Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence (IJCAI-17)

Understanding and Measuring Collective Intelligence Across Different Cognitive Systems: An Information-Theoretic Approach

Nader Chmait
Faculty of Information Technology, Monash University, Clayton, Melbourne, Vic. 3800, Australia
nader.chmait@monash.edu, chmait.nader@gmail.com

Abstract

We develop the idea of collective intelligence by analysing a range of factors hindering the effectiveness of interactive cognitive agents. We give insights into how to explore the potential of collectives across different cognitive systems (human, animal and machine) and research areas. The endeavour is to bridge the different research disciplines in which collective intelligence might occur and apply the studies of intelligence in AI to other fields, thereby cross-fertilising diverse areas of study ranging from business and management to social sciences and fundamental biology.

1 Introduction and Motivations

Collective intelligence [Weschsler, 1971] occurs in a wide range of areas such as social sciences and economics (human societies, polities, and organisations harnessing the wisdom of the crowd), biology (animal herds and insect colonies) and computer science (artificial life and nature-inspired evolutionary computation). Famous applications include crowdsourcing, public policy, recommender systems, social computing, swarm intelligence and complex adaptive systems.

Despite the remarkable advancements in recent years, the vast majority of research on collective intelligence (including the study of collectives in AI) has investigated its emergence in isolation, that is, either within a limited range of disciplines, or at the level of one particular cognitive system. Thus, links are still missing to connect fundamental characteristics that are shared among these studies. The problem of the disparity and absurdity in defining and measuring individual intelligence between, and even within, disciplines extends to collectives, making it difficult to cross boundaries between human, animal and machine entities when it comes to comparing their (collective) intelligence. Another serious limitation inhibiting our understanding of collective intelligence is the lack of *quantitative* analysis of intelligence with regards to task difficulty, and the comparison of individual agent performance to multi-agent or group performance. As a result, it is unclear as to which dynamics and conclusions transcend beyond a particular cognitive system or discipline. Moreover, it is not clear how to quantify and compare the performances of these systems, especially state-of-the-art AI agents [Hernández-Orallo *et al.*, 2017]. Interesting questions such as “how does the effectiveness of a collective quantitatively compare to that of its isolated members?” and, more importantly, “are there some general rules or properties shaping the spread of intelligence across various cognitive systems and environments?” remain somewhat of a mystery.

2 Some Contributions

In this research, we adopt an information-theoretic approach to the measurement of intelligence, using notions from algorithmic information theory and (Solomonoff-)Kolmogorov complexity [Li and Vitányi, 2008]. We develop the idea of collective intelligence by giving insight into a range of factors hindering the effectiveness of interactive cognitive systems. For instance, the intelligence of a collective is known to depend on the communication and observation abilities of its individual members. However, it is not clear which factor has the greater influence. Thus, we have analysed in [Chmait *et al.*, 2015a] the impact of these two factors on the intelligence of agent-based collective. Using *general* intelligence tests [Hernández-Orallo and Dowe, 2010; Chmait *et al.*, 2015b; Chmait *et al.*, 2016b], we evaluated and compared the performance of collaborative agents across different communication and observation abilities of measurable entropies. Our results revealed circumstances under which the effectiveness of a system of agents of low observation or perception abilities could be significantly improved by introducing (higher entropies of) communication between the agents in the system. We further analysed the dependency between observation and communication abilities within a collective of interactive agents.

Earlier studies revealed that collectives can outperform individuals [Woolley *et al.*, 2010], and that their performance is controlled by one or more of their organisational or network structures [Mason and Watts, 2012], the information aggregation details among their individuals, and the diversity between their members [Hong and Page, 2004]. Crowd-computing and crowdsourcing [Bonabeau, 2009] methodologies are good examples of collectives that harness the wisdom of the crowd [Surowiecki, 2005]. After looking at these and other literature on collective intelligence, we filtered a set of important features that we believe to be intimately relevant to the performance of agent collectives [Chmait *et al.*, 2016a]. The identified features, illustrated in [Chmait *et al.*, 2016a, Figure 1], are not coupled to one particular cognitive system, problem or environment. They rather consist of general characteristics such as the number of members in a group, the communication or interaction protocol used, the complexity of the environment, as well as other factors that are often neglected, such as the speed of the agents and the interaction time of the collective as a whole. By conducting a series of experiments on artificial (reinforcement learning, local search and other types of) agents, we empirically demonstrated the (major) influence that each of these features individually has on the intelligence of the group, as well as the simultaneous influence of multiple such features combined. In other words, collective intelligence was shown to be a function of all the examined factors simultaneously, as well as some of them combined. For example, we presented results where one group of agents outperforms another under some values or settings of the studied factors yet failed to do so under others (e.g., after limited vs. extended interaction times between the agents, or under low vs. high problem uncertainties). We showed that a group relying on an expert or super-solver agent does not guarantee its optimal performance. We also measured the effect of introducing more agents into the group, and showed that it is highly influenced by the communication medium implemented between its members. We further demonstrated how the difficulty of the environment in which a group operates (which partly relates to its *uncertainty* and the *algorithmic information-theoretic complexity* of its underlying tasks) is often a major factor controlling the group's capacity for intelligence. Finally, we presented conclusions showing how the network and organisational structures of the members of a group (e.g., flat, hierarchical, subgroup, and other structures) influence its overall performance.

