Refinement: a rigorous description of autonomous adaptive agents

J. Augustus Bacigalupi
Institute for Augmenting Minds, San Francisco, California, USA

Received 2 October 2012
Revised 12 March 2013
Accepted 7 October 2013

Abstract

Purpose – The purpose of this paper is to develop a theoretical framework to empirically test for cognitive behaviour in autonomous adaptive agents.

Design/methodology/approach – This project proposes a theoretical framework, or design parameters, inspired by empirically observed phenomena, cognitive behaviours, and thermodynamics. Success of the framework is measured by its capacity to implement, not just a model of select attributes of cognition, but to implement the foundational physical nature of cognition of which all observed behaviours are based.

Findings – A rigorous mathematical framework, employing only information theory and conventional physics, is hypothesized to empirically measure for cognitive behaviours.

Research limitations/implications – Empirical studies will be conducted on synthetic agents using the theoretical framework described herein to demonstrate whether or not cognitive behaviours have been achieved.

Originality/value – This paper proposes an alternative form of information processing inspired by evolved organisms, distinct from Turing equivalent machines, able to augment existing human and digital systems.

Keywords Cognition, Artificial intelligence, Information theory, AGI, Alternative computing, Reservoir computing, Cyberdization, Singularity

Paper type Research paper

Introduction

Generally, an autonomous adaptive agent is any system that is restructured via internal dynamic constraints, without external design intent or programming, such that the chances of the continued persistence of that system's lineage are increased within its open physical environment. This is demonstrated by both evolution of individuals over many generations and the capacity for single individuals to learn within a generation. The capacity for evolution of agents over generations is a separate, yet related, project being explored with Deacon (2011). This paper, more specifically, presents a hypothesis and empirical test for the latter, i.e. learning. It was developed within a larger research proposal for synthetic cognition, which specifies a buildable draft model of a learning system inspired by animal cognition and thermodynamics. A very concise abstract of that proposal can be found here (Bacigalupi, 2012a), and the complete draft research proposal can be found here (Bacigalupi, 2012b). This paper focuses on the theoretical aspects of this project; and, therefore, will not re-explain the model in Bacigalupi (2012b), but use closely related analogies to focus on the theoretical concepts.

What separates this project from others – such as: Humberto Maturana and Francisco Varela's autopoiesis; second order cybernetics; and, dynamical systems

Kybernetes
Vol. 42 No. 9/10, 2013
pp. 1313-1324

0368-492X
DOI 10.1108/K-10-2012-0065

theory – is the indeterminate nature of thermodynamics. As Stuart Kauffman *et al.* argue in Longo *et al.* (2012), there are no entailing laws or rules capable of fully describing living autonomous systems in their open system surroundings. As Deacon also argues in Deacon (2011), the "function" of an agent is an evolving relation to its open surroundings, and cannot therefore be preconceived a priori. Yet, most existing projects implicitly assume that such closed and determinate functional analysis is both possible and probable. This project does not make that assumption. This proposal intends to ground cybernetics in the physics, so as not to predetermine autonomous agents in a mechanical sense, but to explain how adaptive complexity emerges. And, it is this distinction that also separates this project from countless models implemented upon any Turing equivalent substrate. It is exactly the genius of the Turing machine that processes intended to be closed, or "entailed", are insulated from thermodynamic corruption. This project is unique in treating open system thermodynamics not as "noise" to be filtered, but as necessary structured noise to be leveraged.

In practice, the larger research proposal (Bacigalupi, 2012b) will build novel hardware that will take as inputs patterned stimulation, e.g. visual and auditory patterns. The expected output, or alternative hypothesis, will be an autonomous increase in the internal order of this novel hardware's structure; and that, furthermore, this self-biased structure, i.e. complexity, will increase the work output of the system in the subsequent presence of formerly "learned" patterns. This paper will use a musical analogy to explain, in theory, the behaviour of such a system and what will be measured to justify or deny the Alt. Hypothesis. Again, the physical specifications of the larger research proposal are explained in brief here (Bacigalupi, 2012a), and in detail here (Bacigalupi, 2012b).

The theory describes the measurable capacity of an autonomous agent's work capacity and how that relates to changing internal structural complexity when exposed to external stimulation. This self-bias is called Relevance, which indicates to what extent a system has re-organized its internal structure in the presence of external complex patterns so as to become selectively responsive to those patterns when they re-occur. The change in a system's Relevance over time is called Refinement, a measure of how much a system has learned via experience. The Alt. Hypothesis is that Refinement will have increased on average over time.

