Workshop on
Information Theoretic Incentives for Artificial Life

Christoph Salge Keyan Ghazi-Zahedi Georg Martius Daniel Polani

July 30, 2014

SUNY Global Center, New York, 30th of July

Program

13:00 – 13:05	Christoph Salge	Opening and introduction
13:05 – 13:55	Chis Adami	Keynote "Information-theoretic musings concerning the origin and evolution of life"
13:55 – 14:20	Simon D. Levy	"Wittgenstein's Robot: philosophy, information, and artificial life"
14:20 – 14:35		Coffee break
14:35 – 15:00	Claudius Gros	"The Fisher information as a guiding principle for self-organizing processes"
15:00 – 15:20	Georg Martius	"Predictive information as a drive for self-organizing behavior."
15:20 – 15:35	Tobias Morville	"The homeostatic logic of reinforcement learning"
15:35 – 16:00	Christoph Salge	"Perspective of information theoretic incentives"

Special Issue

The open access journal *Entropy* sponsors this workshop by an open call, special issue on the topic of "*Information Theoretic Incentives for Cognitive Systems*"

More details will be announced to emails received via itialife@gmail.com and over the alife and connectionists mailing lists.

Contact

web: `http://www.mis.mpg.de/ay/index.html?c=workshops/alife14ws.html`
email: `itialife@gmail.com` or `c.salge@herts.ac.uk` or `georg.martius@mis.mpg.de`

Information-theoretic musings concerning the origin and evolution of life

Chris Adami
Michigan State University

June 17, 2014

In this talk, I will discuss two applications of information theory to topics that are important in Artificial Life: an information-theoretic view on the origin-of-life problem, as well as an application of information theory to understanding how brains make decisions.

Research investigating the possible origins of life usually focuses on exploring possible life-bearing chemistries in the pre-biotic Earth, or else on synthetic approaches. Very little work has been done exploring fundamental issues concerning the spontaneous emergence of life, using only concepts (such as information and evolution) that are divorced from any particular chemistry. I advocate studying the probability of spontaneous molecular self-replication as a function of the information contained in the replicator, and the environmental conditions that might enable this emergence, and find that there may be environments that provide "information for free", so that the information gap between the abiotic and biotic regime is minimized.

When evolved brains make appropriate decisions in the environment they are in, they must make use of internal models of the world, called "representations". It is not clear how knowledge about the world is represented in btains, so I propose and information-theoretic measure that not only quantifies how much knowledge about the workd is represented in the brain, but even what concepts are being represented, and how the brain uses these concepts to make decisions. While the analysis pertains to brains evolved in the computer, the methods I describe could be used to analyse the firing patterns of real brains that control behaving anmals.

Wittgenstein's Robot: Philosophy, Information, and Artificial Life

Simon D. Levy[1] and Charles Lowney[2]

[1] Computer Science Department, Washington and Lee University
Lexington Virginia 24450, USA
levys@wlu.edu
[2] Philosophy Department, Washington and Lee University

This paper addresses one of the major themes of the workshop: *the question on how and where the necessary computation to realise agent behaviour is performed.* The traditional view of such computation describes it in terms of symbols and rules of the sort used in writing predicate calculi or formal grammars. The precise details of the symbols and rules do not much matter; what is important here is the hypothesis that any physical system that instantiates the symbols and rules in an explicit, consistent way is a reasonable candidate for being a model of mind [1]. Because of the the brittleness of such rule-based systems and the difficulty of scaling them up to real-world problems, agent-based modeling has shifted its focus in recent years to the use of probabilistic, data-driven models [2]. Relying heavily on mathematically precise notions of information, the learning algorithms for such models [3] have brought information theory once more to the fore.

Advocates of the symbols-and-rules approach have always been quick to point out the limitations of probabilistic / data-driven models (notably, connectionist models): although such models show an impressive ability to learn both rule-like and exception-based patterns, these critics argue that there is little evidence that they are capable of modeling the systematic, compositional nature of language and thought [4]. Without the ability to compose and decompose propositions and other structures in systematic ways there seems to be little reason to expect purely probabilistic models to work for more abstract reasoning such as planning or language.

