An Overlooked Challenge for the "Emergent Multi-Scale Causality" Framework for Modelling Complex Systems. *A Brief Philosophical Note*

Proceedings for the Foundations of Computation Convention of the Society for the Study of Artificial Intelligence and Simulation of Behaviour, July 2026, Sussex

Oriol Roca-Martín[1]

Many of the current world challenges that scientists are trying to tackle are linked by a common thread: they imply the modelling, predicting and controlling of complex systems, which are characterized by having properties that emerge from the non-trivially ordered interactions of large numbers of elements at different scales. Exemplifying but some of the many significant challenging tasks or systems of this kind are the research on climate change and its consequences, effective procedures and policies for mitigating unfair economical trends or global tax evasion schemes, the spread of misinformation and pandemics, or attractor dynamics in gene regulatory networks and the brain.

The classic assumptions of linearity, locally defined dependencies, and reduction have brought us only so far in effectively modelling the world; scientists studying complex systems in many occasions have to work, instead, with non-linearity (e.g., power laws, and phase transitions), strong path-dependencies, and emergent processes at different scales (for examples of the previous contrasts [27, 20, 17, 15]). We are far from a complete formal and conceptual framework that can systematically guide the modelling, understanding, and harnessing of complex processes. Taking into account these meta-theoretical needs, this article focuses on emergence.

Recently, various research strands on this topic have converged towards a set of closely connected ideas that make use of information-theoretic and computational complexity principles in order to account for and formally model emergence and causality across different scales; what I'll call here the *emergent multi-scale causality* (eMSC) approaches. This framework is growing in popularity—in science [32] and in some prominent versions of ontic structural realism in philosophy [12]. However, the methodological and metaphysical considerations that underlie the modelling decisions of the framework remain unclear (as has been noted by numerous authors, e.g., [9, 8, 33, 3]). In the context of this concern, this article shows that the main strategic advantage of eMSC when it comes to defend scientific realism towards effective theories' ontologies—that is, being able to systematically define indispensable ontologies (discussed below)—faces some problems. The reasons for these problems are articulated, and possible natural extensions of the framework are briefly discussed.

Despite significant (and interesting) differences in specificities and intellectual origins (which we will have to ignore here), four pillars indicate a common characteristic thread shared by the approaches that constitute the eMSC framework:

P1. Emergence is thought of in terms of the capacity to construct effective theories from observations of natural systems (Fig.1a) (very similar ideas can be found in [20, 5, 13]). Given a task—solvable if a generalization (from $H1$ to Hi of Hn) robustly obtains with a fidelity down to some noise tolerance—, finding an effective theory means identifying the appropriate abstract or "coarse-grained" scale(s) containing variables and dynamics that robustly predict Hi, avoiding the costs (or the unsolvability) of tracking lower dynamics with more degrees of freedom. We can find tasks and effective theories that solve them "in the wild", in natural sciences—a classic example of a scientific task can be tracking the position of Neptune, we can do so by following its centre of mass (following H dynamics given a generalization) in contrast to having to follow all the degrees of freedom of the particles that compose the planet (L dynamics in a particle-based generalization from another theory). And also in engineering—a digital calculator is cleverly crafted in such a way that it is a system in which L dynamics are robustly securing some specific H dynamics according to some generalizations that solve the tasks of, for example, adding and subtracting with a required fidelity.

P2. In the eMSC framework emergent entities are conceived as patterns compressed from regularities in datasets, and a "more inductive bang for the buck" realist metaphysics of science is adopted towards patterns: the metaphysical force comes from noticing that one cannot just wish into existence neither facts about inductive success—a proposed effective ontology, can either reproduce the required dynamics according to some noise tolerance, or it cannot—nor facts about computational complexity—some successful effective ontologies are informationally more efficient than others (axis A and B in Fig.1b). The framework adopts this stance from computational mechanics' foundational way of reasoning about the topic [22], which is explicitly linked to Daniel Dennett's Real Patterns ideas [7, 12]. (There are characteristic epistemic/evolutionary derivations of the previous metaphyscial points, e.g., [16]).

P3. Just by following the previous criteria, however, a multiplicity of ontologies arises: indefinitely many successful compressions—each yielding its own ontology—can be obtained by systematically finding (even just marginally) different efficient ontologies for a given task. This is considered problematic for the realist aims of the framework, and raises the challenge of providing a scarcity criterion for decreasing ontological density [11, 14, 1].