Central to developing a useful framework for an adaptive autonomous agent in its open system environment is the claim that all events in nature are both distinct relative to their surroundings, yet simultaneously inter-related with other patterns and events. In other words, no event is completely independent of all other events. If it were, it would be transparent to physical reality and therefore non-causal. This raises a paradox in that all events are both distinct and inter-related at the same time. Like interacting particles and their respective extensive fields, each moulds and is moulded by the other, simultaneously. This observation challenges the traditionally discrete and sequential understanding of causality, and shall be the base premise for this project:

> Events in physical reality are *both* distinct *and* inter-related, *simultaneously*.

This premise is illustrated by the ripple analogy in Figure 1.

Living systems must also implement this paradoxical behaviour in that they must simultaneously assimilate relevant patterns from their environment to some degree of accuracy, while simultaneously assimilating relationships between these patterns. These two tasks require quite different strategies: the former isolates and objectifies

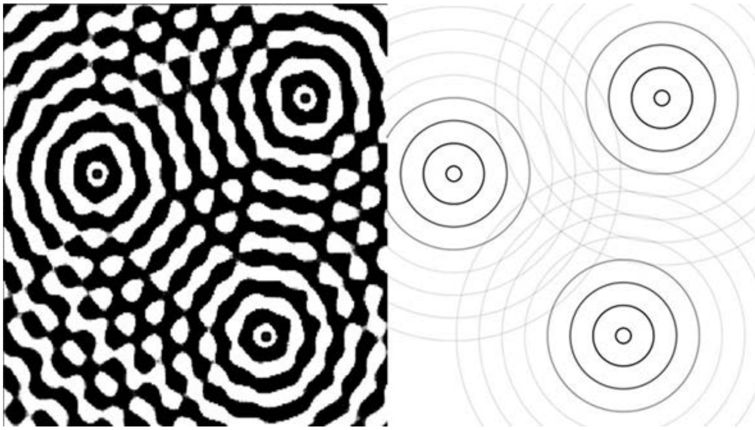

Notes: Left: contrasted image of actual ripple interference pattern illustrating superposition and simultaneity; right: abstraction of ripples used for this paper

Figure 1.
Ripple analogy

the pattern, distinguishing that which separates it from its surroundings; the latter does the opposite by discerning what each pattern has in common with other distinct patterns. And, as with most things in nature, the optimal solution to this apparent contradiction is a balance between both regimes.

When a system can autonomously adapt to its environment so as to increase each of these complementary skills simultaneously, it is called Refinement.

Refinement informally defined

This general adaptive behaviour, called Refinement, entails the capacity to learn in a lifetime both many diverse patterns while also learning useful relations among these patterns. For example, redness, smoothness, and roundness are all useful attributes of food items to perceive; however, these distinct attributes must be inter-related simultaneously in order to conceive of an apple. Knowing infinitely many things is far less useful and adaptive for the organism than knowing only some things plus their inter-relations, not sequentially but simultaneously. Refinement is an increased capacity to change one's internal structure so as to assimilate both distinct patterns and their inter-relations in order to do more adaptive work over time, thereby remaining relevant:

> Refinement: *noun.* A process whereby a system increases its average Relevance (viz., engagement, sensitivity, or "awareness") to its surroundings by increasing its own internal capacity to resolve internal disorder, which is the result of increased Relevance, into a higher-order internal structure able to increase capacity to do useful work.

One difficulty in understanding living systems is that they are never the same system twice. Fundamental to adaptation is the capacity to change. So, when a system changes its internal structure to become sensitive to many diverse patterns, these sensitivities compete inside the agent creating disorder, or "noise". But this noise is not purely uncorrelated random noise; it is structured noise. Living systems have evolved to harvest the sparse order within the so-called "noise". And, as will be shown shortly, it is this evolved response to many distinct, yet superimposed, patterns within a system that is the basis for Refinement.

State of complexity

Refinement reflects an autonomous agent's ability to do adaptive work in recognizing and navigating its often novel environment. A precondition for adaptive work is that a system's internal structure must be commensurately constrained to do that work. This constraint must be commensurately complex. Inspired by extensions of information theory (Roy, 2002), this paper proposes a rigorous statistical measure of complexity, called state of complexity (SoC).

