The Complexity of Cognition

Issues in Emotion and Consciousness Essay
EASY MSc
Department of Informatics
University of Sussex - Brighton - United Kingdom

Abstract

This paper provides a background to the three main approaches to understanding cognition which have dominated the past century; culminating in an overview of newly proposed "complex systems" perspective which seems to bridge some of the gaps between them.

1 *Computationalism*

Theories of cognition and consciousness have been dominated by the combined forces of computation and functionalism for most of the last sixty years. The "mind" had long been thought to be an ethereal realm in which man performed his logical manipulation of facts but when Alan Turing introduced the world to his Universal Computation Machine (Turing, 1946) the abstraction of this domain from the physical world was demystified. For the first time, scientists and philosophers were presented with a device capable of the logical manipulation of symbols within a kind of abstract mathematical domain running parallel to the physical world. Mind had been reduced and cognition was simply the process of informational computation.

The von Neumann architecture of early computers led software engineers and the early practitioners of this new cognitive philosophy down the path of functionalism. Software had to be built by nesting logical functions which called upon symbols either held within a memory or input by some form of electronic sensing device; and philosophy naturally took the same tack. Pain, happiness, hunger were all reduced to mathematical symbols within a computational device. All thoughts, ideas and feelings were considered elements of a hierarchical, functionalistic information processing system (Newell & Simon, 1976).

However, as the functionalistic approach matured during the 1970s it began to show cracks when placed under scientific and philosophical scrutiny.

Neurologists pointed out that the sequential architecture of the physical computing devices bared no semblance to the actually physiology of the brain (Crick, 1990). The brain is a highly distributed parallel processing machine without any discernable functional units such as a central processing unit or a memory block. Parallel processing simply doesn't work in the same functionalistic way as sequential processing. An entirely new kind of software engineering is required and, so too, an entirely new kind of philosophical counterpart.

Many philosophers argued against the completeness of the functionalistic approach to consciousness on the grounds of phenomenology; outlined by David Chalmers as "the Hard Problem" (Chalmers, 1995). John Searle's infamous Chinese Room questioned the capacity for a symbol manipulator to "understand" the symbols that it was manipulating (Searle, 1980). Jackson's colour blind Mary questioned the completeness of physicalistic description of the world (Jackson, 1986). Functionalism, it seemed, severely neglected the experiential aspects of consciousness.

However, the greatest damage to the purely computational theory of mind was done by its most valued ally; Artificial Intelligence herself. In trying to put the theory into practice, it was soon discovered that symbolic information gathering and manipulation suffered from a number of intractable flaws. Real-time dynamics (Fodor, 2000), the frame problem (Pylyshyn, 1987), the symbol grounding problem (Harnad, 1990), and knowledge acquisition all caused major problems; not just due to engineering constraints, but due to feasibility issues. The theoretical framework simply wasn't up to the job.

2 Connectionism

At around the same time as Turing brought us his computational machine, Frank Rosenblatt began to sow the first seeds of the connectionist movement with his model of a perceptron (Rosenblatt, 1958) - later made popular by Marvin Minsky (Minsky M. & Papert S., 1969). The perceptron is a simplified mathematical model of the mechanisms of a neuron. It receives a weighted input and when this input reaches a certain *action potential*

it fires an output. By creating a large interconnected network of perceptrons we reach the heart of the connectionist movement – the neural network (Bryson & Ho, 1969).
A neural network is collection of perceptrons connected in layers. The first layer of neurons receives a weighted input from an external source (e.g. sensory systems). All subsequent layers receive their inputs from the previous layer until an output layer transmits its signal to an external source (e.g. motor systems).

A subtle change of the weights within the network will have a subtle change in the input/output mapping and many algorithms have been devised to allow for the supervised or unsupervised updating of weights in order to "train" a neural network to perform a particular function (Gallant, 1990). As such, a connectionists function is processed in parallel by the entire network and information distributed throughout; unlike computationalism, where functions were engineered using symbolic representations and sequential logic.

Cognition, according to a connectionist, is simply the emergent property of a highly distributed parallel processing system – such as a neural network (Churchland & Sejnowski, 1992).