3 Conclusion and Future Work

This project presents several experimental outcomes that are fundamental to the understanding and prediction of the dynamics and characteristics behind intelligent cognitive systems. The contributions from this research lie at the heart of understanding and potentially devising successful solutions to a variety of complex multi-agent problems. In the following parts of this research, we aim to develop a new mathematical model for predicting the accuracy of agents over problems of different complexities [Chmait *et al.*, 2017]. We also intend to present a new perspective for comparing intelligence between non-uniform types of agents operating in vastly different environments and contexts. Common grounds for evaluation are to be provided using a methodology for abstracting tasks and modelling environments as network graphs showing, at the same time, how to measure their complexities. This will further be used in the endeavour to connect studies of intelligence to other spheres, notably the areas of business decision-making and management.

Overall, this research project provides general guidelines that give insight into how to explore the potential of collectives across different cognitive systems and research disciplines. It also provides initial forays towards bridging different research areas in which collective intelligence might occur, and consequently cross-fertilising diverse fields of study ranging from businesses and large organisations to social sciences and fundamental biology.

References

[Bonabeau, 2009] Eric Bonabeau. Decisions 2.0: The power of collective intelligence. *MIT Sloan management review*, 50(2):45–52, 2009.

[Chmait *et al.*, 2015a] Nader Chmait, David L. Dowe, David G. Green, and Yuan-Fang Li. Observation, communication and intelligence in agent-based systems. In *Proceedings of the 8th International Conference Artificial General Intelligence*, volume 9205 of *Lecture Notes in Artificial Intelligence (LNAI)*, pages 50–59, Berlin, Germany, Jul 2015. Springer.

[Chmait *et al.*, 2015b] Nader Chmait, David L. Dowe, David G. Green, Yuan-Fang Li, and Javier Insa-Cabrera. Measuring universal intelligence in agent-based systems using the anytime intelligence test. Technical Report 2015/279, FIT, Clayton, Monash University, 2015.

[Chmait *et al.*, 2016a] Nader Chmait, David L. Dowe, Yuan-Fang Li, David G. Green, and Javier Insa-Cabrera. Factors of collective intelligence: How smart are agent collectives? In G. A. Kaminka, M. Fox, P. Bouquet, E. Hüllermeier, V. Dignum, F. Dignum, and F. van Harmelen, editors, *Proc. of 22nd European Conference on Artificial Intelligence ECAI*, volume 285 of *Frontiers in Artificial Intelligence and Applications*, pages 542–550, The Hague, The Netherlands, 2016. IOS Press.

[Chmait *et al.*, 2016b] Nader Chmait, Yuan-Fang Li, David L. Dowe, and David G. Green. A dynamic intelligence test framework for evaluating AI agents. In *Proceedings of 1st International Workshop on Evaluating General-Purpose AI (EGPAI 2016), European Conference on Artificial Intelligence (ECAI 2016)*, pages 1–8, The Hague, The Netherlands, 2016.

[Chmait *et al.*, 2017] Nader Chmait, David L. Dowe, Yuan-Fang Li, and David G. Green. An information-theoretic predictive model for the accuracy of AI agents adapted from psychometrics. In *Proc. of the 10th International Conference on Artificial General Intelligence*, Lecture Notes in Artificial Intelligence (LNAI), Melbourne, Australia, 2017. Springer. To appear.

[Hernández-Orallo and Dowe, 2010] José Hernández-Orallo and David L. Dowe. Measuring universal intelligence: Towards an anytime intelligence test. *Artificial Intelligence*, 174(18):1508–1539, December 2010.

[Hernández-Orallo *et al.*, 2017] Jose Hernández-Orallo, Marco Baroni, Jordi Bieger, Nader Chmait, David L. Dowe, Katja Hofmann, Fernando Martínez-Plumed, Claes Strannegård, and Kristinn R. Thórisson. A New AI Evaluation Cosmos: Ready to Play the Game? *To appear in the AI Magazine, Association for the Advancement of Artificial Intelligence*, 2017.

[Hong and Page, 2004] Lu Hong and Scott E. Page. Groups of diverse problem solvers can outperform groups of high-ability problem solvers. *PNAS*, 101(46):16385–16389, 2004.

[Li and Vitányi, 2008] Ming Li and Paul Vitányi. *An introduction to Kolmogorov complexity and its applications (3rd ed.)*. Springer-Verlag New York, Inc., 2008.

[Mason and Watts, 2012] Winter Mason and Duncan J. Watts. Collaborative learning in networks. *PNAS*, 109(3):764–769, 2012.

[Surowiecki, 2005] James Surowiecki. *The Wisdom of Crowds*. Anchor, 2005.

[Weschsler, 1971] David Weschsler. Concept of collective intelligence. *American Psychologist*, 26(10):904, 1971.

[Woolley *et al.*, 2010] Anita Williams Woolley, Christopher F. Chabris, Alex Pentland, Nada Hashmi, and Thomas W. Malone. Evidence for a collective intelligence factor in the performance of human groups. *Science*, 330(6004):686–688, 2010.