Although the conceptual core of information theory is entropy, practical communication systems do not minimize or maximize order. Similarly, living systems do not minimize or maximize internal order. Both pure randomness and order are "death states", because neither extreme is conducive to vital complexity as shown in Figure 2. In order to assimilate both distinct and inter-related patterns from its environment, the agent must embody order in a similarly hierarchical way, i.e. simpler patterns are embodied, which create the foundation for the assimilation of more complex patterns to arbitrary levels of order (Huberman and Hogg, 1986).

Intuitively, imagine the ripple analogy in Figure 1. Extend this analogy to a pond as in Figure 3. In reality, the drops will be uncorrelated with each other (Figure 3(a)). Never the less, imagine that drops local to one another do correlate above chance. Like an unfair

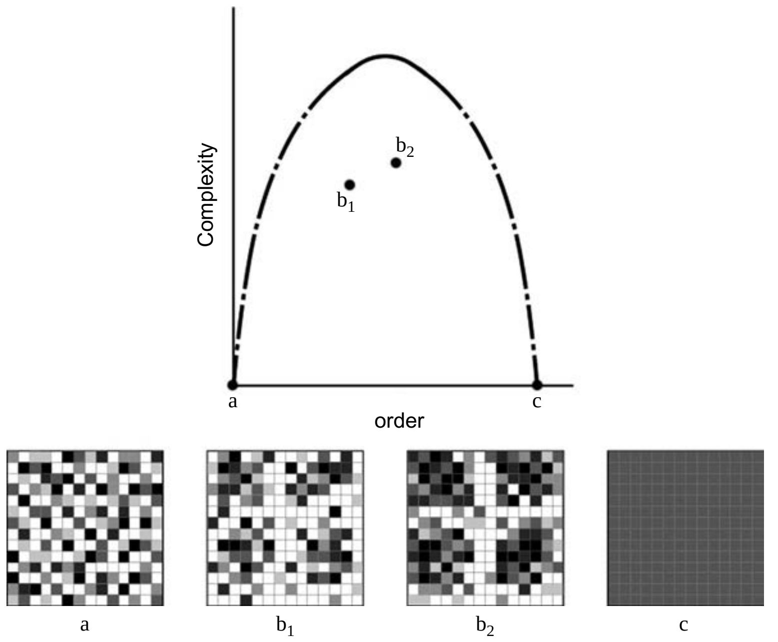

Notes: Order is along the *x*-axis from randomness, at point **a**, to complete order at point **c**, analogous to figure-grounds **a** and **c**; a second dimension is needed to asses levels of order; as adjacent greys correlates and these groups of greys correlate, higher-orders of complexity accrue along the *y*-axis; figure-grounds $\mathbf{b_1}$ and $\mathbf{b_2}$ correspond to points $\mathbf{b_1}$ and $\mathbf{b_2}$ in parabola space; life is in this parabolic balance