3 Behavioural Dynamics

A third cognitive perspective also began to grow from the 1940s, starting with Norbert Wiener notions of feedback and homeostasis (Wiener, 1948). These notions, later put under the banner of the cybernetics movement, were a fundamentally new way of examining control systems. It was shown that, given feedback, a system that changed over time could be made to self-regulate – just as a thermostat, given the feedback of temperature, is able to self-regulate the ambient temperature of your house.

Feedback is now just a small part of a larger mathematical field which studies the way non-linear systems in general change over time – the field of dynamical systems. Every dynamical system has a number of variables which change over time (by some deterministic equation). For example, a pendulum may be considered to be a dynamical system which has the variables of weight, length of cord, longitude, latitude and height. In this instance, the weight (w) of the pendulum changes over time with the equation $w^1 = w^0$ (in other words, it doesn't change at all). However, the height of the pendulum might change at a rate $h^1 = h^0 + sin(x) + \ldots cos(y)^2$. In other words, it is entirely possibly to determine the next value of h by following a simple, deterministic formula. Each and every possible combination of variables (e.g. w = 1, h = 2, x = 3.332, y = 0.232, l = 0.01) is called a *state* of the system; and in a dynamical system, every state will lead deterministically to another.

Such deterministic movement from state to state is said to be a *trajectory* through *state space*, and the imaginary diagram of all of these possible *trajectories* is said to be the *topology* of *state-space.* Within this complex geometric landscape there are troughs and peaks of varying shape and size, known as *attractors,* which are important ways of describing the long-term result of a dynamic system. In the state-space of our pendulum, for example, there is one single attractor to which all states will eventually lead; the state of equilibrium - a simple dangling pendulum.

The dynamics systems approach looks at life in the same light; all animals are a kind of dynamical system flowing deterministically from one state to the next. In fact, according to Maturana and Varela, a living organism is simply defined as a system whose dynamics are autopoietic - or self making (Maturana & Varela, 1980). In such living systems, cognition is the "process" of staying alive; it is the process of life (Maturana & Varela, 1980).

This new view of life and cognition is entirely behavioural, and completely free from internal representations; a stark difference to the computational method. There are no symbols being manipulated. There is no logical functionalism. There is just the mechanical process of staying alive. If an agent looks like a duck, and portrays the dynamical behaviours of a duck then it is a duck and possesses everything that "it is like to be" a duck.

A dynamic agent is seen as a continuous flow of sensory motor actions. A living organism's many senses are linked through to its motor systems via a complex maze of connections (its central nervous system). This complex maze (a dynamic system in its own right) ebbs and flows with respect to the continual and varied waves of stimulation it receives from the environment; these changes in the dynamic flow of the CNS constitute shifts in cognitive trajectories through "mind-space". The environment is, therefore, a key constituent to the behaviour and cognition of the agent (situatedness), and the agents entire body (its collection of autopoietic sensory and motor systems) is the coupled "information processor" (embodiment) (Brooks, 1991a; Brooks 1991b).

The field of Artificial Life (Langton, 1989) is to the dynamic systems approach what Artificial Intelligence was to the computational model; an engineering discipline attempting to implement and examine the many facets of the new theoretical framework. ALife practitioners working with embodied, dynamical agents have produced a wide variety of interesting behaviours from obstacle avoidance (Braitenberg, 1984; Beer, 1995) to more complex planning tasks (Floreano & Mondada, 1996). The plasticity of such systems is remarkable, as is the complete lack of symbolic representation and functional computation. The agents are simple capable of exhibiting the correct behavioural dynamics.

However, the approach is no panacea. Many philosophers (Kirsch, 1991; Kosslyn, 1980) are uncomfortable with the dynamicist's mantra about representational free cognition. For example, it seems possible for a person to rotate an imaginary cube inside their head; a feat which distinctly appears to require the manipulation of an internal representation. Similarly, there are many tasks that humans, and indeed "lower" animals perform which require complex planning (Tarsitano & Andrew, 1999).

