

# Trends in Cognitive Sciences

CellPress  
OPEN ACCESS

25th Anniversary Series: Looking Forward

## Forum

### The computational society

Nick Chater  1,2,\*

**How do individual human minds create languages, legal systems, scientific theories, and technologies? From a cognitive science viewpoint, such collective phenomena may be considered a type of distributed computation in which human minds together solve computational problems beyond any individual. This viewpoint may also shift our perspective on individual minds.**

The central assumption of cognitive science is that the brain is an information-processing system: cognition is computation. Yet, much of human thought and activity is collective. Mathematics, science, technology, legal and ethical systems, and political and economic institutions are inordinately complex cumulative products of interactions between many interconnected minds. Furthermore, the information processing involved in developing a field of mathematics, the patterns of natural language, or systems of norms and laws is distributed across countless individuals and typically far exceeds the capabilities of any one person. How is this possible?

To understand such collective achievements, we need to study how distributed computation arises, how it operates and is organized, and when it fails. Cognitive scientists have long built bridges with the neurosciences. At least as important, though, is extending a computational perspective to the social sciences. In the next 25 years, I hope that cognitive scientists will increasingly explore the implications

of the idea that society is distributed computation.

To make the computational society more than a metaphor, we need conceptual tools and methods to understand social phenomena in information-processing terms. Fortunately, several different, yet complementary, approaches have emerged in recent years. Here I highlight four promising lines of work: (i) social interaction as computation, (ii) the computational Leviathan, (iii) collective self-correction and rationality, and (iv) computation through spontaneous order.

#### Social interaction as computation

The minimal but crucial form of distributed computation is distributed through the social interaction of two or more people. Cognitive science has increasingly viewed social interaction, including communication, as a joint computational process in which people generate common beliefs, plans, and interpretations to coordinate their thoughts and behavior. These joint mental states are open-ended and highly flexible, whether we are engaged in conversation, improvised partner dancing, cooking together, or jointly creating a scientific theory. Taking part in such joint understanding, plans, and actions also presents formidable computational challenges for each individual agent, especially when computing what is or is not common ground between agents, and coordinating thought and behavior by reasoning over that common ground. One promising approach is to see agents as jointly negotiating momentary agreements or pacts [1], which govern joint behavior over short periods but which may then become entrenched over time. But how are such provisional and flexible agreements reached, often with minimal or no communication? One suggestion is that each party is engaged in a process of simulated bargaining, working out what would be agreed about word meanings, objectives, and plans, were negotiation possible – and generated actions and

communicative signals which help coordinate the ongoing negotiation (for this and related approaches, see [2–5]). More broadly, the idea that social life is governed by layers of such implicit bargains, building upon each other over long periods of cultural evolution, provides a possible mechanism for the emergence of social norms, conventions, and ultimately codified rules and legal systems.

#### The computational Leviathan

Extending the scope of distributed computation further, we can consider organized human behavior of all kinds, from momentary interactions to collaborative teams, firms, and even governments as distributed computation, albeit composed of extremely sophisticated processing elements: human minds [6]. Indeed, there is a long tradition of viewing organizations as information-processing systems [7]. But developments in parallel and distributed computing over the past few decades may provide useful insights into questions as diverse as the division of computational labor between human ‘processors’, efficient network structures and protocols for interaction between processors, and questions of stability and security. In short, if organizations are distributed information-processing systems, then our understanding of networked computation should provide the starting point for a cognitive science of organizational behavior.

To illustrate, consider one of the fundamental puzzles of cognitive science: the rough stability of prosocial collaborative behavior (including truth-telling), even though self-interested individuals might benefit from strategically deviating from such behavior. One explanation for the maintenance of prosocial behavior is reputation – those with poor reputations for prosocial behavior will not be chosen as partners in future interactions [8]. But for signals of reputation to be reliable presupposes that communicative signals are mostly trustworthy,

raising the possibility of regress. Very closely related problems are addressed in the long-standing literature in computer science concerning when it is possible to reliably guard against potentially malicious agents or other sources of error [9]. In fact, this literature often uses analogies with problems of social organization to develop and explain both the problem and its solution. Insights from distributed computing may help cognitive scientists develop Thomas Hobbes's vision of society as a Leviathan composed of complex, self-enforcing, and largely mutually beneficial interactions between individuals with primarily local concerns.

### **Collective self-correction and rationality**

We continually correct each other. Mathematicians find holes in each other's proofs; programmers find bugs in their own and each other's programs; science proceeds by a continual cycle of conjecture, refutation, and refinement; lawyers find and close legal loopholes; traders exploit and eliminate arbitrage opportunities in markets, and so forth. Each individual mind is liable to myopia, bias, and error, but when mechanisms of self-correction operate well, our collective efforts can create knowledge, rules, and technologies which can be remarkably sophisticated and consistent.

Indeed, the results of collective cognition are often far more rational than the individual minds of which the collective is composed. This suggests that rational theories, such as logic, probability theory, decision theory, or game theory, need not be applied primarily to individual minds at all. Instead, perhaps formal rational theories provide conditions under which our collective thoughts are (or are not) in equilibrium, highlighting opportunities for correction that can be exploited by individual agents. Thus, rational theories may shape the endpoint of self-correcting distributed cognition (e.g., just as mathematical logic helps

provide foundations for a collectively constructed mathematics), rather than guiding individual minds from within [10]. Equally important is understanding when self-correction fails and why collective cognition can generate irrational outcomes, from cults to conspiracy theories, superstitions, and price bubbles.