P4. Faced with this situation, eMSC proposes an explicit scarcity criterion: indispensability. According to which the relevant effective taxonomy is the one that reproduces a generalization (with a fidelity down to some noise tolerance) at the lowest complexity costs (star in Fig.1b). These 4 pillars have been independently identified and discussed by authors within philosophy of science [14, 26] and scientists

[1] LOGOS, BIAP, Universit de Barcelona; email: o.rocamartin@ub.edu

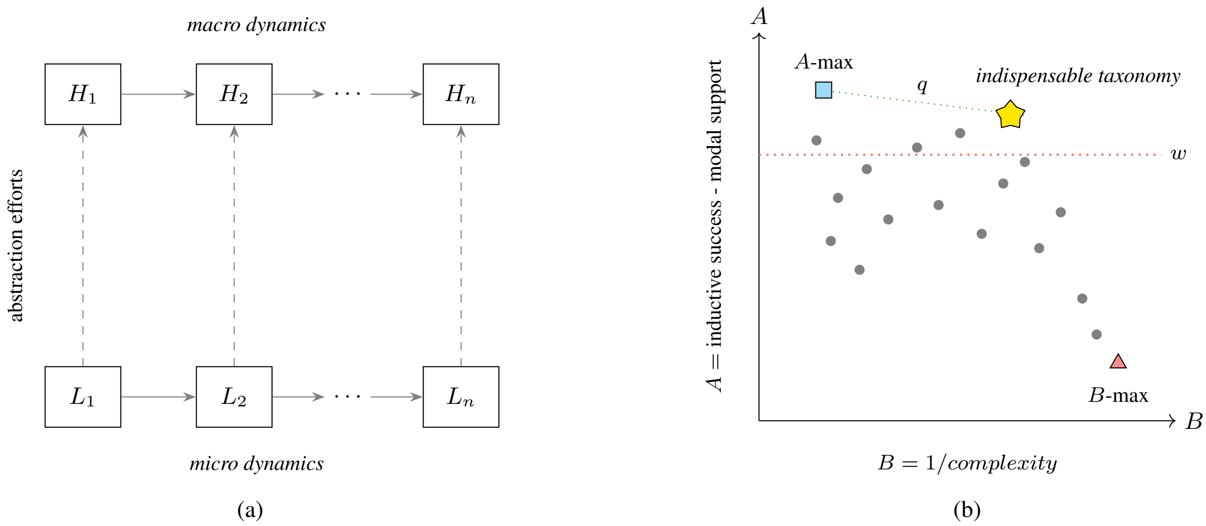

Figure 1: **(a)** Possible paths from microdynamics to a target macrovariable H_i of H_n, through abstraction and macro-/micro-dynamical evolution. **(b)** A task is solvable by different effective taxonomies (each dot) that generate the dynamics of generalizations with a fidelity down to some noise tolerance, here depicted by threshold w; it is a non-trivial, empirical matter whether those taxonomies exist, and how they are distributed. In this toy example, the dots show some trade-offs between inductive success/modal support (A) and inverse complexity (B). The A-max taxonomy (square) maximizes inductive success at relatively maximal complexity costs; the B-max taxonomy (triangle) is relatively minimally complex but inductively idle. The indispensable taxonomy (star) is the effective taxonomy within threshold w that has the lowest complexity; if the example task would be less demanding regarding noise tolerance, both w and, in accordance, the indispensable taxonomy, would change, and the dotted line q, unifying the indispensable taxonomy with A-max, would have a more pronounced negative slope.

modelling complex systems [22, 23, 31, 10].

But despite its clearly objective, mind-independent features (facts about inductive success and complex-theoretic optimality) indispensability is not easy to completely cash out for the methodological and metaphysical debate at hand. The issue is that while indispensability might solve the Local Challenge of providing a scarcity criterion that can choose ontologies for a given solvable task, there is another obvious source that threatens to bring back some form of multiplicity and relativism prima facie unwanted by the realist goals of the framework, the Global Challenge: there is a dense space from which to choose different possibly solvable tasks, each potentially yielding their own new effective taxonomies. Indispensability targets the Local Challenge for a given task, but in its current form cannot address the Global Challenge. Precisely because on many occasions there might be a trade-off between the Local and the Global Challenge. Imagine that by some fluke a particle p is exactly in the centre of mass of Neptune for the next 10 years; we can solve the task of tracking the position of Neptune for the next 10 years by discovering the effective theory "follow this exact particle p"; having an effective theory that implies calculating centres of mass is more complex for the task than simply tracking p; this formally results in particle p as the indispensable ontology; however, having centres of mass would clearly be inductively better (think about the 11th year of trying to track Neptune) and would generalize better (p is not present in other planets). Considerations about robustness in induction and extrapolation seem to have a weight for scientific realist considerations that is not captured by indispensability alone.

Shalizi and Moore wittily comment on what we call here the Global Challenge [23, p.17]:

> Our initial choice of observables is a function of our ability and inclinations to observe the system, and our choices about what to predict. Beyond that initial choice, and the requirement that good macrostates have certain causal properties, the causal states we use are completely out of our control, fixed entirely by the system's objective, microphysical dynamics. A different set of initial observables will, generally, lead to a different set of causal states. Sometimes, but not always, these causal states are related in a hierarchy of emergence. One might put it like this: for every question we ask It, Nature has a definite answer; but Nature has no preferred questions.

This situation gives rise to the following soft **dilemma**: eMSC approaches can **EITHER** (Horn 1) develop accounts that outline how much and what type of taxonomical multiplicity is acceptable while avoiding undermining the ontic and scientific status of effective theories, **OR** (Horn 2) try develop a solution or approximation to the open Global Challenge.

