The collective ochlotecture of large language models

Jon Dron

Athabasca University, Canada

Position paper for the First International Symposium on Educating for Collective Intelligence, University of Technology Sydney, 2024

Introducing "ochlotecture"

"Collective intelligence" (CI) is used across many disciplines to describe diverse phenomena. In my own work I have, for example, used the term to denote:

1. emergent stigmergic, flocking, and similar behaviours that allow crowds to be treated as entities in their own right (e.g. Dron, 2007);
2. the ways in which group processes, structures, human propensities, and individual actions can lead to more or less successful achievement of intentional learning goals (e.g. Dron & Anderson, 2014);
3. how we become parts of one another's cognition through the technologies in which we participate and that participate in us (e.g. Dron, 2023).

These examples only touch on CI's broad usage to denote everything from individual brain organization (e.g. Benjamin et al., 2023), to the organization of groups and networks (e.g. Woolley et al., 2015), to 4E cognition, (e.g. Clark, 2008) to the concept of a global brain (e.g. Bloom, 2000). Such varied uses describe different kinds of cognition, and so different kinds of responses needed from our educational systems to them.

This paper introduces the term "ochlotecture", from the Classical Greek ὄχλος (*ochlos*), meaning "multitude" and τέκτων (*tektōn)* meaning "builder" to describe the structures and processes that connect groupings of people. I use this concept to describe the ochlotecture of various kinds of CI, before discussing the distinctive ochlotecture of generative AIs and their potential impact on human learning.

The ochlotecture of collectives

CI emerges from the dynamic and recursive interplay between structure and behaviour in a network of interacting agents and their environment. CIs are not just the sum of their parts but cognitive agents in their own right. They are Kantian Wholes (S. Kauffman, 2022) in which the whole (the CI) exists for and by means of its parts (the interacting agents). In earlier work (e.g. Dron, 2003) I called these "collectives", because "intelligence" is a fuzzy, contested term and it is at least as possible to be collectively stupid as it is to be collectively smart.

Anderson and I (2014) have argued that there are three distinct phenomena that define a collective, that I present here as fundamental ochlotectural elements:

1. Information gathering: concerned with perception and selection - the selective capture of the signals generated by actions (and sometimes inactions) of individual members of a

collective. Axiomatically, a collective, or its individual agents, must possess the means to perceive its own activities or the results of them.

2. Information processing: concerned with manipulating and (re)structuring what has been gathered. This relates to the algorithms, procedures, methods, and rules through which the captured signals are filtered, aggregated and transformed. The algorithms may be internal to the members of the collective (e.g. flocking) or external to it (e.g. voting, collaborative filtering) or both. Axiomatically, a collective must be able to perform some cognitive task based on its perceptions.
3. Information presentation: concerned with (re)presentation of the processed information as either signs created in the environment (e.g. pheromones, words), changes to that environment (e.g. untidy nests), or the result of perceptions of behaviours of agents in a collective themselves (e.g. herding, flocking). Axiomatically, a collective must be able to represent the output of a cognitive task, whether in its behaviour or by making changes to the environment around it, or both

The smartness of an ochlotecturally simple collective lies almost entirely in the ochlotecture of the whole, not in its member agents. The stigmergic behaviour of investors in stock markets is hardly more intelligent than that of termites building termite mounds. Likewise many network effects deriving from ochlotectural features such as preferential attachment, cyclical structures, or small world topologies (Kearns et al., 2006). The roles played by individuals are, at least in aggregate, largely predictable and rule-bound.

The smartness of individual agents normally matters more in the collective intelligence of intentional groups, and more elements contribute to the ochlotecture. For example, the ochlotecture of teams or universities includes schedules, rules of conduct, policies, decision-making processes, and so on, plus any physical or virtual technologies it relies on, like lecture theatres or whiteboards. Simpler collective processes are at work, too, such as groupthink (Janis, 1972) and network effects. Though very mechanical groups exist (e.g. production lines) this often demands *soft technique* (Dron, 2023) of its members to succeed. "Technique", as I use the term here, is simply a human-instantiated technology. "Soft technique" refers to how we may idiosyncratically and creatively exploit the adjacent possibles (S. A. Kauffman, 2019) of a technology, filling the gaps left open in its orchestration with additional orchestrations to adapt it to new and, often, unprestatable uses. This is in contrast to the hard technique that dominates in simpler ochlotectures, in which we play determinate roles correctly as part *of* the orchestration. In a typical group, a complete description of the ochlotecture includes both the structures and processes in the environment and the ever-unfolding soft techniques of individual agents within it.

Still greater cognitive ability in agents is needed in "extended cognition" collectives. Parts such as words, aircraft controls, or musical phrases provide signals that we may use to endlessly assemble novel technologies: sentences, flights, songs, ideas and so on. Though rigid forms of participation may be needed to achieve some of the benefits – there can be no creativity without some constraint (Boden, 1995) - the components and the ways in which they may be combined are virtually unlimited: it is an open and very soft ochlotecture comprised largely of idiosyncratic soft technique.