Figure 2.
Complexity as levels of order

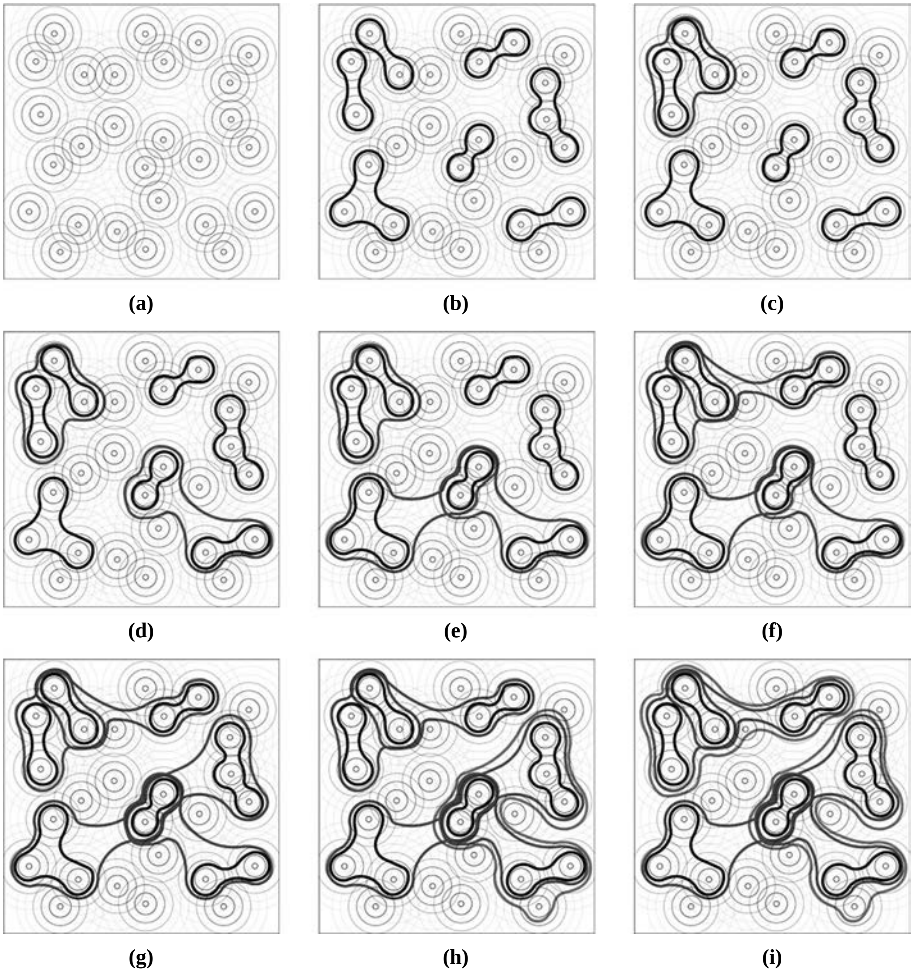

Notes: (a) Abstraction of random drops on a pond (0th order); (b) small encircled groups of drops indicate co-activity above chance (first order); (c)-(g) some of the first order activity co-occurring above chance (second order); (h) and (i) some of the second order activity co-occurring above chance (third order)

Figure 3. Imaginary co-activity of drops on a pond

die, small local groups of drops occur more than is expected (Figure 3(b)). Now imagine that some of these statistically established groups begin to correlate with each other above chance; as such, groups of groups result (Figure 3(c)-(g)). Further imagine groups of groups correlating with other groups of groups to statistically establish groups of groups of groups (Figure 3(h) and (i)), and so on. Although only an analogy, this behaviour is SoC. Distinctness simultaneously co-occurs with inter-relations.

Like a bunch of differently unfair dice, each die has its own distribution of outcomes independent of other dice in the bunch. This is a statistical indication of first order structure, or complexity. But, now imagine that somehow certain pairs of dice not only have their own distinct behaviours, but each die's outcome in the pair are also dependent on each other; this is second order complexity. The Appendix defines SoC more rigorously.

Yet, because SoC is a state variable like entropy, it can indicate nothing about how the system evolved to that state and whether this complexity contributes to work. The populated distributions of fair or unfair dice events says nothing about why it physically came to be that way. And, for any digital or cybernetic model, arbitrary levels of complex behaviour per SoC can be fabricated with no concern for how the system dynamically adapts to its changing environment in order to do adaptive work. To track this dynamic adaptation, we must measure work itself.

Adaptive work potential

Survival requires learning. Learning requires work. At every scale of the biosphere, energy is being converted by living systems to aid in their adaptation to their respective environments. And, in a complex environment, complex work must be done to enable such diverse and dynamic energy transformations. Although SoC is necessary for adaptive behaviour, it is not sufficient. The ability to do adaptive work is what makes complexity physically relevant.

To illustrate, imagine a diversely tuned array of tuning forks in a pool of water as in Figure 4. Now imagine that a C major chord occurs in this system's surroundings. The forks are tuned stochastically, but any tunings near the notes c, e, and g in the C major chord will be stimulated to vibrate causing a very distinct ripple pattern in the water (Figure 4(a)). This ripple pattern is simultaneously composed of both distinct events, individual fork vibrations, and their inter-relations via their superimposed ripples. Imagine further that connections are made between the active forks along the more salient interference (Figure 4(b)) so that these forks will be more directly stimulated by each other when stimulated by C major in the future (Bacigalupi, 2012b).

In addition to C major, other chords can stimulate distinct tunings in the same exact matrix, e.g. E minor in Figure 4(c) and (d). And, since these chords are stimulating the same exact matrix, they can co-occur to create the higher-order chord of C major 7th. In other words, once C major and E minor establish their own inter-connectivity, they exist to create the higher-order pattern C major 7th as in Figure 5. This hierarchical inter-relation is a higher level of complexity, i.e. increased SoC. And, unlike digital implementations, no new resources were required to create each of these chord patterns and higher-order patterns of patterns. Complexity accrues "vertically", i.e. along y-axis in Figure 2, using fixed resources. So, this form of memory and pattern association is efficient, but what about adaptive work?

