Intelligence Is Collective

Emile Servan-Schreiber
Lumenogic, LLC
475 Longview Road, South Orange, NJ 07079, USA
ejss@lumenogic.com

ABSTRACT: How does "collective" intelligence relate to the “individual” kind? A high-level review of the state of the art in cognitive science suggests that both brains and the best artificial intelligence programs can be described as “wise crowds” (in the sense of Surowiecki, 2004). Conversely, the most celebrated instances of collective intelligence bear deep similarities to the human mind’s organization and mechanisms. A tentative conclusion is that intelligence is collective to its core and that Surowiecki's recipe for crowd wisdom provides a framework to unify human, artificial and collective intelligence. Some implications for research, business and individual decision makers are briefly explored.

KEYWORDS: ACT-R, artificial intelligence, brain, collective intelligence, cognitive science, cognitive psychology, crowd wisdom, intelligence, prediction markets, wisdom of crowds

---

Collective intelligence had essentially been the province of insect colonies, flocks of birds, and new age philosophers until the World Wide Web started connecting human brains at an unprecedented scale. Prediction markets, Google and Wikipedia (in that chronological order) suddenly made the concept very real for millions of people on a daily basis, culminating in the wide cultural acceptance of what Surowiecki (2004) artfully coined “the wisdom of crowds”.

To explore this exciting new form of intelligence, a new interdisciplinary field of research draws from computer science, cognitive science, economics and sociology. In 2012, the first conference on collective intelligence was held at MIT. Historically, the situation is reminiscent of the year 1956 when a multidisciplinary bunch of daring scientists gathered at Dartmouth to lay the foundation of the field of artificial intelligence (“AI”). And just as that earlier endeavor shed a bright light on the nature of intelligence itself, we can expect the same from the pursuit of collective intelligence. So this is our topic: what does the study of collective intelligence teach us about intelligence itself?

<u>Machines Who Think</u>

The modern scientific study of intelligence began fifty-five years ago with the invention of the first artificial intelligence program by Newell, Simon and Shaw (Newell & Simon, 1956). This seminal achievement made very tangible the dual notions – still vague and unproven at the time – that machines could be made to “think” and that human thinking itself could be understood with sufficient precision to be simulated by computer programs. The year 1956 thus delivered both the founding of a new interdisciplinary field of artificial intelligence and a “cognitive revolution” in the venerable field of Psychology.

Half a century later, these two intertwined branches of cognitive science have produced richly detailed computational models of how the brain produces the mind, as well as powerful artificial intelligence programs that have cleared iconic milestones, such as besting the best human champions at Chess. However, what is of particular interest to us here is that in each case, as will be shown later, the intelligence of the system, whether natural or artificial, can be described as "collective" in the sense that it arises from the interactions of dumber parts and processes. In fact, all that is needed to make *one* brain or *one* AI program fit Malone et al.'s (2009) broad definition of collective intelligence as "groups of individuals acting collectively in ways that seem intelligent" is to replace "individuals" with "processes".

Of course, the idea that, as Minsky writes in *The Society of Mind* (1986), "you can build a mind from many little parts, each mindless by itself" is not new at all: it is arguably *the* original insight of artificial intelligence. In *Machines Who Think*, McCorduck (2004) recounts how in 1954, more than a year before he would co-invent the first AI program, Newell experienced a life-changing epiphany during a presentation of Selfridge's work on pattern recognition that demonstrated how "working in concert, a set of simple sub-processes that were easy to understand could lead to genuinely intelligent behavior." There is also a rather large consensus in this modern scientific age that the smartest general purpose thinking device we know of, the human brain, must somehow manufacture its higher intelligence from the interactions of billions of rather simple-minded neurons.

But we also know that not all systems involving many interacting parts can produce intelligent behavior. 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 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 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.