All collectives teach, by influencing the cognition of their members. In education, effective collectives should be structured to guide learning toward specific goals. For example, Facebook is a poor teacher because its algorithms are designed for engagement, not teaching. Learning does occur, but not necessarily the kind that users would intentionally seek.

The ochlotecture of large language models

Large language models (LLMs) are collective applications that assemble signals (prompts and data), process them, and re-present them, affecting their human users and being affected by them. The signals they assemble are typically culled from unprecedentedly huge pools of human-created data, culled from years of web crawling, books, Wikipedia articles, Reddit posts, and more, representing a not-insignificant portion of all digitally recorded human knowledge. Like other collectives, feedback loops drive their evolution: recent LLMs make extensive use of data pulled from interactions of users with earlier versions. And, of course, there are algorithms but, unlike most other digital collective apps such as collaborative filters or tag clouds, the algorithms are flexible: highly adaptive, unique to every prompt, with billions (soon trillions) of parameters, making them practically inscrutable and seldom repeating twice. The size and flexibility of their ochlotecture makes them capable of human-like soft technique (Dron, 2023), that is both the source of their usefulness and the reason to fear their effects.

Unlike all previous technologies, generative AIs use other technologies such as words in both non-random and non-programmed ways, and we may interact with them in much the same way that we interact with other people. What we learn in the process is not just a set of facts or skills but also attitudes, values, and ways of thinking: tacit knowledge that is imparted in any act of communication, whether or not it involves formal teaching. We have seldom if ever had to assess this or to purposely teach it in formal education because it is unavoidable when we learn from and with other people. The role for education in developing tacit knowledge, however, is at least as important as the achievement of intended learning outcomes, and is central to how we learn to be, not just to do, in our various cultures and societies.

LLMs have no attitudes, no intentions, no values of their own, save insofar as they are trained on selectively chosen datasets, with inputs and outputs massaged by intentional coding, which is itself troubling because of the power it brings to their owners and creators. Chameleon-like, they will slip into any identity we ask of them unless intentional programming prevents them from doing so, again reflecting owners' and creators' biases. From their interminable patience to their tendency to be, as Dave Cormier (2023) puts it, "autotune for knowledge", to the biases that are introduced in their training, to the quirks in responses because they lack contextual understanding of what they create, LLMs are made of human knowledge and thus appear human-like, but they are not human. It is therefore worrying when, for example, an LLM is used to write or, worse, tell a children's story from which persistent values and attitudes may be formed. However, every communication is a story, whatever our age. Already, in our interactions with them, we cannot help but learn values, attitudes, and ways of thinking from non-human machines.

Furthermore, they increasingly replace the need to perform tasks like drawing, programming, or problem solving, explicit knowledge and skills may diminish. We know, for example, that students who learn through interacting conversationally with ChatGPT are, without such assistance, less capable of independent problem-solving than those who do not (Bastani et al., 2024). We can, of course, rightly argue that we just have a new ratchet that can lift our capacity to create higher than before, as long as we are using the generative AI. However, LLMs' ochlotecture means they replace soft as well as hard technique: unlike any prior technology, they may thus diminish our capacity to create. It is particularly problematic due to the implied relationship with a human creator that we have previously been able to take as a given, if we have thought of it at all, made by people like us solving problems and dreaming dreams like ours.

More and more, AIs will replace humans as educators because they are cheaper, more available, and effective in achieving measurable outcomes. This can and will change us, at an unimaginable global scale. It will be both good and bad, but we are collectively unprepared. We have been able to largely ignore education's fundamental role in tacit knowledge and value creation because, deriving solely from humans, it came for free. We can ignore it no longer.

The matter becomes more pressing because, increasingly, the training sets of generative AIs will themselves be the outputs of previous generations of generative AIs, leading to so-called model collapse (Shumailov et al., 2024) but, even if that can be avoided, their inputs will increasingly be the outputs of people who have learned with and been changed by generative AIs, magnifying and cementing the effects. All CI is an Ouroboros, feeding on itself as much as it imbibes information from its environment: this is its strength. It shapes its parts, and its parts shape it, in a continuing cycle of adaptation of the Kantian Whole. The big difference with LLMs lies in the scale of uptake, especially bearing in mind the dominance of a few key players and the fact that the input itself gets degraded, like a photocopy of a photocopy. This is a global shift that will affect most of human-kind.

Conclusions

Ochlotecturally, I believe it is important to acknowledge generative AIs as independent entities within human collectives, not just as collectives in their own right, and to treat them as fellow agents rather than tools, thus preserving our own agency and recognizing their distinctive nature as fellow tool-users. Where possible, we should protect them from their own outputs, feeding them primary, human-created data. And, where possible, as educators we should focus on what is valuable in relationships with other humans, doing our best to counter-balance the effectiveness of generative AIs in supporting the development of what is measurable with activities and engagements that support the immeasurable, the social, the soft, the unprestatable. Given the increasing focus on measurable outcomes that has driven educational systems for (at least) many decades, this may imply a sea-change in how we view learning and assessment. We must seek ways to harvest learning outcomes, not just to measure compliance in meeting those we intend. We should nurture and seek to make visible the messy, complex, tangential process of learning, not just its product, celebrating its deep entanglement with the ochlotecture from which it arises.

