexity growth is attributed to this, then the source of complexity accumulation is different from the earlier maximum entropy production and is derived from the structure in the agent's environment. Complexity growth would then be considered as a response to capture the regularities (i.e. the predictability) of the environment, while aiming to acquire the environmental state as basis to control the unpredictable part of the environment. This can be considered as outsourcing the unpredictability of controlling the environment back to the dynamics of the environment itself, which has been suggested in recent years by proponents of the concept of "intelligent behaviour without intelligence" and "morphological intelligence" (e.g. (7; 24)).

6 Conclusions

We have argued that there are interesting issues connecting the classical physical principles of maximum entropy, the more novel maximum entropy production and the principles of enforced entropy production and maximum potential entropy production, hypothesized to be relevant for the emergence of complexity and "intelligent" behaviour in living agents. Although the latter are still in a hypothetical stage, the relatedness of the concepts and the similarity of the language used suggests that there are significant possibilities to bridge the gap between micro-level physical and meso-level complex systems. It is clear that our understanding of the emergence of complexity and — possibly — life would significantly profit if one could formulate a connection between these related, yet distinct principles in the future.

Bibliography

[1] Adami, C., *Introduction to Artificial Life*, Springer (1998).

[2] Adami, C., C. Ofria, and T. C. Collier, "Evolution of biological complexity", *Proc. Natl. Acad. Sci. USA* **97** (2000), 4463–4468.

[3] Ashby, William Ross, *Design for a Brain*, Wiley & Sons New York (1952).

[4] Beer, Randall D., "Toward the evolution of dynamical neural networks for minimally cognitive behavior", *From Animals to Animats 4* (Cambridge, MA,) (P. Maes, M. Mataric, J. Meyer, J. Pollack, and S. Wilson eds.), MIT press (1996), 421–429.

[5] Bennett, Charles H., "How to define complexity in physics, and why", *Complexity, Entropy and the Physics of Information* (Reading, Mass.,) (W. H. Zurek ed.), Santa Fe Studies in the Sciences of Complexity, Addison-Wesley (1990), 137–148.

[6] Bennett, Charles H., and Rolf Landauer, "The fundamental limits of computation", *Scientific American* (July 1985), 48–56.

[7] Brooks, Rodney A., "Intelligence without representation", *Artificial Intelligence* **47**, 1–3 (Jan. 1991), 139–159.

[8] Capdepuy, Philippe, Daniel Polani, and Chrystopher Nehaniv, "Maximization of potential information flow as a universal utility for collective behaviour", *2007 IEEE Symposium on Artificial Life*, IEEE (2007).

[9] Cover, Thomas M., and Joy A. Thomas, *Elements of Information Theory*, Wiley New York (1991).

[10] Crutchfield, J. P., "The calculi of emergence: Computation, dynamics, and induction", *Physica D* (1994), 11–54.

[11] Der, Ralf, "Selforganized robot behavior from the principle of homeokinesis", *Proc. Workhop SOAVE '2000 (Selbstorganisation von adaptivem Verhalten)* (Ilmenau,) (H.-M. Gross, K. Debes, and H.-J. Böhme eds.), vol. 643 of *Fortschritt-Berichte VDI, Reihe 10*, VDI Verlag (2000), 39–46.

[12] Der, R., U. Steinmetz, and F. Pasemann, "Homeokinesis – a new principle to back up evolution with learning", *Computational Intelligence for Modelling, Control, and Automation*, (M. Mohammadian ed.) vol. 55 of *Concurrent Systems Engineering Series*. IOS Press (1999), pp. 43–47.

[13] Dewar, Roderick, "Information theory explanation of the fluctuation theorem, maximum entropy production and self-organized criticality in non-equilibrium stationary states", *J. Phys. A: Math. Gen.* **36**, 3 (Jan 2003), 631–641.

[14] Dewar, Roderick, "Maximum entropy production and the fluctuation theorem", *J. Phys. A: Math. Gen.* **38** (5 2005), 371–381.