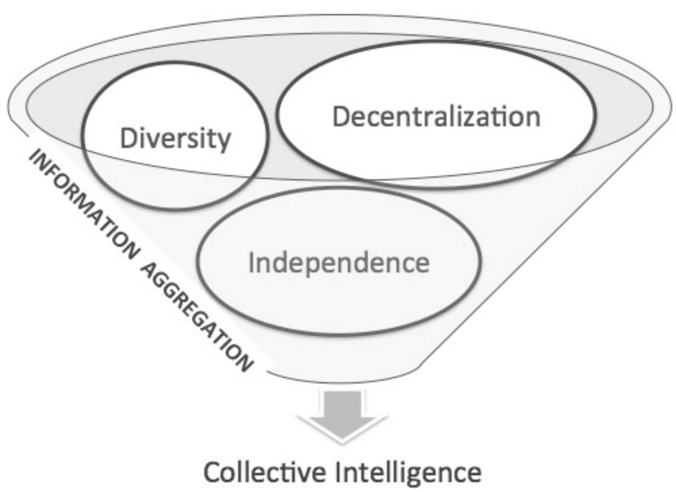

Figure 1: *Surowiecki's (2004) recipe for crowd wisdom: diversity of opinions, decentralized inputs, independent thought, and information aggregation. See text for explanation.*

Surowiecki Machines

For clarity and brevity in the argument that follows, let us name "Surowiecki machine" any information processing system made up of simpler devices that collaborate according to Surowiecki's four principles. This structure is illustrated in Figure 2:

- The devices contribute their input to the machine's aggregation mechanism, which then computes the machine's output.
- 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.
- Each device carries out its processing mostly independently of the others, and sends its output directly to the machine's aggregation mechanism with no interference from other devices.
- The output of the machine may or may not loop back as input to the component devices.

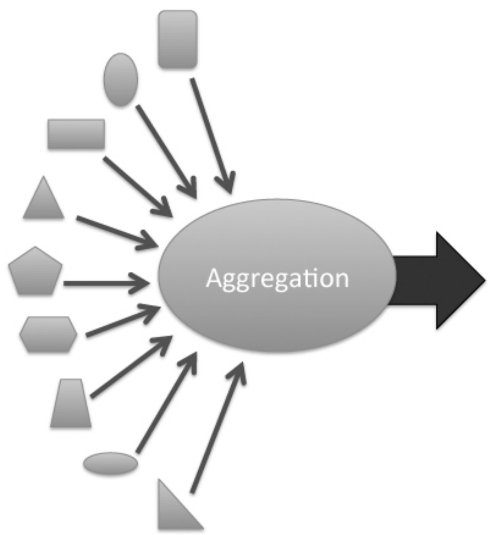

Figure 2: *Overall architecture of a Surowiecki machine. The aggregation module synthesizes its output from the diverse, specialized, independent inputs of a crowd of information processing devices.*

A Surowiecki machine is, by definition, a collective intelligence system. This paper will explore the idea that all intelligent systems, or at least those to which we attribute higher intelligence, can be described as Surowiecki machines. This, in turn, will lead to the conclusion that "regular" intelligence is a form of collective intelligence, or, to put it succinctly, that intelligence is collective.

The rest of this paper 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, artificial intelligence and collective intelligence. It will be shown that the brain, the mind and a variety of best-of-class AI programs can be described as Surowiecki machines. Conversely, today's better-known instances of collective intelligence will be shown to bear deep similarities to the mind's organisation and mechanisms.

THE BRAIN IS A WISE CROWD

The human brain is the most complex object we know in the Universe, and it is fair to say that a lot of its mysteries have yet to be unraveled. But one thing is already clear: the brain is highly modular, anatomically and functionally, at every level of its organization. Furthermore, at various levels, from humble neurons to consciousness itself, the brain seems to operate according to wise-crowds principles.

Each of the 100 billion neurons in a brain receives input from, on average, 1000 other neurons. Taken together, a neuron and those that feed into it make up a simple Surowiecki machine: the target neuron receives on average 1000 *independent* activation or inhibition input signals from other neurons near and far (insuring *diversity* and *decentralization*). Then, it *aggregates* this information by summing up the excitatory and inhibitory signals and, if its activation reaches a certain threshold , firing its own signal to another 1000 neurons.

