

Document Version

Final published version

Citation (APA)

Browne, J. T., & Garnham, I. (2022). Extended, Distributed, and Predictive: Sketches of a Generative Theory of Interaction for HCI. In C. Stephanidis, M. Antona, S. Ntoa, & G. Salvendy (Eds.), *HCI International 2022 – Late Breaking Posters - 24th International Conference on Human-Computer Interaction, HCII 2022, Proceedings* (pp. 17-24). (Communications in Computer and Information Science; Vol. 1654 CCIS). Springer. https://doi.org/10.1007/978-3-031-19679-9_3

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

In case the licence states “Dutch Copyright Act (Article 25fa)”, this publication was made available Green Open Access via the TU Delft Institutional Repository pursuant to Dutch Copyright Act (Article 25fa, the Taverne amendment). This provision does not affect copyright ownership.
Unless copyright is transferred by contract or statute, it remains with the copyright holder.

Sharing and reuse

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

Green Open Access added to TU Delft Institutional Repository

'You share, we take care!' - Taverne project

<https://www.openaccess.nl/en/you-share-we-take-care>

Otherwise as indicated in the copyright section: the publisher is the copyright holder of this work and the author uses the Dutch legislation to make this work public.



Extended, Distributed, and Predictive: Sketches of a Generative Theory of Interaction for HCI

Jacob T. Browne^{1,2}  and Ignacio Garnham³

¹ Philips Experience Design, Eindhoven, The Netherlands

jacob.browne@philips.com

² Delft University of Technology, Delft, The Netherlands

³ Aarhus University, Aarhus, Denmark

igarnham@cc.au.dk

Abstract. This paper blends work in extended mind, distributed cognition, and predictive processing to provide a novel generative theory of interaction. This dovetailing offers an emerging picture of cognition that HCI stands to benefit from: our cognition is extended, distributed, and constantly trying to predict incoming sensory stimuli across social, cultural, and temporal scales. We develop a sketch of a generative theory of interaction for HCI and offer some directions for future work.

Keywords: Extended mind · Predictive processing · Distributed cognition · HCI · Artificial Intelligence · Generative theory of interaction

1 Introduction

HCI has a long history of borrowing concepts from other sciences to guide HCI research. In a synthesis of the histories of HCI Theory, Rogers traces different theories that have come to fruition over the classical, modern, and contemporary stages of HCI [31]. Rogers reflects that importing and building on theories is crucial to HCI: that creating a transdisciplinary agenda will further our impact on society [31].

As our technologies become more complex, it becomes more apparent that the boundaries between mind and technology are blurring. This has spiked concerns in HCI around how to design for and study these complex technologies [12]. Human-centered design and contemporary models in HCI are rendered incompetent to deal with these questions regarding ontological uncertainty: we're struggling to keep up with what is being empirically observed [12]. Frauenberger puts it nicely, "HCI may not yet be in a state of serious crisis, but it is certainly cracking and squealing, struggling to make sense of computers, humans and interfaces in the face of rapid technological progress, coupled with profound social change" [12]. Theoretical and empirical work in predictive processing, extended mind, and distributed cognition similarly point in this direction.

In this paper, we will outline the basics of extended mind, distributed cognition, and predictive processing. A blending of these theories is presented. Then, we will offer some preliminary theoretical implications for HCI, offering an early form of a generative

theory of interaction [1]. This paper makes two contributions: a rendering of a developing picture of cognition for HCI and a direction for a new generative theory of interaction.

2 Extended Mind, Distributed Cognition, and Predictive Processing

This section will present a brief sketch of each body of work. The intent is not to serve as an all-encompassing review or to argue for specific flavors of each respective body of work. Instead, we aim to present what might be most crucial to understand for HCI researchers for each body of work.

2.1 Extended Mind

The extended mind thesis (EMT) maintains that, under certain conditions, cognitive states and processes can be distributed across brain, body, and world [7]. This means that, should an artifact or prop be used in the right way, it would constitute a literal part of one's mind. This thesis stands against cognitivist intracranialism and poses a new "active externalism". Recently, Chalmers (with Clark's approval) presents a refined version of the thesis: "A subject's cognitive processes and mental states can be partly constituted by entities that are external to the subject, in virtue of the subject's sensori-motor interaction with these entities" [2]. Chalmers places a focus on what's most interesting about the thesis: how an intercranial implant wouldn't be much of a surprise to constitute an extension but using a notebook to remember an address or moving tiles on a rack during a Scrabble game do count as extensions [2, 7, 26]. Extensions do not need to replicate neural functioning or have the same dynamics but could complement our cognitive abilities [33]. The extended mind is a soft assembly made up of whatever elements are needed to accomplish the task an agent faces [25]. While this has profound ontological and epistemological implications, we will avoid them for now.