The ability for this simple system to turn energy input into work increases as certain inter-connections are made. Again, the goal of this project is to realize an adaptive agent able to assimilate relevant patterns from its environment to do useful work. In the chord example, chord patterns are assimilated via accrued inter-connections, because they persisted in the environment. This strategy can be extended beyond acoustic frequencies to many diverse patterns and diverse sensory modalities, e.g. sensing different swaths of the electromagnetic spectrum. Once these patterns are instantiated in the system, they increase the capacity to do work by self-connecting to increase available power when patterns re-occur (Bacigalupi, 2012b). Put another way, a chord is "learned" when connections between relevantly tuned forks have accrued via co-stimulation (Figure 6). Now, instead of just each individual fork being sensitive to a small range of

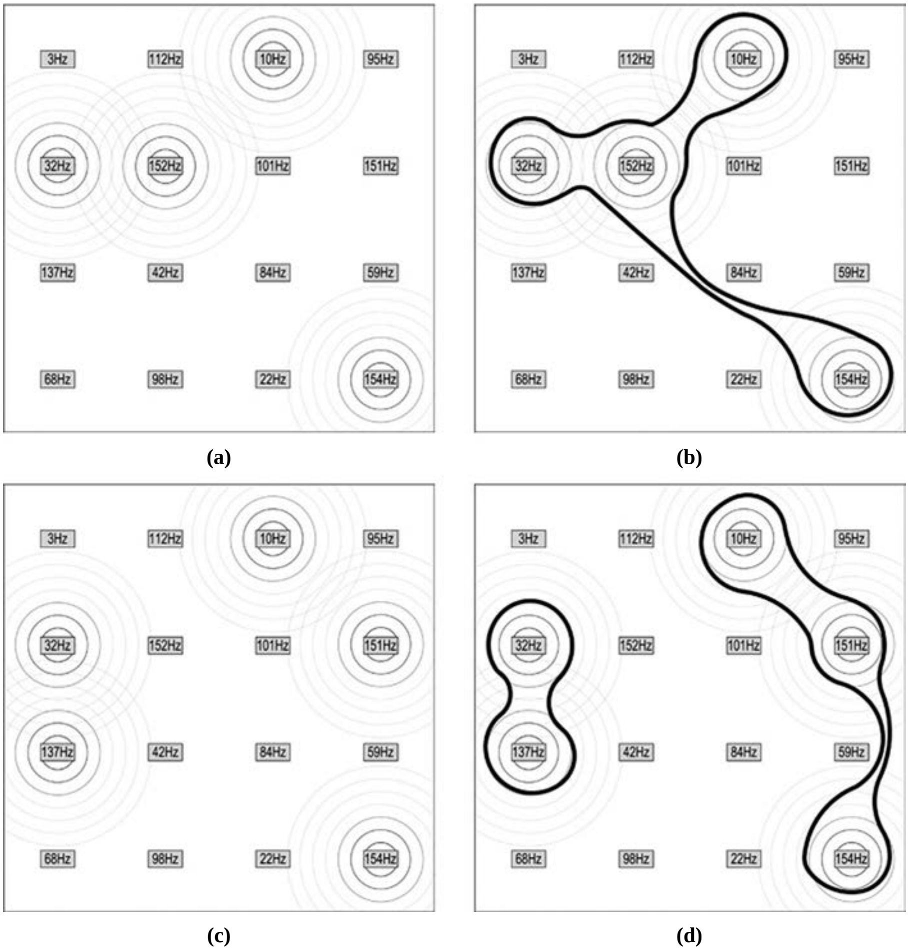

Notes: As groups of sparse nodes remain co-active, a system can be conceived that instantiates connections between such nodes in a common physical matrix; in this case, frequencies found in C major (a) and E minor (c) simultaneously stimulate specifically tuned nodes; (b) and (d) indicate the formation of biasing connections between these co-active nodes

Figure 4.
Instantiating persistent patterns

adjacent frequencies, many diversely tuned forks are jointly coupled and therefore sensitive to a higher-order constellation of frequencies because they previously "learned" the pattern. When this pattern re-occurs in the environment, this sensitivity literally means that the structure of the system is re-wired and therefore biased to respond more powerfully to this particular pattern, e.g. C major, E minor, and/or higher-order C major 7th. This power in turn increases this sparse group of forks' capacity to rise above the din of more weakly stimulated and less correlated forks. And this distinct response is a higher-order constraint, which increases the agent's ability to do higher-order work relevant to patterns in its environment.