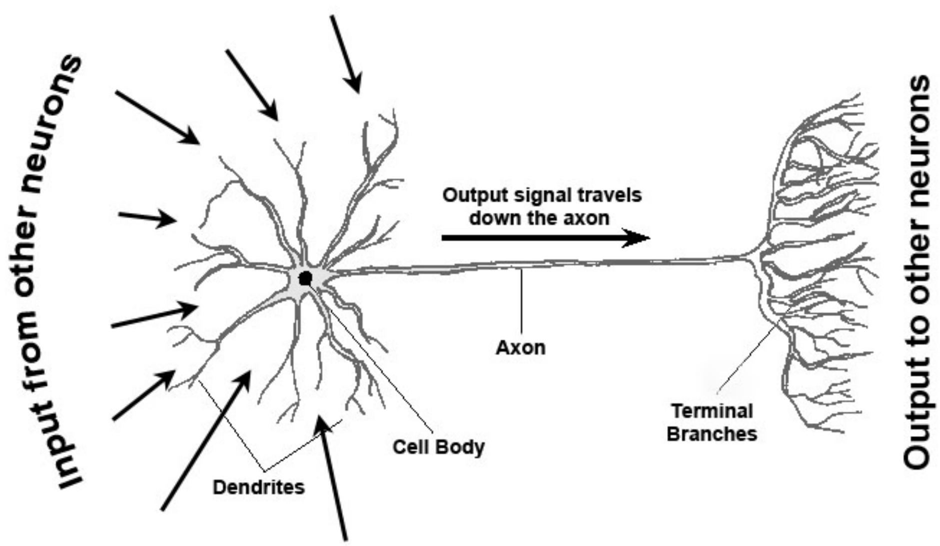

Figure 3: *The structure of a typical neuron closely resembles that of a Surowiecki machine (as in Figure 2). The neuron collects excitatory and inhibitory signals from a great many other neurons through its dendrites. The cell body aggregates the signals and fires its own output signal if the net excitation reaches a threshold. The output signal then travels down the axon to reach other neurons through the terminal branches.*

At a higher level, in the neocortex, functionally specialized neural networks are assembled into lobes with roughly-identified roles such as planning and control in the frontal lobes, memory in the temporal lobes, spatial processing in the parietal lobes and vision in the occipital lobes. The lobes themselves belong to one of two hemispheres that have been shown to process the world quite differently: the Left being more analytical and verbal, the Right being more perceptual and spatial. The extent of this hemispheric specialization can be quite striking, as illustrated by Figure 4 in the case of visual perception: our perceptual system aggregates and fuses the two complementary inputs it receives from the left and right parietal lobes, where one provides the details and the other the overall shape (Robertson & Lamb, 1991). Language provides another classic case of the two hemispheres making very different contributions. Whereas, for the vast majority of people, the ability to understand speech and produce grammatical sentences is located in various areas of the left hemisphere, their counterparts in the right hemisphere play a complementary role: they are concerned not with the content of speech (as encoded by grammar and vocabulary) but with its prosody, i.e., the rhythm, stress and intonation that convey the emotional subtext of an utterance (Kandel, 1997).