Clark has stated that HCI can be seen as a sort of nascent science of the extended mind [3]. In this direction, Heersmink offers us a multidimensional framework for understanding the degree of extendedness to guide empirical research [14]. The different dimensions include information flow, reliability, durability, trust, procedural transparency, informational transparency, individualization, and transformation [14]. Systems can ebb and flow between degrees of cognitive extendedness, varying in time and by need. The higher the scores across dimensions, the denser the integration [14].

2.2 Predictive Processing

Predictive processing (PP) is a "framework involving a general computational principle which can be applied to describe perception, action, cognition, and their relationships in a single, conceptually unified manner" [36]. Our brain is seen as a prediction machine [5]. It's constantly attempting to match incoming sensory information with top-down predictions through a probabilistic generative model [5]. This is to reject the standard passive, stimulus-driven view of perception, action, and the brain, instead viewing it as highly active and predictive of incoming stimuli, a controlled, online hallucination [5, 8, 36].

Any unpredicted elements are propagated as information back into the predictions in the form of prediction errors [5]. Through active inference, actions are then recruited to improve the state of information within the environment and reduce future prediction error, reducing free energy (known as epistemic actions) [5, 13, 27, 30]. Prediction error minimizing is constantly modulated by precision weighting, adjusting the gain on specific error units, or unexplained sensory units, until an error signal is suppressed [30]. There is a great deal of traffic between interoceptive and exteroceptive information as well, where bodily states (e.g., hunger) can modulate what information in the environment is salient (e.g., food) [8]. PP is particularly alluring given it unifies the mechanics of action, perception, attention, emotion, language, and learning into one unified framework, while bridging empirical and theoretical work [6, 30, 36].

2.3 Distributed Cognition

Distributed cognition (DCog) is a branch of cognitive science, a perspective that seeks to understand cognitive systems by extending the unit of analysis for cognition beyond the boundaries of the skin or skull [16–18]. DCog takes the body, social, and material world as fundamental to understanding cognition: “a [functional] system that can dynamically configure itself to bring subsystems into coordination to accomplish various functions... delimited by the function; relationships among the elements that participate in it, rather than by the spatial colocation of the elements” [16, 17, 21]. DCog has historically been adopted to study human activity ‘in the wild’, noting how cognitive processes are distributed across people, time, and internal/external structures [16, 17]. It recognizes that “humans create their cognitive powers in part by creating the environments in which they exercise those powers”, as ecological assemblies [17, 22]. Cognition doesn’t only happen in the brain, rather it is distributed between people and technology, across time in the development of social and material contexts.

A profound element of DCog is its realization of culture. DCog renders culture as “a process that accumulates partial solutions to frequently encountered problems”, and “a complex cognitive ecosystem that includes, in addition to the brain, a large number of somatic and extrasomatic processes” [16, 19]. Cultural practices are “the things people do and their [learned] ways of being in the world... [that] organize the action in situated action. Emergent products of dynamic distributed networks of constraints” [20]. Culture serves to reduce entropy at multiple scales [22]. Our cultural practices shape our perception as active ways of “seeing” the world by indicating what to pay attention to and what to perceive [19]. For instance, seeing a constellation of stars is a process enacted via cultural practices which enable you to give visual attention in that specific way to see the constellation [23]. Our environment itself consists of dynamic products nearly entirely due to prior cultural activity [20]. High-level cognitive outcomes are born through the coordination of the mechanisms of distributed cognitive systems within these cultural practices [23].

The questions distributed cognition asks concern “the relations among the elements [of the cognitive system], and how cognitive processes arise from interactions among those elements” [22]. It looks for principles that might apply across scales and cognitive systems (e.g., all cognitive systems are characterized by “local regions of high interconnectivity separated by regions of lower interconnectivity”) [22]. DCog doesn’t assume

a center for any cognitive system. Instead, the centers are determined by the “relative density of information flow across a system” [22]. This is to say that some systems have an obvious center, others several centers, others no center. It all depends on the scaling in question.