Once again, neurologists too are beginning to enter the critical debate. Recent work (Quiroga Kreiman Koch & Fried, 2005) has indicated that our brains do indeed contain internal representations, of a kind. Using electrical implants positioned within the brains of seizure patients, they found that some neurons fired only when the patient was presented with a rather specific visual stimulus (their tests used images of famous faces). One patient had a particular neuron fire exclusively to images of the actress "Halle Berry". Astonishingly, this didn't just include clear photographs, but included more obscure images[1], pencil-drawn pictures and even just to the written words of her name. "*This neuron is responding to the concept, the abstract entity, of Halle Berry*" (Quiroga, 2005).

In fact, a deeper examination of some of the work put forward by dynamicists as examples of representational free planning (Floreano & Mondada, 1996) actually shows similar neural correlates. The dynamics systems approach, it seems, is slowly being redefined as a mere extension of connectionism.

4 *Complexity*

A sub-field to the mathematics of dynamic systems is the field of complexity science. In the early 1990s, a number of pioneering mathematicians began to study the behavioural patterns of dynamical systems in general and discovered that systems tend to display one of three types of behaviour - ordered, chaotic or "complex"[2] (Wolfram, 1984; Langton, 1990; Wuensche, 1996). The field of complexity was born out of the study of this sub-domain of "complex" dynamical systems (and its relation to the other sub-domains).

"Complex" systems have one extremely evocative property in relation to our examination of cognition; they are dynamic systems which show a capacity for emergent computations (Berlekamp Conway & Guy, 1982; Mitchell Crutchfield & Das, 1997; Mitchell Hraber & Crutchfield, 1993; Crutchfield & Mitchell, 1995).

A Cellular Automata (CA) is a particularly useful tool for gaining an understanding of dynamical systems in general and is widely used within the field of complexity science. John Conway's, now famous, Game of Life demonstrates quite beautifully this notion of emergence; simple underlying rules resulting in complex dynamical particles such as gliders, glider guns and blinkers. In fact, it has been mathematically proven that the

[1] Her dressed up as Catwoman, wearing a mask which covered most of her face.

[2] When referring to "complex" systems in this sense, it is meant Wolfram Class IV systems

emergent particles of the Game of Life could be arranged in such a fashion as to act as a Universal Computation machine (Berlekamp Conway & Guy, 1982).

In 1995, James Crutchfield and Melanie Mitchell, of the Santa Fe Institute, evolved CAs to perform simple computational tasks and then analysed the emergent *particles* and the dynamics of these particles; explaining it in terms of a kind of "*particle logic*" (Mitchell Crutchfield & Das, 1997; Mitchell Hraber & Crutchfield, 1993; Crutchfield & Mitchell, 1995). For the dynamic system to perform any kind of computation, it seemed that it generated a number of emergent symbols and a logic of symbolic manipulation.

"Dynamics, computation and adaptation are beginning to be viewed in a more unified framework" (Mitchell, 1998).

Complexity theory is still a fledgling field and the nature of emergence yet to be fully understood; but the "*creation*" of these "*informational particles*" might invoke new and novel insights into phenomenology of conscious experience.

5 Conclusions

In studying the various disciplines it has become apparent that distinctions between them may not be so clean cut after all. A brand of connectionists' neural network is usually the tool of choice for the dynamicists and in fact many authors debate whether or not the distinction between the two fields actually exists (Eliasmith, 1996). Complexity theory, on the other hand, seems to show how a highly distributed dynamical system can perform a kind of emergent symbolic computation. In fact, for human cognition in general, it seems that a distributed, connectionist's architecture may provide the physical substrate for a "complex" dynamical system performing emergent computations.

Can our ethereal realm of "mind" be reconsidered as the emergent realm of "*particles*"? Is the flow of thought and logic isomorphic to the flow and interaction of these non-physical particles?

6 References

Beer R. (1995) "*On the dynamics of small continuous-time recurrent neural networks.*" Adaptive Behavior (3) 469-509.

Berlekamp E. R., Conway J. H., Guy R. K. (1982) "*Winning Ways for Your Mathematical Plays, Vol. 2: Games in Particular*", Ch 25, London: Academic Press

Braitenberg V. (1984) "*Vehicles: Experiments in Synthetic Psychology*", Bradford

Brooks R. (1991a) "*Intelligence without Representation*", Artificial Intelligence Journal (47), pp. 139–159.

Brooks R. (1991b), "*Intelligence without Reason*" Proceedings of 12th Int. Joint Conf. on Artificial Intelligence, pp. 569–595.

