Commentary on: Michael Gilead, Yaacov Trope, and Nira Liberman, "Above and Beyond the Concrete: The Diverse Representational Substrates of the Predictive Brain"

Word counts:

Abstract: 60
Main text: 984
References: 300
Entire document: 1,425

Scale-free architectures support representational diversity

Chris Fields[1] and James F. Glazebrook[2]

[1] Independent, 23 Rue des Lavandières, 11106 Caunes-Minervois, FRANCE
fieldsres@gmail.com; +33 6 44 20 68 69; https://chrisfieldsresearch.com
[2] Department of Mathematics and Computer Science, Eastern Illinois University, Charleston, IL 61920, USA jfglazebrook@eiu.edu; +1 217-581-5328; https://faculty.math.illinois.edu/~glazebro/

Abstract: Gilead, Trope and Liberman propose an ontology of abstract representations based on folk-psychological conceptions of cognitive architecture. There is, however, no evidence that the experience of cognition reveals the architecture of cognition. Scale-free architectural models propose that cognition has the same computational architecture from sub-cellular to whole-organism scales. This scale-free architecture supports representations with diverse functions and levels of abstraction.

Gilead, Trope and Liberman propose an "ontology" of representation types, argue that this ontology captures "meaningful diversity in the representational substrates of the mind," and criticize the architectural assumptions of predictive-coding models as "overly simplistic" (both p. 54). It is never entirely clear what the elements of this ontology are – a table would have helped – but the following all seem to be included: beliefs, desires, intentions, conditions of satisfaction, subjectively distinguishable objects, features and relations represented as modality-specific, multimodal, or categorical abstractions, episodes, "lemmas" defining words, semantic and temporal networks, hierarchies, "predicators" functioning as "mentalese" verbs, models, scripts, and simulations. The distinctions between these various entities are localized to Marr's computational level of analysis (p. 8); however, the critique of "sub-symbolic" architectures and focus on a "layer of language-like mental representations" (p. 53) suggest an implementation-level analysis. This distinction is critical, as few would argue that different types of representations at different levels of abstraction do not have different roles in cognition. Neuroimaging results demonstrating functional localization, for example, support functional but not architectural distinctions between types of representations.

An unstated assumption of this ontology appears to be that the structure of conscious experience is a reliable guide to the architecture of the neurocognitive system that implements this experience, including the structure of its representations. The "rich and intricate theoretical conceptualizations" that predictive-coding models are claimed to have ignored (p. 5) are conceptualizations of the structure of a particular kind of experience, the experience of *thinking*. Hoffman and colleagues (2015; 2018) have argued on evolutionary grounds that perceptual experience is an "interface" onto the external world that supports the prediction of fitness consequences of actions but provides no reliable guide to the structure or dynamics – the architecture – of the external world. This argument can easily be inverted: conceptual experience is an interface that provides no reliable guide to the architecture of cognition. Just as humans have, in general, no need to know how computers work to use them effectively, humans have no need to know how their minds work to operate effectively in the world. A simplified folk "theory of mind" on the interface is good enough. Similar points have been made before, e.g. by Chater (2018).

Relinquishing the assumption that the experience of cognition constrains the architecture of cognition is, we argue, the key to making significant progress in cognitive science. It enables asking: how is the experience of cognition produced as an *output*, and what inputs and inferential processes are needed to produce that output? Assuming the Church-Turing thesis, all computation is platform-independent: any collection of diverse representations can be generated, in principle, by any Turing-complete virtual machine. The central claim of AI, often rendered just as "cognition is computation," is actually that cognition is *platform-independent*. It remains far from obvious, however, how to implement cognition on *any* platform, including the human brain-body system. Nor is it obvious that understanding one implementation of cognition would provide useful hints toward understanding other implementations.

The claim that cognition is scale-free is far stronger than platform independence: it is the claim that a single computational architecture works "all the way down" – describing every virtual machine at every useful level of analysis. Gilead, Trope and Liberman recognize this when they describe the theory of active inference, the dominant current scale-free proposal, as claiming that "the complexity of cognition can naturally arise from a canonical computation repeated across different layers of a single continuum of representational abstractness" (p. 52). The theory of active inference is scale-free because the free-energy principle on which it is based is scale free (Friston, 2013; Friston et al., 2015), with its underlying basis, the existence of Markov blankets, derivable from classical (Kuchling et al, 2019) and even quantum (Fields and Marcianò, 2019) physics. We have proposed an alternative, category-theoretic, scale-free formulation in which inferential coherence is enforced by commutativity between within-scale and between-scale mappings (Fields and Glazebrook, 2019); our proposal is in the spirit of Goguen's (1991) dictum, within computer science, that abstraction *always* corresponds to the construction of a category-theoretic "cocone" as a maximal representation of inferential coherence.