3 Dovetailing Sketches of Cognition and Implications for HCI

HCI isn’t new to many of these ideas. For instance, distributed cognition has been used to understand complex sociotechnical systems in a variety of contexts and was borne from researchers in HCI [16, 31]. These different theories complement each other and are relatively consistent in principle. According to the authors, this would be the first attempt to mesh these ideas together to form a coherent generative theory of interaction. So, what is this composite depiction of cognition?

3.1 Extended, Distributed, and Predictive: A Low-Fidelity Prototype

Recent compositions of extended mind, predictive processing, active inference, and cognitive niche construction are seen by Constant et al. in their work on “extended active inference” and Veissière et al. in their work on “thinking through other minds” (TTOM) [9, 34]. A complete review of these works is beyond the scope of this paper, and instead we’ll point out some relevant aspects.

Constant et al. depict a cognitive niche (externally realized cause-and-effect models) where epistemic cues and affordances are proliferated throughout our inherited environments, reducing uncertainty through ecological legacies, as a shared generative model [9]. Veissière et al. depict “regimes of attention”, which highlight different contextually relevant actions as particularly important and aid agents in learning to adapt to their local niche, shaping these epistemic cues [34]. Cognitive niche construction is shown as a “shared cognitive function enabling organisms to track—often implicitly and at low cost—cause–effect relationships otherwise difficult” [9]. This is especially apparent in cases where causal structures are volatile or too complex to be learned by the agents own sensory mechanisms [9, 34].

Learning to use these epistemic cues reduces the complexity of our own generative models, increasing our performance. Constant et al. call the leveraging and optimizing of this shared generative model “extended active inference” (EAI), where active inference is “optimizing an organism’s generative model about the cause of its sensations” through action [9]. This leads to the maximization of model evidence through perception, and selective sampling of expected sensory sensations through action. Cognitive niche construction is a cognitive function: uploading and offloading into the niche. Once uploaded into the environment, cognitive extensions can be shared by other agents, as the “scaffolding of complex networks of shared expectations encoded across brains, bodies, constructed environments, and other agents, which modulate attention, guide action, and entail the learning of patterned behaviours” [34]. The environment can be seen as a generative model of the agent. As Constant et al. state, “...one can treat the environment as inferring the cause of the “sensations” it receives from being acted upon by its denizens”, or as Veissière puts it, “that the affordances of the–environment and the

capacities of an individual are inextricably interwoven, and co-determining” [9, 34]. The cognitive niche and the agent are constantly trying to optimize their generative models of each other.

3.2 Implications for a Generative Theory of Interaction in HCI

HCI can make use of this blending of theories through a new generative theory of interaction [1]. Following Beaudouin-Lafon et al., a generative theory of interaction is:

- Grounded in a theory of human activity and behavior with technology [1]
- Involves analytical (“a description of current use and practice”), critical (“assesses both the positive and negative aspects of a system”), and constructive lenses (“inspires new ideas relative to the critique”) [1]
- Provides tools that allow examination of the design space related to a research problem [1].

In the next section, we’ll briefly sculpt out how blending predictive processing, extended mind, and distributed cognition might fit into a generative theory of interaction.

Theoretical Grounding

We can see that our blending of predictive processing, distributed cognition, and extended mind is well grounded in theory. Through predictive processing and active inference lenses, we’re able to render distributed cognition and extended mind through similar principles. Extended mind provides the basis for blurring the boundary between technology and the brain, laying the groundwork for extending the mind beyond the boundaries of the skull. Distributed cognition provides the temporal and cultural lenses to view these extensions. Predictive processing provides the computational principles which govern the different scales. This is to render culture, environments, tools, brains, bodies, and other generative models as mechanisms for uncertainty reduction [32]. Cultural practices can be seen as shaping epistemic cues: highlighting affordances, uploading cognitive functions, extending precision weighting, shared across time and people through cognitive niches. Our cognitive capacities constitute our tools, culture, environment, etc.: “ontologically inseparable from the start” [24]. As Herbert Simon is said to remark, “Most human intelligence is artificial intelligence anyway”.

We can surface some early concepts (cognitive extensions, temporality) and principles (as adopted from Heersmink: information flow, reliability, durability, trust, procedural transparency, informational transparency, individualization, and transformation) to examine analytical, critical, and constructive lenses [1, 14].