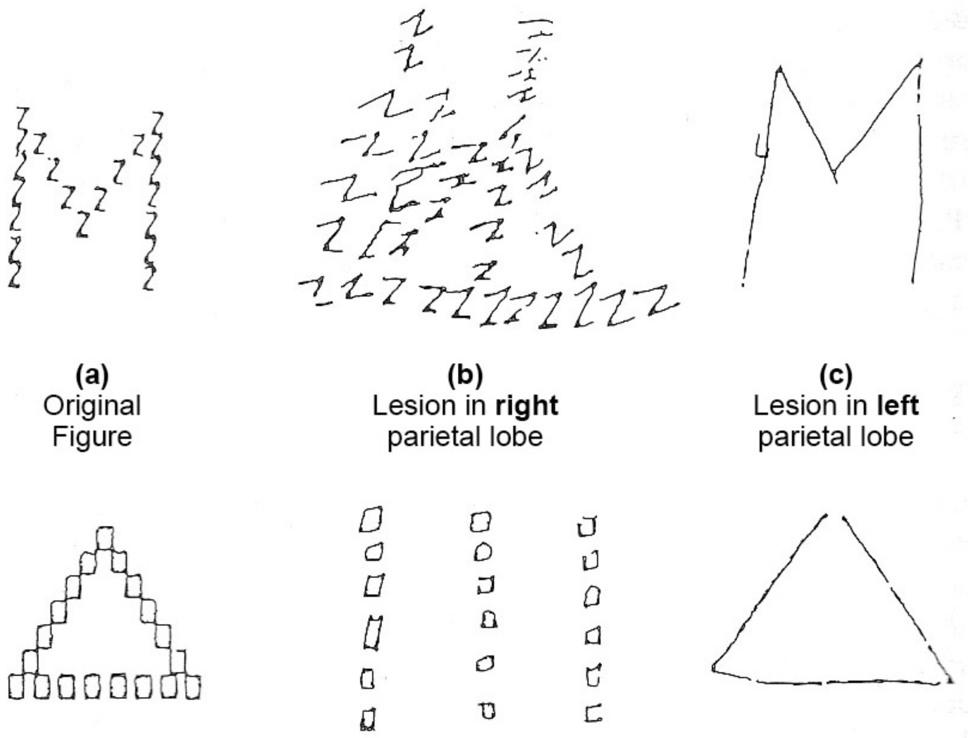

Figure 4: *An example of how the left and right hemispheres process the world differently, as revealed by patients with lesions to the right or left parietal lobes. (a) Figures that patients are asked to reproduce. (b) Patients with lesions in the right parietal lobe get the parts right, but are confused about the whole. (c) Patients with lesions to the left parietal lobe can reproduce the whole, but ignore the parts. Adapted from Robertson & Lamb (1991).*

Only a relatively small band of fibers, the corpus callosum, connects the two hemispheres, and spectacular experiments with so-called “split-brain” patients have shown that when the corpus callosum is severed, the two hemispheres can operate quite independently, each with a consciousness of its own (Gazzaniga, 1998). Furthermore, the language-enhanced left hemisphere seems to include a peculiar kind of information aggregation module, which Gazzaniga calls “the interpreter”, whose function is to retrospectively feed our conscious mind with rational interpretations of our behavior. The existence of the interpreter is revealed in a split-brain patient when her left hemisphere earnestly spins creative verbal explanations for behavior that was in fact driven by the right hemisphere. So, at the highest level of the brain’s organization, consciousness itself operates like a Surowiecki machine that aggregates the different views of the world contributed by the two hemispheres.

What about the organization of the mind? While far from complete, our understanding of human cognition is now such that a richly detailed computational model of its “architecture” offers a unified theory of memory, learning, problem solving, perception, and action: the ACT-R theory of mind can not only explain complex thinking at an abstract “symbolic” level, but also predict *where* and *when* it happens in the brain (Anderson, 2007). For instance, it can account for the detailed time course 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, correlates of which can be precisely localized in the brain (see Figure 5a). A non-exhaustive list includes modules for visual and aural perception, 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. To coordinate these various aspects of cognition, a central “Procedural” module serves as the gateway through which all the other modules communicate (see Figure 5b.) Anatomically, the Procedural module maps onto the *basal ganglia*, a set of subcortical structures deep in the brain that receive input from almost all areas of the cortex and project back to them either directly or indirectly through various frontal areas (Middleton & Strick, 1996; 2000). At each cycle of cognition, the Procedural module reacts to the pattern of input it receives from the other modules by matching to it a large set of condition-action rules, called “productions”. The production whose condition best matches the input pattern then “fires” and sends commands, requests and data to the various modules for another round of processing. Importantly, only one production may fire at each cycle, which takes approximately 50 ms and sets the tempo of cognition.

ACT-R’s unified theory of mind, and its mapping to brain, is a great scientific achievement,[1] but what is of most interest to us here is that a theory which represents the absolute state of the art in cognitive science’s effort to explain how the brain produces the mind is so congruent with viewing the mind as Surowiecki machine: while each module independently processes and contributes a very specific type of information, insuring diversity and specialization, the Procedural module is a central cognitive bottleneck which aggregates it all to decide in which direction our thinking should flow next. (In this case, the output of the machine - the “action” part of the winning procedure - loops back to the component devices, i.e., the input modules.)

[1] John R. Anderson, the founder and leader of the ACT-R research community (currently 205 researchers at 94 institutions in 15 countries) was awarded the 2011 Benjamin Franklin Medal in Computer and Cognitive Science "for the development of the first large-scale computational theory of the process by which humans perceive, learn and reason".