Bryson A.E., Ho Y. (1969) "*Applied optimal control*". New York: Blaisdell.

Chalmers D. (1995) "*Facing up to the problem of consciousness*", Journal of Consciousness Studies 2 (3), pp. 200-219

Churchland P. S., Sejnowski T. (1992) "*The computational brain"*, Cambridge, MA, MIT Press

Crick F., Koch C. (1990) "*Towards a Neurobiological Theory of Consciousness*", Seminars in the Neurosciences. vol 2 pp 263-275.

Crutchfield J.P., Mitchell M. (1995) "*The Evolution of Emergent Computation*", Proceedings of the National Academy of Sciences, USA 92:23 10742-10746.

Eliasmith C. (1996). "*The third contender: A critical examination of the dynamicist theory of cognition*." Philosophical Psychology. Vol. 9 No. 4 pp. 441-463

Fodor J.A. (2000) "*The Mind Doesn't Work That Way"*, Cambridge, MA: MIT Press

Floreano D., Mondada F. (1996) "*Evolution of Homing Navigation in a Real Mobile Robot*". Cybernetics, 26(3), 396-407.

Gallant S. I. (1990) "*Perceptron-based learning algorithms*", IEEE Transactions on Neural Networks, vol. 1, no. 2, pp. 179-191.

Harnad S. (1990) "*The Symbol Grounding Problem*", Physica D 42: 335-346

Jackson F. (1986) "*What Mary Didn't Know*", Journal of Philosophy (83) 291-295.

Kirsch D. (1991), "*Today the earwig, tomorrow man?*" In Artificial Intelligence, V 47, pp 161-184

Kosslyn S.M. (1980) "*Image and mind*", Cambridge, MA: Harvard University Press.

Langton C.G. (1989), "*Artificial Life*" In Artificial Life. Addison-Wesley. pp. 1-47.

Langton C.G. (1990) "*Computation at the edge of chaos: Phase transitions and emergent computation*", Physica D vol. 42, pp. 12-37.

Maturana H., Varela F. (1980) "*Autopoiesis and Cognition: The Realization of the Living",* Vol. 42: Boston Studies in the Philosophy of Science

Minsky M., Papert S. (1969) "*Perceptrons*", Cambridge, MA: MIT Press

Mitchell M., Crutchfield J.P., Das R. (1997) "*Evolving Cellular Automata to Perform Computations*"

Mitchell M., Hraber P.T., Crutchfield J.P. (1993) "*Revisiting the Edge of Chaos: Evolving Cellular Automata to Perform Computations*", Complex Systems 7, pp. 89-130.

Mitchell M. (1998) "*A Complex-Systems Perspective on the "Computation vs. Dynamics" Debate in Cognitive Science",* Twentieth Annual Conference of the Cognitive Science Society

Newell A., Simon H. (1976) "*Computer science as empirical enquiry: Symbols and Search*", Communications of the Association for Computing Machinery, 19(3), 113-126

Pylyshyn, Z.W. (1987), "*The Robot's Dilemma: The Frame Problem in Artificial Intelligence"*, Norwood, NJ: Ablex.

Quiroga R., Kreiman G., Koch C., Fried I. (2005) "*Invariant visual representation by single-neurons in the human brain*", Nature, 435: 1102-1107

Rosenblatt F. (1958) "*The Perceptron: A Probabilistic Model for Information Storage and Organization in the Brain*", Cornell Aeronautical Laboratory, Psychological Review, v65, No. 6, pp. 386-408.

Searle J. (1980) "*Minds, Brains, and Programs*", the Behavioural and Brain Sciences, vol. 3

Tarsitano M.S., Andrew R. (1999) "*Scanning and route selection in the jumping spider Portia labiata*", Animal Behaviour, Volume 58, Number 2, pp. 255-265(11)

Turing A. (1946) "*Proposed Electronic Calculator*", National Physical Laboratory

Wiener N. (1948) "*Cybernetics, or Control and Communication in the Animal and the Machine",* New York: Wiley

Wolfram S. (1984) "*Universality and complexity in cellular automata*", Physica D vol. 10

Wuensche A (1996) "*Attractor Basins of Discrete Networks*", D.Phil thesis, University of Sussex