Analytical, Critical, and Constructive Lenses

Analytically, we can begin to look at how well certain extensions extend cognition and how this might vary across time by [14, 15]. How well do these cognitive functions embed into the cognitive niche? How does it increase uncertainty reduction across different time scales? Critically, we can begin looking at who gets the opportunity to extend their minds with different extensions and how has that access propagated across society, in what forms? Who has more power to design the cognitive niches and what types

of cognitive niches are most dominant [10]? Constructively, we can begin to consider how we could better extend people’s minds, increase access to different extensions, or increase people’s ability to design their cognitive niche. We believe that Heersmink’s dimensions of integration offer a good starting point to consider different principles to evaluate cognitive systems upon (information flow, reliability, durability, trust, procedural transparency, informational transparency, individualization, and transformation) [14].

4 Conclusion

In this paper, we sought to present a developing generative theory of interaction through blending extended mind, distributed cognition, and predictive processing into a single account. We chose these theories given how well they complement each other, and how they lack directly actionable principles to guide HCI research. For future work, we’ll more closely develop this generative theory of interaction by offering more insight into a theoretical grounding, offering well-defined principles and concepts to look at this theory through, and develop questions that can be asked by looking at different lenses (as depicted by Beaudouin-Lafon et al.) [1]. We hope that we’ve offered a starting point for others to more directly consider how we can bring these theories into HCI and open channels for transdisciplinary work with these fields.

Acknowledgements. This work is part of the DCODE project. The project has received funding from the European Union’s Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No. 955990.

References

1. Beaudouin-Lafon, M., et al.: Generative theories of interaction. *ACM Trans. Comput. Hum. Interact.* **28**(6), 1–54 (2021). <https://doi.org/10.1145/3468505>
2. Chalmers, D.J.: Extended cognition and extended consciousness. In: Andy Clark and His Critics, pp. 9–20. Oxford University Press (2019)
3. Clark, A.: Memento’s revenge: the extended mind, extended. In: *The Extended Mind*, pp. 43–66. The MIT Press (2010)
4. Clark, A.: *Natural-Born Cyborgs: Minds, Technologies, and the Future of Human Intelligence*. Oxford University Press, New York (2004)
5. Clark, A.: *Surfing Uncertainty: Prediction, Action, and the Embodied Mind*. Oxford University Press, New York (2019)
6. Clark, A.: Whatever next? Predictive brains, situated agents, and the future of cognitive science. *Behav. Brain Sci.* **36**(3), 181–204 (2013). <https://doi.org/10.1017/s0140525x12000477>
7. Clark, A., Chalmers, D.: The extended mind. *Analysis* **58**(1), 7–19 (1998). <https://doi.org/10.1111/1467-8284.00096>
8. Clark, A.: Consciousness as generative entanglement. *J. Philos.* **116**(12), 645–662 (2019). <https://doi.org/10.5840/jphil20191161241>
9. Constant, A., et al.: Extended active inference: constructing predictive cognition beyond skulls. *Mind Lang. Mila.* 12330 (2020). <https://doi.org/10.1111/mila.12330>