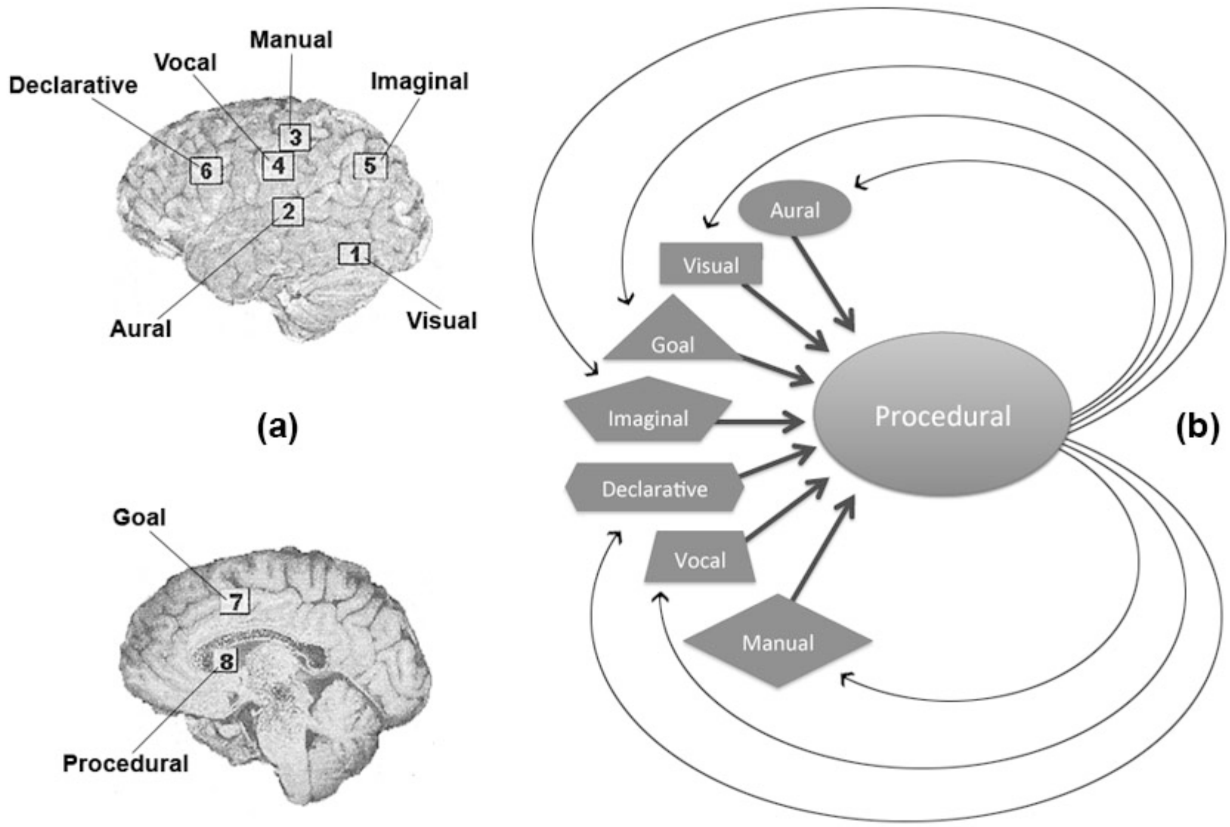

Figure 5: *How the brain does the mind according to ACT-R: (a) Localization in the brain of some neural correlates of the various ACT-R modules (homologues in the opposite hemisphere are not shown). Each region has been chosen because it systematically "lights up" when the corresponding module is active, but ACT-R certainly does not claim that the module's entire neural network must fit within that particular region. (b) Cognitive flow among the ACT-R modules. At each cycle of cognition (approx. 50 ms), the central Procedural module attempts to recognize patterns in the input it receives from the various modules in order to select a rule that issues commands, requests and data back to the various modules. Each module contributes its specialized, independent input for aggregation by the Procedural module. This diagram shows how similar ACT-R's cognitive flow is to a Surowiecki machine (see Figure 2).*

THE BEST ARTIFICIAL INTELLIGENCE 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 more recent dominance in the Jeopardy game show (Thompson, 2010). Interestingly, what these two digital high achievers have in common, apart from being IBM creatures, is that, at a high level, they can both be described as Surowiecki machines.

The core of Deep Blue's artificial intelligence is it's 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 independently at four complementary 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.

For its part, Jeopardy-champ Watson exploits collective wisdom at much larger scale. Facing the incredibly hard problem of understanding questions asked in natural language, a task at which no single algorithm consistently succeeds, Watson's breakthrough strategy was to run *thousands* of different natural language processing algorithms in parallel. Each algorithm would independently generate possible Jeopardy answers depending on what it understood was being asked. Watson then aggregated all these answers by selecting the one answer that was returned by a plurality of the algorithms. It is another clear case of a Surowiecki machine.

