

Emergent Artificial Intelligence*

Graeme Taylor

November 26, 2004

Note: there are two interpretations of “emergent” with respect to AI. One relates to theories of technological singularity- and in particular the birth of AI through sheer weight of connectivity and processing power thrown up by such an event. Since speculation on post-singularity science and society seems even more fraught with hazards than general AI discussion, this is not the focus of this writeup. Rather, it concentrates on the less fantastical (but still far reaching) view of success in strong AI being dependant on holistic systems whose intelligence is a function of their perception of, and interaction with, the real world.

Prelude: Knowledge-based Artificial Intelligence

As AI research began to organise into a coherent subject in the 1950s and 60s it was (unsurprisingly) unclear as to what route offered the most promise in creating machine intelligence, or even what such intelligence would look like. The essence of a knowledge-based approach was captured by John McCarthy in 1956- “*making a machine behave in ways that would be called intelligent if a human were so behaving.*” This encompasses tasks such as playing chess, proving theorems, or engaging in convincing conversation. Such a framework disconnects intelligence (in the sense of doing smart things) from the broader human experience of mental activity- an AI that interacts only in a virtual environment could be entirely bereft of emotion, for instance.

Emergent Artificial Intelligence

Instead, a holistic view can be argued- that the only example we have of intelligence operating in the real world is ourselves, and that whilst our intelligence comprises of numerous separate abilities- from reflex reaction through to day-dreaming- it is their combination into perception and a model of the world that drives our conscious experience. Thus to isolate any one aspect of intellect from the mind as a whole is to miss the point.

*First appeared on Everything2, at http://www.everything2.com/index.pl?node_id=1686321

The emergent view of AI is by necessity strongly tied to the real world, in two ways- technologically, through A-life research and robotics; and through research into biological intelligence, striving for greater understanding of the human brain. It is a younger field, since the restrictions of early-era computer kit meant that such equipment was more suited to the functional decomposition of knowledge-based AI into individual high-level tasks rather than juggling many low-level tasks towards creating an intelligent whole.

A glass of water, please

In a recent lecture I attended, Steve Grand (more on him later) offers a good illustration of this fundamentally different approach to AI, which has picked up over the last decade or so. He offers the following example- that you are trying to build a robot to fetch a glass of water. Try a google image search for “glass”, “water” and “glass of water” (as any self-respecting and net-savvy AI would) and you’ll get quite a range of material to interpret: it’s unlikely that presenting a vase of water (recently minus its flowers) would win a robotic bartender many advocates (well, right now even that might impress, but you see the point). But set aside for the moment the difficulty of understanding the physics of containers or the mechanics of flowing water, and just contemplate the much simpler task of picking up a random glass from a cluttered table in the first place. If we offer (reasonably) a robotic arm with 22 degrees of freedom, with each capable of 10 different positions, then there is probably enough mechanical flexibility to physically accomplish the task. But can its completion really be reasoned out through logic? The maths is frightening- to try all the permutations of positions of your robot arm would require longer than the current history of the universe to complete at a rate of one change per second. It seems hopeless.

Yet you can usually reach for a glass unthinkingly. How? Your binary vision lets you gauge depth, whilst the movement of an outstretched arm confirms the mental prediction made. Over time, repeated interactions of these two systems have reinforced one another into a body of experience about best practice- but you’d be hard pressed to describe that experience mathematically.

Perhaps this is the defining feature of an emergent view of AI. We would all credit a chess grandmaster with intelligence. But the chess-playing alone doesn’t suffice: there’s a default level of competence assumed through their being human, which if absent would be considered signs of unintelligence. Emergent AI aims for these core competencies- whilst you may look foolish losing a game of chess, you’d look immeasurably more so if you fell over on attempting to leave your chair at the board due to a complete lack of a notion of balance.

Biological Inspiration

These pre-requisites for intelligence have emerged throughout biology, and are manifest even in creatures that wouldn’t be considered intelligent themselves. For instance, much A-life research

revolves around recreating the focused efficiency of insects in both robots and digital simulations. But the greatest inspiration is inevitably our own brains. I've discussed before the difficulties of trying to mimic such a marvel- ironic given emergent AI's philosophy that it might try to run before it can walk.