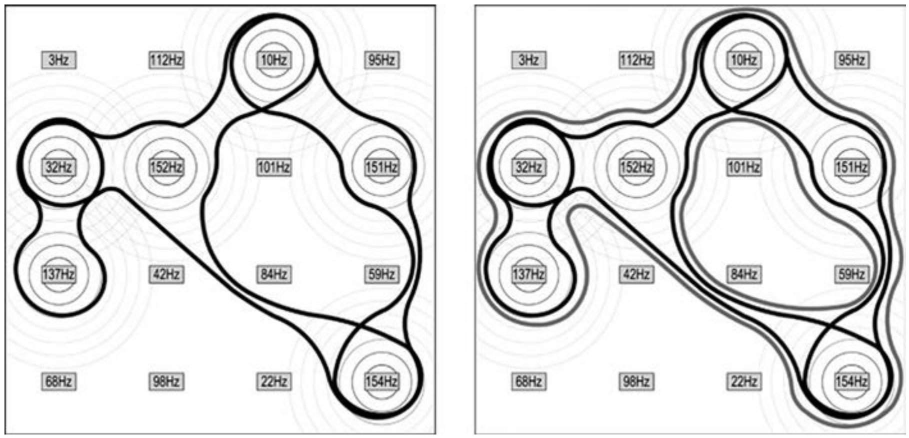

(a) **(b)**

Figure 5.
Higher-order patterns

Notes: Given that the system is already biased by the chords C major and E minor, these two patterns can now co-occur as in Figure 5(a) to create a higher-order pattern; the persistence of these two patterns, which is C major 7th, can also be instantiated (Figure 5(b)); C major and E minor are first order patterns, while C major 7th in Figure 5(b) is a second order pattern

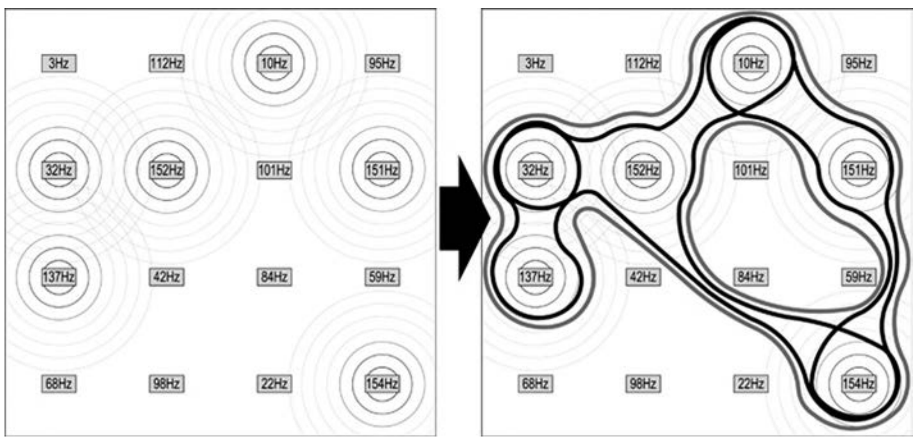

(a) **(b)**

Figure 6.
Adaptive work potential

Notes: As inter-connections are formed in response to patterns in the environment the system adapts to its environment by becoming biased at different levels of SoC (second order in this example); this bias increases the response power of the system when formerly instantiated, or “learned”, patterns re-occur; due to this power response, the potential to do work is larger with nodal inter-connections as in Figure 6(b) than without as in Figure 6(a); see Appendix for a more rigorous description of AWP

These salient "islands" of distinct activity above the noise floor are dynamic constraints that have accrued based on complex inter-relations among distinct patterns in the system's environment (Bacigalupi, 2012b). The work of recognition, association, and relevant adaptive behaviour is built upon a system's capacity to instantiate not only the distinct parts, but also how those parts relate to each other in time and space. It is the measure of both complexity and physical work within the agent that goes beyond current information theories. Refinement, a new understanding of how information evolves in autonomous adaptive agents, is a contribution to a new information theory inclusive of living systems.