Scale-free models have the advantage of being rigorously testable at every experimentally-accessible level of analysis, from those of basic physics, intracellular, and cellular processes up to the whole-organism scale and beyond. At every level, they must specify explicitly what inputs are required and what outputs are produced; indeed the role of the Markov blanket is to provide an explicit encoding of these inputs and outputs. This requirement for theoretical explicitness illuminates a key question that Gilead, Trope and Liberman appear to have missed: what is it about experiences of "mental travel," whether in time or across social relations, that identify them as such? What experientially distinguishes a memory from an imagined future? What distinguishes another's imagined thought from one's own? What *makes the distinction* between Gilead, Trope and Liberman's "ontologically" distinct representations?

In scale-free models, such distinctions can only be made by scale-dependent *inputs*, the sources of the experienced "epistemic feelings'' of reality, memory, and imagination that distinguish the functions of representations that may have the same "propositional" content. Hence identifying these inputs is a crucially important theoretical and experimental task. Considerable progress has been made in understanding how inputs from the body and the external world are combined to locate the experienced self in the here and now (Craig, 2010) and describing these processes within a predictive-coding framework (Seth and Tsakiris, 2018). The signals identifying memories, future projections, and thoughts of others are less well-characterized, though it is clear that specific activities in rostral PFC (Simons, Garrison and Johnson, 2017) and the insula – cingulate salience network (Uddin, 2015) are involved. The disruption of these signals in pathology and their potential for therapeutic modulation, e.g. with entheogens (Thomas, Malcolm and Lastra, 2017), give their mechanistic understanding clinical urgency. Such understanding cannot be accomplished if the distinctions they signal are simply taken as given.

References

Chater, N. (2018). *The Mind is Flat*. London: Allen Lane.

Craig, A. D. (2010). The sentient self. *Brain Structure and Function* 214, 563-577.

Fields, C. and Glazebrook, J. F. (2019). A mosaic of Chu spaces and Channel Theory II: Applications to object identification and mereological complexity. *Journal of Experimental and Theoretical Artificial Intelligence* 31, 237-265.

Fields, C. and Marcianò, A. (2019). Markov blankets are general physical interaction surfaces. *Physics of Life Reviews* in press. https://doi.org/10.1016/j.plrev.2019.08.004

Friston, K. (2013). Life as we know it. *Journal of the Royal Society Interface* 10, 20130475.

Friston, K., Levin, M., Sengupta, B. and Pezzulo, G. (2015). Knowing one's place: a free-energy approach to pattern regulation. *Journal of the Royal Society Interface* 12, 20141383.

Goguen, J. A. (1991). A categorical manifesto. *Mathematical Structures in Computer Science* 1, 49-67.

Hoffman, D. D. (2018). The interface theory of perception. In Serences, J. (Ed.) *The Stevens' Handbook of Experimental Psychology and Cognitive Neuroscience, Vol. II.* Hoboken, NJ: Wiley (Ch. 16).

Hoffman, D. D., Singh, M. and Prakash, C. (2015). The interface theory of perception. *Psychonomic Bulletin & Review* 22, 1480-1506.

Kuchling, F., Friston, K., Georgiev, G. and Levin, M. (2019). Morphogenesis as bayesian inference: A variational approach to pattern formation and control in complex biological systems. *Physics of Life Reviews* in press. https://doi.org/10.1016/j.plrev.2019.06.001

Seth, A. K. and Tsakiris, M. (2018). Being a beast machine: The somatic basis of selfhood. *Trends in Cognitive Sciences* 22, 969-981.

Simons, J. S., Garrison, J. R. and Johnson, M. K. (2017). Brain mechanisms of reality monitoring. *Trends in Cognitive Sciences* 21, 462-473.

Thomas, K., Malcolm, B. and Lastra, D. (2017). Psilocybin-assisted therapy: A review of a novel treatment for psychiatric disorders. *Journal of Psychoactive Drugs* 49, 446-455.

Uddin, L. Q. (2015). Salience processing and insular cortical function and dysfunction. *Nature Reviews Neuroscience* 16, 55-61.