Intelligence Is Collective

EMILE SERVAN-SCHREIBER, Lumenogic LLC

How is collective intelligence different from the regular kind? The cognitive sciences have produced richly detailed computational models of how the brain produces the mind, as well as powerful artificial intelligence programs ("AIs"). It is intriguing that in each case – natural or artificial – the intelligence of the system can be described as "collective", not only in the sense that it arises from the interactions of dumber parts and processes – that isn't new – but that it does so according to specific principles better known for extracting wisdom from crowds.

In *The Wisdom of Crowds*, Surowiecki (2004) identified the four criteria that separate smart groups from dumb ones: members of the group hold a variety of opinions (*diversity*) that draw on their own specialized or localized knowledge (*decentralization*), and are able to express them without being unduly influenced by others (*independence*). The final requirement is an appropriate mechanism to put it all together and determine the group's behavior (*information aggregation*). Surowiecki was of course referring to groups of *people*, but I will argue that his elegant four-ingredients recipe for crowd wisdom can be generalized to groups of information processing devices of any kind, collaborating to produce intelligence in, for instance, a brain or a computer. In other words, it is a recipe for intelligence in general.

SUROWIECKI MACHINES

Let us call "Surowiecki machine" any information processing system made up of simpler devices that collaborate according to wisdom-of-crowds principles: (1) the devices contribute their input to the machine's aggregation mechanism, which then computes the machine's output; (2) the contributions are diverse because each device is processing information in different ways (specialization), or because each device is processing different information (localization), or both; (3) each device carries out its processing mostly independently of the others, and sends its output directly to the machine's aggregation mechanism with little interference from other devices.

I will argue that the human mind and a variety of best-of-class AI programs can be described as Surowiecki machines. I will attempt to support this claim, not through formal proof, but by examining selected examples of intelligent systems that capture the state of the art in cognitive neuroscience, and artificial intelligence.

THE MIND IS A WISE CROWD

"Minds are simply what brains do," famously wrote Minsky (1986), so it is suggestive that neurons, the basic building blocs of cognition, are themselves set up as Surowiecki machine: Each neuron receives *independent* activation or inhibition input signals from a large number of other neurons near and far (insuring *diversity* and *decentralization*). Then, it *aggregates* this information by summing up the excitatory and inhibitory signals, and decides, if its activation reaches a threshold, to transmit its own signal to another large number of neurons.

But what about the organization of the mind itself? While far from complete, our understanding of human cognition is now such that at least one richly detailed model of its “architecture” offers a unified theory of memory, learning, problem solving, perception, and action: the ACT-R theory of mind can explain complex thinking not only at an abstract “symbolic” level but also predict *where* and *when* it happens in the brain (Anderson, 2007; Borst & Anderson, in press). For instance, it can simulate the timing and localization of activation in the brain while a subject solves an algebraic equation.

According to ACT-R the mind is organized in independent functional modules, each of which can be localized in the brain. Among them are sensory modules for visual and aural perception, motor modules for manual and vocal response, and other modules dedicated to internal processing such as mental representation (Imaginal), memory for facts (Declarative), and goal management. A central *Procedural* module is the gateway through which all the other modules communicate: it coordinates the cognitive flow in the system through a large set of “If-Then” rules that respond to other modules’ stimuli by sending them appropriate requests and available information. Importantly, only one rule may be triggered at a time, which takes approximately 50 ms and determines the tempo of cognition.

It is striking, and non-trivial, that a theory which represents the state of the art in cognitive neuroscience’s quest to understand of the architecture of cognition should be so congruent with a Surowiecki machine: The central Procedural module plays the information-aggregation role – by selecting *the* rule that fires in response to the input from all other modules – while each module independently processes and contributes a very specific type of information, insuring diversity, and specialization. The output of the machine then loops back to the contributing modules, and a new aggregation is performed every 50 msec.

THE BEST AI IS COLLECTIVE

While it took longer than expected by the pioneers of the field, artificial intelligence has achieved iconic milestones, not least of which was Deep Blue’s victory over the world’s chess champion Gary Kasparov in 1997 (Hsu, 2002), and Watson’s dominance on the Jeopardy game show (Thompson, 2010). Interestingly, while artificial intelligence has in principle no need to imitate natural intelligence, it turns out that these digital high achievers are also Surowiecki machines, just like the brain is.

The core of Deep Blue’s artificial intelligence is a so-called “evaluation function”, the process by which it computes the quality of a chess position. In the regulation 3 minutes it has to make a move, Deep Blue uses 256 processors working in parallel to assess 60 billion Chess positions. Then it will make the move that takes it to the position that received the highest score. The evaluation function is itself a combination of four sub-processes that look separately at four ways to score a Chess position based on (1) the relative worth of the pieces on the board, (2) the ability of various pieces to go on the attack, (3) the safety of the King, and (4) the rate at which one is making progress against the opponent. The evaluation function aggregates the diverse, specialized, independent input it gets from its four sub-processes, making Deep Blue a Surowiecki machine.