Refinement rigorously defined

Refinement is the capacity for autonomous adaptive agents, like living organisms, to embody the complex patterns experienced in their environment so as to do adaptive work. This adaptive attribute increases the tendency for an agent's structural properties to persist in that environment, therefore remaining relevant.

Relevance is the ratio of the work done by all nodes in a system, e.g. tuning forks, neurons, etc. to the energy input into the system, multiplied by SoC. The work/energy ratio will always be less than one due to the second law of thermodynamics. Assimilated patterns create a work savings for the learned system in terms of pattern recognition, and potentially adaptive behaviour, which is beyond the scope of this paper, see Bacigalupi (2012b). As long as SoC increases on average over time and its structure is predominantly moulded on patterns in the agent's environment, Relevance of the agent to its environment will tend to increase on average over time.

The *x*-axis in Figure 7 is the development span of the agent from conception to death, whereas the *y*-axis is Relevance. Note that Relevance oscillates around an average trajectory. This trajectory is Refinement.

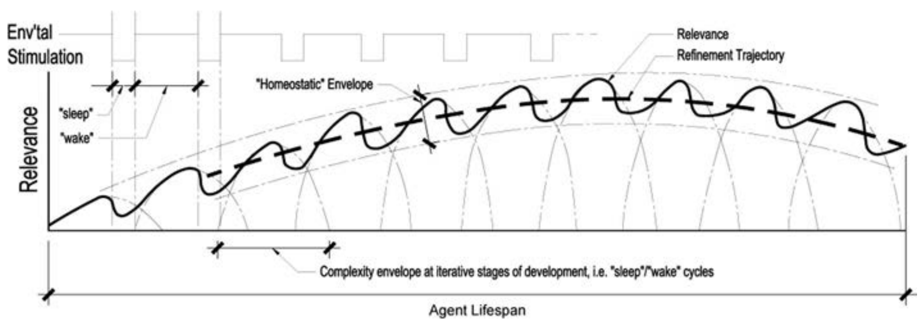

Notes: As the internal structure of the agent is biased by complex patterns in its environment it too becomes more complex; this structural complexity increases non-linearly in phase with the agent's iterative stimulation by its environment; each new "complexity envelope" creates the foundation upon which an agent may perceive and assimilate new patterns from its environment; the curve of assimilation is Relevance; environmental patterns that exist beyond the agent's current complexity envelope are imperceptible to that agent; the average trajectory of Relevance is the Refinement trajectory

Figure 7.
Refinement trajectory

It is hypothesized that Relevance will tend to oscillate on a regular cycle in order for Refinement to increase. A natural oscillatory feedback is created between formerly assimilated patterns and new patterns. A system must unlearn a little in order to learn more. It must un-wire itself a little to accommodate new patterns and their more complex inter-dependencies. New patterns will necessarily compete with existing order via the internal interference patterns. And it is the differentiable regions, or sparse order, of this structured noise that can be harvested and structurally instantiated to increase the potential of doing relevant work in the future. Paradoxically, this temporary disorder creates the opportunity for new, formerly unknown, orders. Living systems "ratchet" levels of order upon previously assimilated patterns. However, in order to maintain cohesion, or "meaning", subsequent patterns must literally weave their inter-connections among previous inter-connections. In other words, circuits must necessarily re-wire existing connections. In a very real way, death at every scale of pattern is a necessary precondition for the increased persistence of life.

Refinement, although a rigorous framework (see the Appendix), is not a theory as yet. It is a schematic design for building and empirically testing synthetic cognition. This short paper introduced the theoretical basis to test for behaviours necessary for autonomously learning complex patterns about an agent's environment. Once actual hardware is built, it will be exposed to both simple and more complex composite patterns. If repeated exposure to a diversity of such patterns does in fact increase measured Refinement, as defined, the alternative hypothesis can be selected. If not, the theory and possibly the hardware in Bacigalupi (2012b) will need to be modified.

References

Bacigalupi, A. (2012a), *An Abstract for Synthetic Cognition: Implementing a Testable Hypothesis* (*Harvesting Sparse Order from Local Interference*), available at: http://iofam.net/research/publications#SynthCog_Abstract (accessed 28 July 2013).

Bacigalupi, A. (2012b), *Synthetic Cognition Implementing a Testable Hypothesis. A Research and Development Proposal*, available at: http://iofam.net/research/publications#SynthCog_Proposal (accessed 28 July 2013).