Hybrid models, such as probabilistic context-free grammars [5] combine features of both approaches and have achieved significant practical results. But any sort of hybrid leads us to wonder about the possibility of a single, unifying framework for representing and acquiring structured knowledge that might generate behavior. Although it would be naive to think that such a framework by itself could end up being the source of all useful agent behavior, we see a challenge in providing a principled characterization of the source of behavior that can be considered "inside" the agent.

In response to this challenge, we have spent the past decade or so developing connectionist models that support the acquisition of systematic, compositional behavior from a small number of positive examples and provide plausible, scalable models for language, behavior, and thought. The general term we use for these models – Vector Symbolic Architecture, or VSA [6] – describes a class of connectionist networks that use high-dimensional vectors of low-precision numbers to encode systematic, compositional information as distributed representations. VSAs can represent complex entities such as multiple role/filler relations or attribute/values pairs in such a way that every entity – no matter how simple or complex – corresponds to a pattern of activation distributed over all the elements of the vector. The biggest advantage of VSA representations over other connectionist approaches is that a single association (or set of associations) can be quickly recovered from a set (or larger set) of associations using the same simple operator that creates the associations, in a time that is independent of the number of associations.

In developing VSA and explaining it to colleagues, we have been struck by the extent to which it is also the kind of model of mental activity that we arrive at if we take seriously Ludwig Wittgenstein's critiques of the philosophy of language and mind. We note, for example, that the uniform nature of representation in VSA eliminates the sort of problems that early Wittgenstein criticized in Russell's theory of types and related formalisms. The idea that all symbols are of the same type coordinates well with VSA representation, which, in turn, coordinates well with Wittgenstein's later view: meaningful activity (linguistic and non-linguistic) emerges from practices and not from the artificial characterizations we impose by designating symbols and manipulating them with formal rules [7]. This insight leads us to question and perhaps amend some of the underlying assumptions of the workshop regarding the existence of intrinsic motivations producing behavior.

First, the concept, symbol, or representation would not be considered an intrinsic or primal cause that motivates behavior. VSA representations can emerge from experiences the system encounters, so the "meanings" represented in Wittgenstein's robot would be a product of its interactions and not something independent, like a function that is programmed in ahead of time. Second, the robot's concept of "boredom" for instance, would not be the same as ours even though we might be able to generate what we would want to call "bored behavior" in the robot. As Wittgenstein famously noted, because we do not have the same form of life and practices as a lion, "if a lion could speak, we could not understand him" [7, p.223].

Recent work by our research group has shown how using VSA can support inferencing of analogical structure through implicit maximum-likelihood estimation [8]. We have also constructed a simple VSA-based model for behavior-based control of a simulated robot [9]. We are currently investigating the possibility of combining VSA with Sparse Distributed Memory [10], to model the acquisition of sequential structure by an agent navigating an environment. We plan to present all three models at the workshop, demonstrating the progress we have made towards constructing "Wittgenstein's robot."

References

[1] Newell, A.: Physical symbol systems. Cognitive Science **4**(2) (1980) 135–183

[2] Thrun, S., Burgard, W., Fox, D.: Probabilistic Robotics (Intelligent Robotics and Autonomous Agents). The MIT Press (2005)

[3] Murphy, K.P.: Machine learning: a probabilistic perspective. The MIT Press, Cambridge, MA (2012)

[4] Fodor, J., Pylyshyn, Z.W.: Connectionism and cognitive architecture: A critical analysis. Cognition **28** (1988) 371

[5] Jelinek, F., Lafferty, J., Mercer, R.: Basic methods of probabilistic context free grammars. In Laface, P., Mori, R., eds.: Speech Recognition and Understanding. Volume 75 of NATO ASI Series. Springer Verlag, Berlin (1992) 345–360

[6] Gayler, R.: Vector symbolic architectures answer jackendoff's challenges for cognitive neuroscience. In Slezak, P., ed.: ICCS/ASCS International Conference on Cognitive Science. CogPrints, Sydney, Australia, University of New South Wales (2003) 133–138

[7] Wittgenstein, L.: Philosophical Investigations. Basil Blackwell, Oxford (1958) trans. G.E.M. Anscombe.