10. Costanza-Chock, S.: *Design Justice: Community-Led Practices to Build the Worlds We Need*. MIT Press, London (2020)
11. Engel, A.K., et al.: Dynamic predictions: oscillations and synchrony in top-down processing. *Nat. Rev. Neurosci.* **2**(10), 704–716 (2001). <https://doi.org/10.1038/35094565>
12. Frauenberger, C.: Entanglement HCI the next wave? *ACM Trans. Comput. Hum. Interact.* **27**(1), 1–27 (2020). <https://doi.org/10.1145/3364998>
13. Friston, K., et al.: Active inference and agency: optimal control without cost functions. *Biol. Cybern.* **106**(8–9), 523–541 (2012). <https://doi.org/10.1007/s00422-012-0512-8>
14. Heersmink, R.: Dimensions of integration in embedded and extended cognitive systems. *Phenomenol. Cogn. Sci.* **14**(3), 577–598 (2014). <https://doi.org/10.1007/s11097-014-9355-1>
15. Heersmink, R.: The cognitive integration of scientific instruments: information, situated cognition, and scientific practice. *Phenomenol. Cogn. Sci.* **15**(4), 517–537 (2015). <https://doi.org/10.1007/s11097-015-9432-0>
16. Hollan, J., et al.: Distributed cognition: toward a new foundation for human-computer interaction research. *ACM Trans. Comput. Hum. Interact.* **7**(2), 174–196 (2000). <https://doi.org/10.1145/353485.353487>
17. Hutchins, E.: Cognition, distributed. In: *International Encyclopedia of the Social & Behavioral Sciences*, pp. 2068–2072. Elsevier (2001)
18. Hutchins, E.: *Cognition in the Wild*. MIT Press, Cambridge (1994)
19. Hutchins, E.: Cognitive ecology. *Top. Cogn. Sci.* **2**(4), 705–715 (2010). <https://doi.org/10.1111/j.1756-8765.2010.01089.x>
20. Hutchins, E.: Enculturating the supersized mind. *Philos. Stud.* **152**(3), 437–446 (2011). <https://doi.org/10.1007/s11098-010-9599-8>
21. Hutchins, E.: Imagining the cognitive life of things. <https://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.104.6035&rep=rep1&type=pdf>. Accessed 18 Mar 2022
22. Hutchins, E.: The cultural ecosystem of human cognition. *Philos. Psychol.* **27**(1), 34–49 (2014). <https://doi.org/10.1080/09515089.2013.830548>
23. Hutchins, E.: The role of cultural practices in the emergence of modern human intelligence. *Philos. Trans. R. Soc. Lond. B Biol. Sci.* **363**(1499), 2011–2019 (2008). <https://doi.org/10.1098/rstb.2008.0003>
24. Introna, L.D.: Towards a post-human intra-actional account of sociomaterial agency (and morality). In: Kroes, P., Verbeek, P.-P. (eds.) *The Moral Status of Technical Artefacts*. PET, vol. 17, pp. 31–53. Springer, Dordrecht (2014). https://doi.org/10.1007/978-94-007-7914-3_3
25. Kirchhoff, M.D., Kiverstein, J.: *Extended Consciousness and Predictive Processing: A Third Wave View*. Routledge, London (2019)
26. Kirsh, D.: The intelligent use of space. *Artif. Intell.* **73**(1–2), 31–68 (1995). [https://doi.org/10.1016/0004-3702\(94\)00017-u](https://doi.org/10.1016/0004-3702(94)00017-u)
27. Kirsh, D., Maglio, P.: On distinguishing epistemic from pragmatic action. *Cogn. Sci.* **18**(4), 513–549 (1994). https://doi.org/10.1207/s15516709cog1804_1
28. Marshall, P., Necker, E.H.: Theories of embodiment in HCI. In: *The SAGE Handbook of Digital Technology Research*, pp. 144–158. SAGE Publications Ltd., London (2014)
29. Menary, R. (ed.): *The Extended Mind*. The MIT Press, Cambridge (2010)
30. Nave, K., et al.: Wilding the predictive brain. *Wiley Interdiscip. Rev. Cogn. Sci.* **11**(6), e1542 (2020). <https://doi.org/10.1002/wcs.1542>
31. Rogers, Y.: *HCI Theory: Classical, Modern, and Contemporary*. Morgan & Claypool, San Rafael (2012)
32. Sterelny, K.: Minds: extended or scaffolded? *Phenomenol. Cogn. Sci.* **9**(4), 465–481 (2010). <https://doi.org/10.1007/s11097-010-9174-y>
33. Sutton, J.: Exograms and interdisciplinarity: history, the extended mind, and the civilizing process. In: *The Extended Mind*, pp. 189–225. The MIT Press (2010)

34. Veissière, S.P.L., et al.: Thinking through other minds: a variational approach to cognition and culture. *Behav. Brain Sci.* **43**, e90 (2020). <https://doi.org/10.1017/s0140525x19001213>
35. Verbeek, P.-P.: Beyond interaction: a short introduction to mediation theory. *Interactions* **22**(3), 26–31 (2015). <https://doi.org/10.1145/2751314>
36. Wiese, W., Metzinger, T.K.: Vanilla PP for philosophers: a primer on predictive processing. (2017). <https://doi.org/10.15502/9783958573024>
37. Wilson, R.A., Clark, A.: How to situate cognition: Letting nature take its course. In: Aydede, M., Robbins, P. (eds.) *The Cambridge Handbook of Situated Cognition*, pp. 55–77. Cambridge University Press, Cambridge (2009)