One difficulty is that neurobiological research and related cognitive study can tell us much about the mechanical workings of the brain- the plumbing, so to speak- without giving any indication of the underlying principle upon which it is organised. Evidently the brain has a purpose- the signal delay between sensation and action alone is enough to warrant an evolutionary advantage through having a device to predict the future- of being able to ask, What next? It's harder to explain why we should be able to explore mental models of What If?, but thinking about the limits of imagination itself suggests that it all ties back to perception- for who could imagine what they couldn't perceive?

Nuts and Bolts

So armed with an emergent interpretation of AI, it becomes apparent that to mimic human intelligence, a machine needs to be able to do more than manipulate virtual environments, but cross reference those with experiences of and actions in the real world. Robot building can be seen as a hindrance in the classical approach to AI- at first when computing resources were limited, the capability for processing the environment wasn't there; then advances in processing power made simulation far more economical, as experiments can be tweaked in terms of program parameters instead of costly rebuilds. Furthermore, any shot at emergent AI needs to master not just the intelligence but the mechanics of motion, vision, balance and so on.

A daunting task for sure, but attempts are being made. Two notable examples lie at either end of the spectrum. One is Cog, the product of MIT AI Lab's wealth (both financial and intellectual) under the direction of Rod Brooks, a strong advocate of emergent AI and such "situated robotics". As with program-based AI ventures, robots still only prosper in very restricted domains of experience, and there is the problem of scaling low-level successes into general intelligence. Whilst Cog can explore the connections between perception and action by existing in the real world, it nonetheless currently fails to reliably perceive the difference between a mobile phone and a glasses case.

At the other end of the funding spectrum, and hence free of the restrictions of board meetings and other academic trappings is Steve Grand's robot Lucy- who, at her peak, could grasp the invariant properties of a banana well enough to perform her party trick- pointing at a banana, regardless of its orientation or the extent to which it had been peeled. This from a machine built in a garage in Somerset by one lone researcher, proof as good as any that an AI breakthrough could come from anywhere- formal academia, informal research, military projects or commercial ventures.

Meeting in the middle

Emergent AI is intertwined with biological research in a symbiotic way- Grand admits that Lucy is not an end in itself but simply a tool for better understanding human minds; whilst greater understanding of the machinations of the brain can point the way to better creation of synthetic minds. Ultimately, developments in cybernetics, through mechanical improvements to the body or our ever-increasing use of intellectual augmentation (stop and consider how much more useful you are with google or E2 at your fingertips) may blur the lines between man and machine to the extent that once AI is achieved the 'artificial' distinction has become irrelevant. Perhaps this merger would need to be three-fold- of symbolic AI, situated robotics and biological components. Or emergent AI may prove to just be another waymarker on the journey to strong AI, rather than the entire path. At this stage its too early to tell, and such speculation, like singularity theory, may best be left to science fiction.

Reference and related reading

- On 22/xi/04 I attended a lecture by Steve Grand entitled *Machines Like Us* (part of the millenium lectures series), which is the source of many of the examples in this writeup and also the inspiration to write it in the first place.
- Steve's website is at <http://www.cyberlife-research.com>, and has some discussion of AI and A-life
- "*Oneworld Beginner's Guides- Artificial Intelligence*": Blay Whitby, ISBN 1-85168-322-4
- <http://encyclopedia.thefreedictionary.com/Artificial%20intelligence>- for the McCarthy quote, but looks a good jumping-off point for further reading.
- *Artificial Life- The Quest for a new creation*: Steven Levy, ISBN 0-14-023105-6 explores the bottom-up, biologically inspired approach to AI.
- *Ghost in the Shell* explores the boundaries between man and machine and raises the question of what a cyborg could cling to as a sense of humanity.
- William Gibson's *Neuromancer* explores the other type of emergent artificial intelligence- of AI getting loose in the 'net.