[8] Gayler, R., Levy, S.: A distributed basis for analogical mapping. In: Proceedings of the Second International Analogy Conference, NBU Press (2009)

[9] Levy, S.D., Bajracharya, S., Gayler, R.: Learning behavior hierarchies via high-dimensional sensor projection. In Pickett, M., ed.: Learning Rich Representations from Low-Level Sensors: Papers from the 2013 AAAI Workshop, AAAI Press (2013)

[10] Kanerva, P.: Sparse Distributed Memory. Cambridge, Massachusetts, MIT Press (1988)

The Fisher information as a guiding principle for self-organizing processes

Claudius Gros and Rodrigo Echeveste
Institute for Theoretical Physics, Goethe University Frankfurt, Germany
gros07[at]itp.uni-frankfurt.de

Information theoretical principles can be employed to derive many known fundamental laws of nature. Maximum entropy estimates may by used (Jaynes, 1957), on a macroscopic level, to derive statistical mechanics, and the Schrödinger equation of quantum mechanics is equivalent (Reginatto, 1998), on a microscopic level to the minimization of the Fisher information of the probability density, together with the continuity equation.

In the same way, information theoretical guiding principles can be considered as incentives for self-organizing processes. In this sense, the role of predictive information has been studied in the context of autonomous robots (Ay et al., 2008) and within the perception-action cycle (Tishby and Polani, 2011).

The two most important classes of information-theory based generating principles are Shannon's information entropy, and variants thereof, and the Fisher information

$$F(\theta) \;=\; \int p(y,\theta)\left(\frac{\partial}{\partial\theta}\,\ln(p(y,\theta))\right)^2 dy \;=\; \int \frac{1}{p(y,\theta)}\left(\frac{\partial p(y,\theta)}{\partial\theta}\right)^2 dy$$

of a probability distribution function $p(y,\theta)$, dependent on some external parameter θ. The Fisher information plays an important role when estimating a given parameter θ (Brunel and Nadal, 1998; Prokopenko et al., 2011). Here, alternatively, we propose to consider $F(\theta)$ as an incentive function for the generation of equations of motion governing the time evolution of a slow variable $\theta = \theta(t)$ (Echeveste and Gros, 2014).

Learning is an important task for both biological and artificial cognitive systems and involves the adaption of slow variables, or parameters, a process also termed meta-dynamics. It also involves synaptic plasticity in the context of neural networks and generating principles for synaptic updating rules are hence at the base of learning processes in general. Here we propose to consider the Fisher information of the post-synaptic firing-rate distribution $p(y)$ with respect to the synaptic flux operator, compare Fig. 1, as an information-theoretical based objective function for Hebbian-type learning.

The purpose of adapting synaptic weights is to encode the information present in the statistics of the afferent inputs. The statistics of the output neural activity will hence become stationary whenever this task is completed. At this point the sensitivity of the activity of the post-synaptic neuron, with respect to changes in the synaptic weights, will become minimal. Minimizing the Fisher information with respect to the synaptic flux is hence a natural way to generate synaptic plasticity rules.

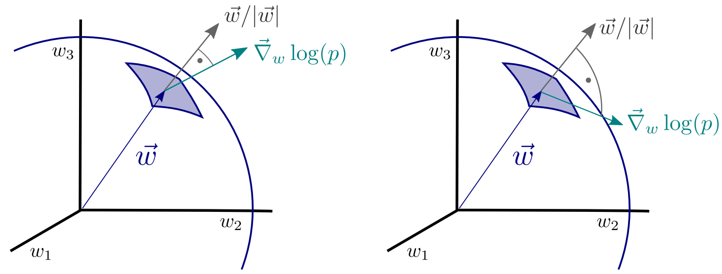

Figure 1: The synaptic flux operator $\partial/\partial\theta \equiv \vec{w} \cdot \nabla_{\vec{w}} \log(p)$ is minimal when the gradient $\nabla_{\vec{w}} \log(p)$ of the firing-rate distribution $p = p(y)$ is orthogonal to the synaptic weight field $\vec{w}$.