Artificial intelligence has in principle no need to imitate natural intelligence. Indeed, brains and computers excel at different tasks, such as patter recognition for the former, and deep search for the latter. However, what these two examples of best-of-class AI seem to suggest is that in order to achieve higher intelligence – by which I mean the capacity to solve problems that require hard thinking from human brains – the AI has to be organized as a Surowiecki machine, just like the human mind is.

THE BEST COLLECTIVE INTELLIGENCE IS BRAIN-LIKE

The field of collective intelligence is comparatively recent, yet millions of people are already interacting daily with its more powerful incarnations such as Google's search engine, Wikipedia, or prediction markets. Interestingly, each of these incarnations bears deep similarities to the human mind's organization and mechanisms.

To a cognitive psychologist, Google's PageRank algorithm for retrieving the most relevant web pages in response to a search query seems very familiar (Brin & Page, 1998). Indeed, as others have observed (e.g., Akim et al., 2011), PageRank implements a conservative form of the "spreading activation" mechanism that has been used in cognitive psychology since the 1970's to account for the behavior of human associative memory (e.g., Collins & Loftus, 1975; Anderson, 1983). Essentially, PageRank assigns higher potential relevance to pages that have many links to them, or links from other highly relevant pages, just like the human brain retrieves more easily memories that are associated with many others, or associated with other strong memories.

Wikipedia is the other poster child of collective intelligence on the web, and it too has an intriguing relationship to human cognition. The clustering of object names ("table", "chair", "kitchen") in Wikipedia pages so closely mirrors the clustering of those same concepts in human brains that researchers were recently able to use it to literally read the minds of their subjects. Pereira et al. (2011) recorded fMRI images of brains reading common object names while looking at simple line drawings. They used half the data to train a pattern detection algorithm, and kept the other for testing. The algorithm was trained to map brain activation patterns evoked by particular objects onto concept clusters extracted from relevant Wikipedia pages. It was then able to "look" at the fMRI images in the test set and correctly guess the category of the corresponding objects.

Prediction markets predate both Google and Wikipedia and were perhaps the first incarnation of collective intelligence on the Web. Their ability to consolidate the guesses of many individuals into accurate predictions in all sorts of subject domains – sports, politics, movies, business, and more – has been well documented over more than 20 years (see Servan-Schreiber, in press, for a full review). In a prediction market, participants trade stocks that will be worth for instance $1 if a particular event happens, or nothing if it doesn't. Those who make correct guesses will make money at the expense of those who don't.

To understand how this relates to brain mechanics, consider that a trader who guesses right and makes a profit will have more money to invest in the next round, while a trader who guesses wrong will have less money to invest in the next round. In this way, knowledge is rewarded with ever more influence on the market, while ignorance and hubris are progressively weeded out. This, of course, is reminiscent of how brains learn by repeatedly strengthening the neuronal connections that lead to desired or successful behaviors while weakening those that lead to undesirable behaviors. Play-money prediction markets are particularly good at improving their performance in this way because, in contrast to real-money markets, they maintain a strong correlation between individual wealth and individual accuracy. Servan-Schreiber et al. (2004) proposed this as an explanation for why play-money markets can demonstrate as much accuracy as real-money markets despite the lack of hard currency in the system.

CONCLUSION & IMPLICATIONS

This overview of the state of the art in cognitive neuroscience and artificial intelligence suggests that intelligence, in its highest forms, is organized according to the collective intelligence principles identified by Surowiecki in *The Wisdom of Crowds*. Conversely, the best examples of collective intelligence appear related in intriguing ways to the workings of the brain itself.

It seems that Surowiecki's recipe for crowd wisdom provides an elegant framework for unifying all forms of intelligence: human, artificial and collective. *Intelligence is collective to its core.* 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 (i.e., independent, specialized) parts at every level. This line of 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, especially companies. Without any pretense at exhaustivity or thoroughness, let's briefly discuss a few.

Research Implications

For cognitive science, one implication is that groups of brains 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 small groups of people possess a collective intelligence that is as tangible a construct as individual IQ, and furthermore that it is only weakly correlated with the average or maximum individual IQ of group members. Cognitive psychologists should not shy away from studying these exciting new types of intelligent entities. Interestingly, collective IQ seems driven by the ability of groups members to listen earnestly to each other, 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).

