Local experts: Coordination and the albatross of domain-general HLC*

Eugene Philalithis[1,2]

[1] School of Mathematical and Computer Sciences, Heriot-Watt University
[2] School of Informatics, University of Edinburgh

Keywords: coordination · cognition · expert systems · distributed reasoning

1 The Myth of the AGI

Domain-general reasoning, inference and knowledge coordination mark a totemic goal for human-like computing (HLC). An artificial general intelligence (AGI) that can solve problems across an arbitrary range of domains, transfer expertise developed for one domain to another, and coordinate its solutions with both human and artificial collaborators, is ostensibly the most human-like of all. Far in the future a singular, ideally human-like agent teaches a course, designs a new product, diagnoses a patient, or conducts an orchestra without being specifically designed to do so. Yet in the present, artificial agents are still far from this ideal.

The current expertise of even the best AI systems, capable of beating humans in games or trivia, or matching humans in driving skill, is so narrow that it has recently been argued their brand of intelligence is fundamentally less than human [1]. Presently, I consider whether such thinking is justified: (a) whether artificial agents with narrow expertise are as far removed from domain-general intelligence as suggested, and equally (b) whether human cognition and human agents themselves typically meet the criteria demanded of an AGI in many cases. Across both these lines of enquiry, my intent is not to claim that an AGI is close at hand, or that human intelligence is not in principle domain-general. Instead, my aim is to highlight that a singular AGI-equivalent intelligence is often a poor fit for the human case, compared to a self-coordinated network of local experts. I then discuss the lessons future modelling work can learn from this comparison.

1.1 Less Than Human

As summarised above, current cutting-edge AI systems are, at best, local experts. The term 'local' as I use it here may refer to a particular domain of knowledge, e.g. possible chess moves, or to a particular application, e.g. image classification. Whatever their ultimate competence in that domain, something setting them apart from human agents is their inability to transfer that prowess into another task or area of knowledge entirely - something humans are eminently capable of.

* The author acknowledges the support of EPSRC Platform Grant EP/N014758/1.

The above received wisdom is not something I intend to challenge here. What I instead challenge is the further inference that, because human intelligence is *in principle* capable of such broad skill and effortless knowledge transfer, human achievement is mainly (if not exclusively) the result of such reasoning *in practice*. In defence of this deflationary line of thought, I review three different examples of human problem-solving: visual perception, joint language use, and pre-industrial goods manufacturing. Each of these tasks calls for domain-general expertise when taken as a whole, but for each task this is routinely approximated through the efficient coordination of local expertise across (or within) human problem-solvers.

The first of these examples, visual perception, illustrates a cognitive system with immense local expertise in classifying images geometrically [2]. At the same time however, many of the most critical, and most routine classification decisions made by this system integrate information from other cognitive processes, such as semantic memory [3]. Where classifying an image based on geometry may be inefficient or insufficient, the conglomerated expertise of two otherwise separable systems successfully, efficiently identifies e.g. a dog shape based on a noisy image. There is no one system which achieves this, only a coalition of processes (cf. [1]).

Where perception is a more abstract example highlighting the separate parts comprising a human intelligence, the second example concerns the competence of each human agent taken as a whole. Humans routinely and efficiently coordinate their use of language [4], and one consequence of this is the joint use of the same names or other labels to describe the same objects, as part of a common language. At the same time, it is not the case that every user of a common label is fully aware of its meaning, except through vague connotations (e.g. radium is 'a green rock to avoid'). Rather than each language user holding the full information, and general expertise, required to define all the terms in their language, the burden of defining each label is distributed to local experts (e.g. radiologists) and diffused into a community of language users, in order to shape their later behaviour [5]. Collectively the language is full of information, but no one user can access it all.

A final example where the individual expertise contributed by a human agent (or cognitive system) is less than the 'ideal human' standard aspired to by an AGI is far removed from cognition. Combining the insights identified in perception and joint language use, this last case in point is the pre-industrial manufacturing of commercial goods [6]. In stark contrast to the uniform and pre-planned nature of manufacturing in industrial times, craftsmen from earlier periods faced much uncertainty. Typically, an item such as a book required a broad range of expertise to produce, such as paper-making, printing, and textile production for binding them; but the relevant experts worked without explicit coordination at each step. Instead, every expert assumed and was ready to deal with large variation in the available input for their part of the chain. Despite having no expertise beyond their narrow role, their task included locating the best available materials, as a result of a previous expert's contribution, to then assemble a haphazard product. Neither all the expertise required (as with language) nor all the input necessary (as with perception) was available to a single agent in this sequence. Nonetheless, the craftsmen succeeded at their tasks to assemble varied, sophisticated products.

2 The Promise of Coordination

I have considered three examples of systems, one within and two between human agents, where an impressive result demonstrating expertise along several domains of industry or knowledge is not the product of any one contributor. Despite the hypothetical potential of any one contributor to master all of the necessary skills and acquire all the resources required, humans instead routinely distribute their expertise; even where this chain of interlinked expertise is unsteady and variable.

As a result, it is possible to argue that the coordination of partial, domain-specific solutions may represent an equally human-like alternative and far more feasible ideal than a singular AGI, for routinely domain-general problem solving. This in turn suggests that modelling coordination is a pressing goal for HLC to pursue, especially where such coordination is mediated by inference. Whether gauging which factor is relevant for classifying an image, or which detail is most important to remember from the expert definition of radium, or how to sustain a flexible manufacturing chain based on the work of others, the above examples may contest (as I have argued) the AGI-like understanding of human intelligence; yet they nonetheless demonstrate a sophisticated, flexible brand of coordination.

This in turn motivates two distinct conclusions. Firstly, that an emphasis on developing individual artificial agents with domain-general expertise may run contrary to the practical goal of developing human-like computing capability. In so far as humans are the basis for comparison, the most efficient problem-solving often uses less than the idealised human potential, relying more on flexible links between local experts much like current only-chess-playing, only-car-driving AI.

Secondly, that the requirements for such flexible coordination do nonetheless outpace existing multi-agent systems solutions comparable to planned industrial manufacturing, e.g. swarm robotics [7], and offer another target for future work. A network of agents (or subsystems) with local expertise collectively approaches domain-generality through flexible coordination - so that the goal of engineering coordination, mediated by inference, might take precedence over expertise itself.

References

1. Bundy, A. (2017). Smart machines are not a threat to humanity. *Communications of the ACM*, *60*(2), 40–42.
2. Marr, D. C. (1982) *Vision: A computational investigation into the human representation and processing of visual information.* Cambridge, MA: The MIT Press.
3. Vetter, P. & Newen, A. (2014). Varieties of cognitive penetration in visual perception. *Consciousness and Cognition, 27*, 62–75.
4. Pickering, M. J. & Garrod, S. (2006). Alignment as the basis for successful communication. *Research on Language and Computation, 4*(2-3), 203–228.
5. Burge, T. (1979). Individualism and the mental. *Midwest Studies in Philosophy, 4*(1), 73–122.
6. Reynard, P. C. (2000). Manufacturing quality in the pre-industrial age: Finding value in diversity. *The Economic History Review, New Series, 53*(3), 493–596.
7. Gautan, A. & Mohan, S. (2012). A review of research in multi-robot systems. In: *Proceedings of the 2012 IEEE 7th International Conference on Industrial and Information Systems (ICIIS)*, pp. 1–5. Chennai, India.