We find that the plasticity rules obtained by minimizing the Fisher information for the synaptic flux have a set of attractive features. They incorporate standard Hebbian updating, are naturally self-limiting, perform a principal component analysis and, whenever possible, a binary classification of input features. We also investigate the time course of the learning process and observe the self-organized emergence of a fading memory.

References

Nihat Ay, Nils Bertschinger, Ralf Der, Frank Güttler, and Eckehard Olbrich. Predictive information and explorative behavior of autonomous robots. *The European Physical Journal B*, 63(3):329–339, 2008.

Nicolas Brunel and Jean-Pierre Nadal. Mutual information, fisher information, and population coding. *Neural Computation*, 10(7):1731–1757, 1998.

Rodrigo Echeveste and Claudius Gros. Generating functionals for computational intelligence: The fisher information as an objective function for self-limiting hebbian learning rules. *Frontiers in Robotics and AI*, (in press), 2014.

Edwin T Jaynes. Information theory and statistical mechanics. *Physical review*, 106(4):620, 1957.

Mikhail Prokopenko, Joseph T Lizier, Oliver Obst, and X Rosalind Wang. Relating fisher information to order parameters. *Physical Review E*, 84(4):041116, 2011.

Marcel Reginatto. Derivation of the equations of nonrelativistic quantum mechanics using the principle of minimum fisher information. *Physical Review A*, 58:1775–1778, 1998.

Naftali Tishby and Daniel Polani. Information theory of decisions and actions. In *Perception-Action Cycle*, pages 601–636. Springer, 2011.

Predictive information as a drive for self-organizing behavior

Georg Martius
Max Planck Institute for Mathematics, Leipzig, Germany
martius@mis.mpg.de

May 12, 2014

Autonomy is a puzzling phenomenon in nature and a major challenge in the world of artifacts. A key feature of autonomy in both natural and artificial systems is seen in the ability for independent exploration. In animals and humans, the ability to modify its own pattern of activity is not only an indispensable trait for adaptation and survival in new situations, it also provides a learning system with novel information for improving its cognitive capabilities, and it is essential for development.

We propose to implement the exploration as a deterministic law derived from maximizing an information quantity. We have studied the use of predictive information of the sensor process (of a robot) as a driving force for behavior generation[1]. In order to obtain an update rule (exploration dynamics) for the controller parameters and be adequate in robotics application the non-stationary nature of the underlying time-series have to be taken into account, which can be done by a time-local predictive information. Importantly the exploration dynamics is derived analytically and by this we link information theory and dynamical systems. Without a random component the change in the parameters is deterministically given as a function of the states in a certain time window. For an embodied system this means in particular that constraints, responses and current knowledge of the dynamical interaction with the environment can directly be used to advance further exploration.

Another aspect of information quantities is the possibility to measure behavior, for instance by quantifying the complexity and the attractor dimension of the sensor process. These measures can be obtained in a length-scale dependent fashion allowing us to make statements about different levels of detail and to compare different behaviors objectively.

References

[1] Georg Martius, Ralf Der, and Nihat Ay. Information driven self-organization of complex robotic behaviors. *PLoS ONE*, 8(5):e63400, 2013.

The Homeostatic Logic of Reinforcement Learning

Avoiding death is the ultimate challenge. For an organism to stay alive, it must regulate the internal states of its body within narrow bounds, and it must do so in the face of stochastic and often volatile environments [1,2]. For instance, only small deviations in the systemic temperatures, hydration, and oxygen-levels, can be the difference between life and death [3]. As such, biological agents have evolved sophisticated, flexible, and highly robust, mechanisms for physiological regulation. However, for mobile organisms, autonomous physiological mechanisms are necessary, but not sufficient to maintain homeostasis. Overt behavior is critical for maintaining homeostasis – escaping predators, foraging, finding water, and keeping warm – are all examples of how behavior, and therefore the cognitive systems supporting them, are pivotal in mediating homeostatic regulation. Although obvious, and despite a wealth of knowledge in physiological and cognitive sciences, the question of how living and artificial cognitive systems, optimize homeostasis via behavior has been largely neglected.