[15] Filyukov, A. A., and V. Y. Karpov, "Method of the most probable path of evolution in the theory of stationary irreversible processes", *Inzhenerno-Fizicheskii Zhurnal* **13**, 6 (1967), 798–804.

[16] Haken, H., *Advanced synergetics*, Springer-Verlag Berlin (1983).

[17] Jaynes, E. T., "Information theory and statistical mechanics", *Phys. Rev.* **106**, 4 (1957), 620–630.

[18] Jaynes, E. T., "Information theory and statistical mechanics II", *Phys. Rev.* **108**, 2 (171-190 1957).

[19] Kaplan, F., and P.-Y. Oudeyer, "Maximizing learning progress: an internal reward system for development", *Embodied Artificial Intelligence*, (F. Iida, R. Pfeifer, L. Steels, and Y. Kuniyoshi eds.) vol. 3139 of *LNAI*. Springer (2004), pp. 259–270.

[20] Klyubin, Alexander S., Daniel Polani, and Chrystopher L. Nehaniv, "All else being equal be empowered", *Advances in Artificial Life, European Conference on Artificial Life (ECAL 2005)*, vol. 3630 of *LNAI*, Springer (2005), 744–753.

[21] Klyubin, Alexander S., Daniel Polani, and Chrystopher L. Nehaniv, "Empowerment: A universal agent-centric measure of control", *Proc. IEEE Congress on Evolutionary Computation, 2-5 September 2005, Edinburgh, Scotland (CEC 2005)*, IEEE, (2005), 128–135.

[22] Lungarella, Max, and Olaf Sporns, "Mapping information flow in sensorimotor networks", *PLoS Computational Biology* **2**, 10 (2006).

[23] MARTYUSHEV, L. M., and V. D. SELEZNEV, "Maximum entropy production principle in physics, chemistry and biology", *Physics Reports* **426** (2006), 1–45.

[24] PFEIFER, Rolf, and Josh BONGARD, *How the Body Shapes the Way We think: A New View of Intelligence*, Bradford Books (2007).

[25] PRIGOGINE, Ilya, and G. NICOLIS, *Self-Organization in Non-Equilibrium Systems: From Dissipative Structures to Order Through Fluctuations*, J. Wiley & Sons New York (1977).

[26] PROKOPENKO, Mikhail, Vadim GERASIMOV, and Ivan TANEV, "Measuring spatiotemporal coordination in a modular robotic system", In ROCHA et al. (27).

[27] ROCHA, Luis M., Mark BEDAU, Dario FLOREANO, Robert GOLDSTONE, Alessandro VESPIGNANI, and Larry YAEGER eds., *Proc. Artificial Life X*.

[28] SHALIZI, Cosma Rohilla, *Causal Architecture, Complexity and Self-Organization in Time Series and Cellular Automata*, PhD thesis University of Wisconsin-Madison (2001).

[29] SPORNS, Olaf, and Max LUNGARELLA, "Evolving coordinated behavior by maximizing information structure", In ROCHA et al. (27).

[30] STEELS, L., "The autotelic principle", *Embodied Artificial Intelligence*, (I. FUMIYA, R. PFEIFER, L. STEELS, AND K. KUNYOSHI eds.) vol. 3139 of *Lecture Notes in AI*. Springer Verlag Berlin (2004), pp. 231–242.

[31] TOUCHETTE, Hugo, and Seth LLOYD, "Information-theoretic limits of control", *Phys. Rev. Lett.* **84** (2000), 1156.

[32] TOUCHETTE, Hugo, and Seth LLOYD, "Information-theoretic approach to the study of control systems", *Physica A* **331** (2004), 140–172.

[33] VIRGO, N., and I. HARVEY, "Entropy production in ecosystems", *Proceedings of European Conference on Artificial Life (ECAL 2007, Lisbon)*, Springer (2007).