For artificial intelligence engineers, a general heuristic might be to try to build in as much collective intelligence 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.

For the field of collective intelligence, one implication is that what we know of the sophisticated architectures of brains and minds might be good sources of inspiration for designing new powerful forms of collective intelligence. Learning and development are particularly interesting areas to explore for this purpose, for, as is the case with brains, it will be surely be easier to achieve high-performance collective intelligence through learning than through static design. This is already evident with prediction markets, as discussed earlier, but there are also indications that properly designed feedback mechanisms can increase the efficiency of crowd-sourced idea markets over time (Huang, & al., 2012).

Implications for Decision Makers

As reasoning homo sapiens, we are all hopelessly afflicted by what psychologists call 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 to this major flaw in one's individual decision-making apparatus 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. Seeking out like-minded thinkers won’t help in this regard. Casting a wide net for a diversity of opinions is the surest way to harvest counter-biases.

It is natural for novices to seek advice, but even experts can improve their performance by turning to collective intelligence. For instance, in a recent study of experienced attorneys trying to predict civil jury verdicts, Jacobson et al. (2011) found that one's estimate could be improved by 25% if averaged with a single other attorney's estimate, or it could be improved by almost 50% if averaged with the estimates of three fellow attorneys. In a famous study of expert political forecasters, Tetlock (2005) found that no matter where they fell on the political spectrum, 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 takeaway is that turning yourself into a hub for collective intelligence is a particularly efficient way to become a smarter decision maker.

Implications for Organizations and Companies

The same lesson holds true for organizations and companies. With business fitness in the information age increasingly depending on smarts rather than brute capital, the past decade has seen more and more leading companies seizing on collective intelligence in attempts to raise their ability to forecast and innovate to the next level (e.g., Pethokoukis, 2004; Lohr, 2008; Stier, 2011; Lesser et al., 2012).

Traditional companies are poorly organized to behave as Surowiecki machines and capitalize on their collective intelligence potential: diversity and decentralization are thwarted by geographic separation or by deep silos maintained by incentives. Independence of thought is quashed by politics. Culture is often the only aggregation mechanism. Addressing these failings can have dramatic effect on company IQ, and leading companies are often those already doing this.

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". Similarly, a number of companies have reported improvements in their forecasting accuracy when their regular forecasting process is combined with a prediction market (e.g., Llera, 2006; Hopman, 2007).

Even at a smaller scale - without creating prediction or idea markets that tap hundreds or thousands of employees - companies may reap the profits of collective intelligence by heeding Woolley et al's (2010) and Jacobson et al's (2011) results on group IQ discussed earlier. If intelligence is collective, then teams of 2 to 5 people, rather than individuals, should be the basic decision-making units in the company. At all levels of the organization, from the cubicle to the C-suite, teams, rather than individuals, should be rewarded, promoted, and held accountable. HR departments should focus on assembling smart teams whose members are recruited as much for their ability to communicate, listen, and collaborate, as for their high IQ or individual achievements.

ABOUT THE AUTHOR

Emile Servan-Schreiber is a Managing Director of Lumenogic, LLC, where he consults with leading companies worldwide to design and implement prediction markets and other collective intelligence applications aimed at improving forecasting and accelerating innovation. He co-founded NewsFutures.com (2000-2010), a pioneering prediction market that received worldwide coverage and is repeatedly referenced in James Surowiecki's *Wisdom of Crowds*. Trained in cognitive psychology at Carnegie Mellon University (Ph.D., 1991), he is the author of several research publications about human learning, prediction markets and collective intelligence, and co-author of two award-winning multimedia CD-Roms featuring the collective wisdom of dozens of world-class scientists: *The Challenge of the Universe* (1996, Oxford University Press) and *Secrets of the Mind* (Montparnasse, 1998; Focus Multimedia, 2007). He can be reached at **ejss@lumenogic.com**.

REFERENCES

Akim, N. Md., Dix, A., Katifori, A., Lepouras, G., Shabir, N., Vassilakis, C. (2011). “Spreading Activation for Web Scale Reasoning: Promise and Problems,” In *Proceedings of WebSci'11*, Koblenz 15-17 June 2011.