Deacon, T. (2011), *Incomplete Nature*, W.W. Norton & Company, New York, NY.

Huberman, B.A. and Hogg, T. (1986), "Complexity and adaptation", *Physica D*, Vol. 22 Nos 1/3, pp. 376-384.

Longo, G., Montevil, M. and Kauffman, S. (2012), *No Entailing Laws, But Enablement in the Evolution of the Biosphere*, arXiv:1201.2069v1 [q-bio.OT], available at: http://arxiv.org/abs/1201.2069v1 (accessed 28 July 2013).

Roy, J.E. (2002), "The neurophysics of consciousness", *Brain Research Reviews*, Vol. 39 No. 1, pp. 1-28, Appendix B.

Appendix. Mathematical development of refinement

State of complexity

SoC, a state variable, is defined by not only summing the normalized variance of distributions of individual events, First order distinctions, but also summing the variance of the distribution of inter-relations between multiple distributions, higher-order "patterns of patterns". SoC is increased in three main ways:

(1) more events above chance at each level of order;
(2) more levels of order; and
(3) the increased distinctness of events in (1) and (2) above chance.

Satisfying these relations, SoC is defined as follows.

Definition 1:

$$\mathrm{SoC} \equiv \sum_{\mathrm{1st}}\left(1/\sigma_{\mathrm{i}}^{2}\right)+\sum_{\mathrm{2nd}}\left(1/\sigma_{\mathrm{i}}^{2}\right)+\sum_{\mathrm{3rd}}\left(1/\sigma_{\mathrm{i}}^{2}\right)+\cdots+\sum_{\mathrm{nth}}\left(1/\sigma_{\mathrm{i}}^{2}\right)$$

Adaptive work potential

For a circuit-based implementation, as in Bacigalupi (2012a, b), work is as follows:

$$\mathbf{Work} = \int \mathbf{I}\times\mathbf{V}d\boldsymbol{\tau}; \tag{A1}$$

whereby:

- **I** and **V** are phasors, **I** is alternating current, **V** is alternating voltage, and **I × V** = real power.

For the project specified in Bacigalupi (2012b), the work done by an adaptive system over time interval [a, b] is:

$$\mathrm{Work}|_{\mathrm{b}}^{\mathrm{a}} = \sum_{i}^{n}\int_{\mathrm{b}}^{\mathrm{a}} \mathrm{I_i}\times\mathrm{V_i}d\tau; \tag{A2}$$

whereby:

- **I × V** is power as in equation (A1), and its integral is the amount of work done for the ***i***th node over the time interval [a, b]. The work done for ***n*** nodes is then summed over that same interval.

AWP is exhibited by an agent of fixed nodes, ***n***, when that agent tends to do more work over interval [**x, y**], due to accrued internal structural connections, than the previous interval [**a, b**] of the same duration and given the same environmental stimulations:

$$\mathrm{Work}|_{\mathrm{b}}^{\mathrm{a}} = \sum_{i}^{n}\int_{\mathrm{b}}^{\mathrm{a}} \mathrm{I_i}\times\mathrm{V_i}d\tau < \mathrm{Work}|_{\mathrm{y}}^{\mathrm{x}} = \sum_{i}^{n}\int_{\mathrm{y}}^{\mathrm{x}} \mathrm{I_i}\times\mathrm{V_i}d\tau \tag{A3}$$

Relevance

Relevance is defined as follows.

Definition 2:.

$$\mathbf{Relevance} \equiv \mathbf{SoC}\times\left(\sum \mathbf{W_i}/\mathbf{E_{input}}\right);$$

whereby:

- **SoC** is state of complexity in Definition 1, assessed at the latter bound of AWP, e.g. at time **y** of the interval [**x, y**].
- $\sum\mathbf{W_i}$ is the summed work of ***n*** nodes over specified interval as in equation (A2).

Refinement

Refinement is defined as follows.

Definition 3:

- Refinement ≡ curved regression of Relevance in Definition 2, as shown in Figure 7.

About the author

J. Augustus Bacigalupi studied environmental sciences and physical chemistry at UC Santa Barbara. He then received an MArch from CU Denver. His design practice and science education cross-fertilized to inspire a decade of independent research in the field of cognitive science. He is currently working with Professor T. Deacon who has independently developed similar ideas. J. Augustus Bacigalupi can be contacted at: augustus@iforam.org