Our approach is centred on the core assumption that behaviour is driven ultimately by the need to optimise homeostasis. Specifically our work builds on the Homeostatic Reinforcement Learning (HRL) framework [4] (Fig. 1), which provides a generalised framework for how homeostatic states should be valued, and how actions can be learnt to predictively optimize homeostatic states in living and artificial agents. It was previously demonstrated that, by integrating a homeostatic definition of reward (as drive-reduction) into reinforcement learning algorithms, then reward maximization and homeostatic regulation are optimized simultaneously. Fig. 1 shows a schematic of the original HRL model in an exemplary 2-dimensional homeostatic space for glucose and temperature. Upon performing an action, the agent receives an outcome *Kt* from the environment. The rewarding value of this outcome depends on its ability to make the internal state, *Ht*, closer to the homeostatic setpoint, *H**, and thus reduce the drive level (the vertical axis). This experienced reward, denoted by *r(Ht,Kt)*, is then learnt by Reinforcement Learning algorithms [5], and via value computation, is the basis upon which actions are selected. The form of this model mandates several behaviors not assumed in the model including - risk aversion, anhedonic effects of irrelevant drives, excitatory effects of drive, and temporal discounting.

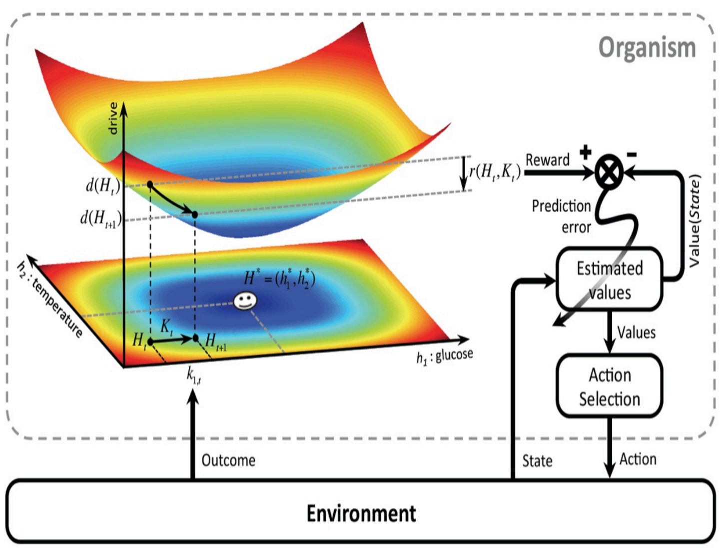

Here we make a number of advances on the original HRL model, and provide an extended set of predictions in biological and artificial domains. We motivate the general form of the drive function from basic evolutionary principles, such that the general functional form of survival probabilities over homeostatic states, has the effect of mandating drive functions to be both smooth and unimodal. This generalizes the model beyond its specific reliance on arbitrary parameters, but without losing the any of the central behavioral predictions. Further, the model is extended to include the full space of surplus and deficit states, and is discussed in the context of additional behavioural phenomena predicted including loss aversion, valence grounding, and risk aversion for losses. From this foundation, we advance the argument that homeostatically-grounded computational theories of cognition are central to understanding, modeling, and building complex living systems.

Fig. 1 | Schematic of Homeostatic Reinforcement Learning Framework (figure from presentation by Keramati at DRCMR 2013, use permitted by author)

Litterature:

[1] Cannon, W. B. (1929). Organization for physiological homeostasis. *Physiological Reviews*, *IX*(3), 399–431.
[2] Sterling, P. (2012). Allostasis: a model of predictive regulation. Physiology & Behavior, 106(1), 5–15. doi:10.1016/j.physbeh.2011.06.004
[3] Schmidt-Nielsen, K. (1997). Animal physiology: adaptation and environment.
[4] Keramati, M., & Gutkin, B. S. (2011). A Reinforcement Learning Theory for Homeostatic Regulation., 82–90.
[5] Sutton, R. S., & Barto, A. G. (1998). Reinforcement Learning: An Introduction. IEEE Transactions on Neural Networks, 9(5), 1054–1054. doi:10.1109/TNN.1998.712192