Watson, another very public I.B.M. success, displays smarts that are even more massively collective. Facing the incredibly hard problem of trying to understand questions asked in natural language, it’s breakthrough strategy was to run *thousands* of different imperfect natural language processing algorithms in parallel, resulting in various Jeopardy answers depending on what each algorithm *independently* understood was being asked. Watson’s confidence in each possible answer then simply increased with the number of algorithms that returned it.

CONCLUSION & IMPLICATIONS

This selective overview of the state of the art in cognitive neuroscience and artificial intelligence suggests that intelligence, in its highest forms, tends to be organized according to the collective intelligence principles identified by Surowiecki to extract wisdom from crowds. Thus the wisdom-of-crowds recipe might well be a recipe for intelligence in general. It would seem that there is nothing fundamentally different between what might be called “regular” intelligence and the collective kind, apart from a preference to use the latter appellation when people (or insects), rather than algorithms or neurons, are collaborating successfully. As Simon (1969) suggests, there are good theoretical reasons why any complex system that is designed or evolved should necessarily be organized as a hierarchy of functionally and/or physically nearly-decomposable parts at every level. Such reasoning might help explain why intelligence has to be collective.

Viewing intelligence as collective is a mind shift with widespread implications for research and development in related fields, but also for individual decision makers and organizations. Without any pretense at exhaustivity or thoroughness, let's briefly discuss a few.

For cognitive science, one implication is that groups are legitimate objects of study at one extreme of an extended neurons-to-brains-to-crowds continuum of naturally intelligent systems. Woolley et al. (2010) recently showed that groups of people possess a collective IQ that is as tangible a construct as individual IQ, and furthermore that it is not strongly correlated with the average or maximum individual IQ of group members. Cognitive psychologists should not shy away from studying these exciting new types of thinking entities. Interestingly, collective IQ seems driven by the quality of the communication between group members, which is reminiscent of the connectionist maxim that "knowledge is in the connections," rather than in the nodes of a neural network (Rumelhart & McClelland, 1986, p. 132).

When designing an artificial intelligence, a general heuristic might be to try to build as much collective intelligence into it as possible, that is, to try to bring forth a diversity of independent contributions from many different algorithms to inform the ultimate decision-making module. Structuring an AI as a Surowiecki machine also enables it to make the best of parallel computing resources, as we have seen with the Deep Blue and Watson examples.

As individual decision makers, we are all hopelessly afflicted by the "confirmation bias": the natural tendency to seek and interpret evidence "in ways that are partial to existing beliefs, expectations, or a hypothesis in hand" (Nickerson, 1998, p. 175). An efficient remedy is to turn oneself into a Surowiecki machine by earnestly seeking and aggregating the equally-biased opinions of others. Of course, this strategy only works to the extent that the biases of others effectively run counter to our own, hence the need to cast a wide net for a diversity of opinions, no matter how "expert" you already are. In a famous study of expert geopoltical forecasters, Tetlock (2005) found that those willing to entertain several viewpoints were more likely to be right than those clinging to a dearly-held worldview. As Tetlock concludes: "what experts think matters far less than *how* they think" (p. 2).

The same lesson holds true for organizations and companies. When intelligence is collective, the obvious path to becoming a smarter company is to embrace a flatter organizational chart with enhanced collaboration over silos, a more widely shared power pie, and more employees tapped for various strategic decisions. As a case in point, a recent Booz & Co. study found that "most of the top companies ranked by their peers as "innovative" weren't among the top five spenders on research and development" (Korn, 2011). Instead of spending big, "the biggest innovators involve employees company-wide to help generate ideas".

REFERENCES

Anderson, J.R. (2007). *How Can the Human Mind Occur in the Physical Universe?* Oxford University Press.

Borst, J.P. & Anderson, J.R. (in press). Using the ACT-R Cognitive Architecture in combination with fMRI data. In B.U. Forstmann, & E.-J. Wagenmakers (Eds.), *An Introduction to Model-Based Cognitive Neuroscience.* Springer: New York.

Hsu, FH. (2002). *Behind Deep Blue: Building the Computer that Defeated the World Chess Champion*. Princeton University Press.

Korn, M. (2011). Top 'Innovators' Rank Low in R&D Spending, *Wall Street Journal*, October 24, 2011.

Minsky, M. (1986) *The Society of Mind*. Simon & Schuster.

Nickerson, R. (1998). Confirmation Bias: A Ubiquitous Phenomenon in Many Guises, *Review of General Psychology*, 2(2):175–220

Rumelhart, D. and McClelland, J. (1986). PDP Models and General Issues in Cognitive Science. In Rumelhart, D. and McClelland, J. (Eds) *Parallel Distributed Processing. Volume 1. Foundations*.

Simon, H. (1969). *The Sciences of the Artificial*. MIT Press.

Surowiecki, J. (2004). *The Wisdom of Crowds*. Doubleday.

Tetlock, P. (2005). *Expert Polti:cal Judgement: How Good Is It? How Can We Know? Princeton University Press.*

Thompson, C. (2010). What Is I.B.M. Watson? *The New York Times*, June 16, 2010

Woolley, A., Chabris, C. Pentland, A., Hashmi, N., and Malone, T. (2010). Evidence for a Collective Intelligence Factor in the Performance of Human Groups, *Science*, 330(6004), 686-688.