### **Computation via spontaneous order**

Which distributed computations does an organization perform? Does the goal of collective computation need to be set by the objectives of a leader? A long tradition in the social sciences suggests not. The Scottish Enlightenment philosopher Adam Ferguson famously saw social complexity as 'the result of human action, but not the execution of any human design' [11]. This view aligns with the idea that social complexity arises through spontaneous order, as developed by Michael Polanyi and Friedrich Hayek. An important challenge for future research is to understand such processes from a computational point of view.

The construction of a language, a banking system, or a legal and political order seems an information-processing problem par excellence. Each is surely a problem that exceeds the capacity of a single human mind, requiring trial-and-error learning and continual adjustment and innovation. Moreover, the space of possible innovations is vast, and the purposes to which they are put seem to coevolve with the mechanisms themselves, rather than solving a prespecified computational problem. For example, the invention of money might initially allow the exchange of goods, but it has evolved and developed to underpin the astonishing complexity and multiple purposes of the modern financial system. How can we understand such distributed, spontaneous processes in computational terms? This is perhaps the most challenging of our four examples.

The cultural evolution of language may provide a promising case study. Improvised communication and grammaticalization, over many generations, appear to explain the complex mixture of order and disorder that arises in human language, without postulating an innate universal grammar [12]. Moreover, parallels with self-organization in the natural world are intriguing. There have been tentative suggestions that self-organizing structures in nature arise from optimizing various quantities related to entropy. This raises the possibility that aspects of social structure can be seen in the same terms. The attempt to connect entropy with economic and social order has a long and somewhat unhappy history, littered with loose parallels that collapse under scrutiny. The mathematician Robert Mackay and I have recently been studying the very simplest case: mapping simple exchange economies into classical thermodynamics. Perhaps the entropy-like ideas may usefully be extended to other social phenomena.

### **Concluding remarks**

Cognitive science may stand on the brink of a new revolution, seeing social, organizational, and cultural processes as distributed computation. If so, we will need to look afresh at the computational role of individual minds. For example, rather than seeing each developing child as a lone minilinguist or a scientist-in-the-crib, we may, following Adam Ferguson, see humans as primarily learning to contribute to collective computations beyond the understanding of individual understanding.

### **Acknowledgments**

These ideas have been developed over conversations with many friends and colleagues, especially Morten Christiansen, George Loewenstein, Robert Mackay, Tigran Melkonyan, Jennifer Misyak, Mike Oaksford, Adam Sanborn, Robert Sugden, and Hossam Zeitoun. This work was supported by the ESRC Network for Integrated Behavioural Science (grant number ES/P008976/1).

### **Declaration of interests**

The author has no interests to declare.

<sup>1</sup>Behavioral Science Group, Warwick Business School, University of Warwick, Coventry, CV4 7AL, UK  
<sup>2</sup><https://www.wbs.ac.uk/about/person/nick-chater/>

\*Correspondence:

[Nick.Chater@wbs.ac.uk](mailto:Nick.Chater@wbs.ac.uk) (N. Chater).

<https://doi.org/10.1016/j.tics.2022.07.004>

© 2022 The Author(s). Published by Elsevier Ltd. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0>).

## References

- Brennan, S.E. and Clark, H.H. (1996) Conceptual pacts and lexical choice in conversation. *J. Exp. Psychol. Learn. Mem. Cogn.* 22, 1482–1493
- Chater, N. *et al.* (2022) The paradox of social interaction: shared intentionality, we-reasoning and virtual bargaining. *Psychol. Rev.* 129, 415–437
- Goodman, N.D. and Frank, M.C. (2016) Pragmatic language interpretation as probabilistic inference. *Trends Cogn. Sci.* 20, 818–829
- Misyak, J. and Chater, N. (2022) Instantaneous systems of communicative conventions through virtual bargaining. *Cognition* 225, 105097
- Wang, R. *et al.* (2021) Too many cooks: coordinating multi-agent collaboration through inverse planning. In *Human-like Machine Intelligence* (Muggleton, S. and Chater, N., eds), pp. 152–170, Clarendon Press
- Goldstone, R.L. and Theiner, G. (2017) The multiple, interacting levels of cognitive systems perspective on group cognition. *Philos. Psychol.* 30, 334–368
- March, J.G. and Simon, H.A. (1958) *Organizations*. Wiley
- Baumard, N. *et al.* (2013) A mutualistic approach to morality: the evolution of fairness by partner choice. *Behav. Brain Sci.* 36, 59–78
- Lampert, L. *et al.* (1982) The Byzantine generals problem. *ACM Trans. Program. Lang. Syst.* 4, 382–401
- Nau, R.F. and McCardle, K.F. (1991) Arbitrage, rationality, and equilibrium. *Theory Decis.* 31, 199–240
- Ferguson, A. (1767) In *An Essay on the History of Civil Society* (Oz-Salzberger, F., ed.), pp. 1995, Cambridge University Press
- Christiansen, M.H. and Chater, N. (2022) *The Language Game*. Bantam Books/Basic Books