Anderson, J.R. (1983), “A Spreading Activation Theory of Memory,” *Journal of Verbal Learning and Verbal Behavior*, **22**, 261–295.

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

Brin, S. and Page, L. (1998), “The Anatomy of a Large-Scale Hypertextual Web Search Engine,” *In: Seventh International World-Wide Web Conference (WWW 1998)*, April 14-18, 1998, Brisbane, Australia.

Collins, Allan M.; Loftus, Elizabeth F. (1975) "A Spreading-Activation Theory of Semantic Processing", *Psychological Review*, **82**(6) 407-428

Gazzaniga, M. (1998), *The Mind’s Past*. University of California Press.

Hopman, J. (2007). "Using forecasting markets to manage demand risk," *Intel Technology Journal*, 11(2)

Huang, Y., Singh, P., and Srinivasan, K. (2012) “Crowdsourcing New Product Ideas under Consumer Learning”. Working Paper available at http://ssrn.com/abstract=1974211

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

Jacobson, J., Dobbs-Marsh, J., Liberman, V., and Minson, J. (2011) "Predicting Civil Jury Verdicts: How Attorneys Use (and Misuse) a Second Opinion," *Journal of Empirical Legal Studies*, **8**(S1), 99-119

Kandel, E. (1997), “Anatomy of the Brain,” in Servan-Schreiber, E. and Levy, B. (Eds) *Secrets of the Mind*. Focus Multimedia.

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

Lesser, E., Ransom, D., Shah, R., and Pulver, B. (2012) "Collective Intelligence: Capitalizing on the Crowd," *IBM Global Business Services Executive Report*.

Llera, J. (2006). "Sparc: Arcelor’s Prediction Market," *Presentation at the First European Prediction Market Summit,* Vienna, November 2006.

Lohr, S. (2008) "Betting to Improve the Odds," *New York Times*, April 9, 2008.

Malone, T.W., Laubacher, R., and Dellarocas, C.N. (2009), “Harnessing Crowds: Mapping the Genome of Collective Intelligence,” *MIT Sloan Research Paper No. 4732-09*.

McCorduck, P. (2004), *Machines Who Think: A Personal Inquiry into the History and Prospects of Artificial Intelligence*. A K Peters.

Middleton, F., & Strick, P. (1996). “The Temporal Lobe is a Target of Output From the Basal Ganglia,” *Proceedings of the National Academy of Sciences of the USA*, **93**, 8683-8687.

Middleton, F., & Strick, P. (2000). “Basal Ganglia and Cerebellar Loops: Motor and Cognitive Circuits,” *Brain Research Reviews*, **31**, 236-250.

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

Newell, A. & Simon, H.A. (1956), “The Logic Theory Machine: A Complex Information Processing System,” *Rand Corporation Report*.

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

Pereira, F., Detre, G., and Botvinick, M. (2011), “Generating Text from Functional Brain Images,” *Frontiers in Human Neuroscience*, **5**.

Pethokoukis, J. (2004) "All seeing, all knowing," *US News & World Report*, Aug 22, 2004.

Robertson, L.C., and Lamb, M.R. (1991), “Neuro-Psychological Contributions to Theories of Part/Whole Organization,” *Cognitive Psychology*, **23**, 299-330.

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*. London: MIT Press, 110-46.

Servan-Schreiber, E. (in press), “Trading Uncertainty for Collective Wisdom,” in Landemore, H. and Elster, J. (Eds) *Collective Wisdom: Principles and Mechanisms*, Cambridge University Press.

Servan-Schreiber, E., Wolfers, J., Pennock, D., and Galebach, B. (2004), “Prediction Markets: Does money matter?” *Electronic Markets*, **14**(3).

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

Stier, K. (2011) "Lumenogic Mines Worker's Opinions for U.S. Air Force, Multinationals," *Bloomberg*, August 30, 2011.

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

Tetlock, P. (2005) *Expert Political 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.

LUMENOGIC

Solving Business Problems With Collective Intelligence

Using online prediction markets, innovation competitions and other crowdsourcing techniques, Lumenogic provides solutions that harness the collective intelligence of employees or customers to improve forecasting, innovation, customer insights and strategic planning.

http://www.lumenogic.com

Lumenogic, LLC
475 Longview Road
South Orange, NJ 07079
U.S.A.

Tel: +1 203 702 6224