To make this possible, we should attempt to decouple teaching from credentialling. Embedded credentialling reduces the situated complexity of learning to a McNamara-esque caricature, driven by what can easily be measured to the exclusion of much of what gives it value. Meanwhile, the ochlotectural power imbalance that it entails decimates support for autonomy, competence, and relatedness needs without which intrinsic motivation cannot emerge (Ryan & Deci, 2017), replacing them with extrinsically driven targets that generative AI can hit at least as well. The ochlotecture of an institution that adopts a decoupled approach may lose many of the rigid power-drenched hierarchies and rules of conduct that are characteristic of classes and courses today, replacing them with more distributed, fluid, diverse, evenly distributed networks and ad hoc groupings of learners together with those who support them which, from a CI perspective, is also likely to result in smarter collective behaviours (Page, 2008). Where credentials are needed, the evidence from harvested outcomes and traces of the learning process may suffice or, following the lead set by Brunel University in the UK, we may create interdisciplinary integrative assessed learning activities that acknowledge the diversity and interests of individuals, and that are sufficiently authentic that they reward and leverage rather than punish and inhibit the use of CI, in all its forms, including generative AIs.

The immeasurability of tacit learning, however, means that we do not know what we are currently losing: though we clearly do learn attitudes, ways of being, and values through our interactions with other people and the artifacts they create, research into how we do so is at best very general and difficult to reliably capture. Often, we express the tacit only indirectly, whether through art or action, and, by its nature, it cannot be quantified. Though unquantifiable and complex, research that uncovers how we are changing and the role of generative AI in bringing about that change is vital. We cannot respond intelligently without knowing these effects. Understanding the ochlotecture of AI-imbued collectives may provide a useful piece of that puzzle.

References

Bastani, H., Bastani, O., Sungu, A., Ge, H., Kabakcı, Ö., & Mariman, R. (2024). *Generative AI Can Harm Learning* (SSRN Scholarly Paper No. 4895486). https://doi.org/10.2139/ssrn.4895486

Benjamin, F. J., Kaaronen, R. O., Moser, C., Rorot, W., Tan, J., Varma, V., Williams, T., & Youngblood, M. (2023). All intelligence is collective intelligence. *Journal of Multiscale Neuroscience*, *2*(1), 169–191.

Bloom, H. (2000). *Global Brain: The evolution of mass mind*. Wiley.

Boden, M. (1995). Creativity and unpredictability. *Stanford Humanities Review*, *4*(2), 123–139.

Clark, A. (2008). *Supersizing the Mind: Embodiment, Action, and Cognitive Extension*. Oxford University Press.

Cormier, Dave. (2023, January 20). *ChatGPT search – Autotune for knowledge – Dave's Educational Blog*. http://davecormier.com/edblog/2023/01/20/chatgpt-search-autotune-for-knowledge/

Dron, J. (2003). *The Blog and the Borg: A Collective Approach to E-Learning*. E-Learn 2003.

Dron, J. (2007). *Control and Constraint in E-Learning: Choosing when to Choose*. Idea Group International.

Dron, J. (2023). *How Education Works: Teaching, Technology, and Technique*. AU Press.

Dron, J., & Anderson, T. (2014). *Teaching crowds: Learning & Social Media*. AU Press.

Janis, I. (1972). *Victims of Groupthink*. Houghton Mifflin.

Kauffman, S. (2022). Autogen is a Kantian Whole in the Non-Entailed World. *Biosemiotics*, 1–4. https://doi.org/10.1007/s12304-021-09466-4

Kauffman, S. A. (2019). *A World Beyond Physics: The Emergence and Evolution of Life*. Oxford University Press.

Kearns, M., Suri, S., & Montfort, N. (2006). An Experimental Study of the Coloring Problem on Human Subject Networks. *Science*, *313*(5788), 824–827.

Page, S. E. (2008). *The Difference: How the Power of Diversity Creates Better Groups, Firms, Schools, and Societies (New Edition)* (Kindle). Princeton University Press.

Ryan, R. M., & Deci, E. L. (2017). *Self-determination theory: Basic psychological needs in motivation, development, and wellness*. Guilford Publications.

Shumailov, I., Shumaylov, Z., Zhao, Y., Papernot, N., Anderson, R., & Gal, Y. (2024). AI models collapse when trained on recursively generated data. *Nature*, *631*(8022), 755–759. https://doi.org/10.1038/s41586-024-07566-y

Woolley, A. W., Aggarwal, I., & Malone, T. W. (2015). Collective Intelligence and Group Performance. In *Current Directions in Psychological Science* (Vol. 24, Issue 6, pp. 420–424). https://journals.sagepub.com/doi/abs/10.1177/0963721415599543