The strategies constitute a dilemma when they are defended in their stronger versions where one excludes the other; but softer versions of both positions can be developed by conceding points to the contrary strategy. Regarding Horn 1, it is still philosophically underexplored whether the realist goals of eMSC are compatible with a form of multiplicity that would be satisfactory for the framework. This is an open question, which is being indirectly and partially explored in different but related contexts (e.g., in discussions over the compatibility of scientific realism with the abandonment of the natural/non-natural kinds distinction: [24, 6]). Regarding Horn 2, the philosophers Ladyman [12] and Ross [14] have indirectly pointed to a way to have a further scarcity criterion, that seems to target the worries of the Global Challenge: we should aim at effective taxonomies that can be part of a systematic process of scientific unification, and coherent triangulation between the ontic commitments of different domains. This move, however, has not been clearly developed anywhere in the philosophical literature (this shortcoming has been highlighted by [3, 18]). Other intellectually related projects, such as William Wimsatt's "Robustness Analysis" [29], are also underexplored in the sense we care about here.

After providing this framing that underlines the Global Challenge, this article asks: while its philosophical side remains only timidly explored, has the scientific side of eMSC addressed it with more success? The short answer is: no, but there are tools from the broader computational modelling field that directly fit the Challenge. The different scientific modelling proposals within the eMSC frameworkñ try to provide objective criteria for choosing taxonomies at some scales over others by analysing the dynamics of interaction between the different possible scales of specific systems. As a result, interesting compositional properties between scales have been defined (e.g., causal emergence, effective information, synergistic information, and computational closure), and different empirical techniques have been proposed [32, 11, 20]. However, while formal developments around scales and mereology might be a crucial part of the necessary work for the eMSC project to be successful, a focus on isolated systems does not address the dilemma. The dilemma can be addressed by exploring the possibilities and limits of Horn 1, and/or by evaluating and developing concepts and tools that potentially tackle Horn 2 by exploring when and why patterns that are found in very distinct systems at different scales are modally equivalent. On this latter point, eMSC's framework could integrate active discussions regarding multiple realizability [19], clustering and machine learning methods [28], as well as analyses of the limits of induction [21, 30, 25] and across-system extrapolation (see [2]; or, especially relevant for eMSC, the latest extensions of computational mechanics [4]).

In sum, the point raised by this article is that the goals of eMSC approaches may call for complementary work that currently receives insufficient attention. We might discover that there are principled reasons why the Global Challenge is unsolvable, either in general or for particular fields. This would imply that there is no scientifically principled justification for their effective taxonomies, and—importantly—that extra-theoretical, practical considerations would be needed to carry those scientific inquiries forward.

ACKNOWLEDGEMENTS

I would like to thank the BM group, and especially my supervisor Manolo Martínez, for stimulating discussions on this topic. I am also grateful to James Ladyman, David Wolpert, and Abel Jansma for their welcoming engagement with my questions. This research was made possible by a BIAP Maria de Maeztu FPI grant (CEX2021-001169-M-20-1).

REFERENCES

[1] Holly K. Andersen, 'The density of structure', *Synthese*, **206**(5), 237, (2025).

[2] Shai Ben-David, John Blitzer, Koby Crammer, Alex Kulesza, Fernando Pereira, and Jennifer Wortman Vaughan, 'A theory of learning from different domains', *Machine Learning*, **79**(1–2), 151–175, (2010).

[3] Majid Davoody Beni, 'Structural realism, metaphysical unification, and the ontology and epistemology of patterns', *International Studies in the Philosophy of Science*, **31**(3), 285–300, (2017).

[4] Nicolas Brodu and James P. Crutchfield, 'Discovering causal structure with reproducing-kernel hilbert space ϵ-machines', *Chaos: An Interdisciplinary Journal of Nonlinear Science*, **32**(2), (2022).

[5] Sean M. Carroll and Achyuth Parola, 'What emergence can possibly mean', in *Dennett's Real Patterns in Science and Nature*, eds., Tyler Millhouse, Steve Petersen, and Don Ross, MIT Press, Cambridge, MA, (2026).

[6] Anjan Chakravartty, 'Last chance saloons for natural kind realism', *American Philosophical Quarterly*, **60**(1), 63–81, (2023).

[7] Daniel C. Dennett, 'Real patterns', *The Journal of Philosophy*, **88**(1), 27–51, (1991).

[8] Joe Dewhurst, 'Causal emergence from effective information: Neither causal nor emergent?', *Thought: A Journal of Philosophy*, **10**(3), 158–168, (2021).

[9] Frederick Eberhardt and Lin Lin Lee, 'Causal emergence: When distortions in a map obscure the territory', *Philosophies*, **7**(2), (2022).

[10] Peter Grassberger. Some comments on computational mechanics, complexity measures, and all that, 2018.

[11] Abel Jansma and Erik Hoel. Engineering emergence. arXiv preprint, 2025.

[12] James Ladyman, 'Patterns all the way up: Prolegomena to a future naturalized metaphysics', in *Dennett's Real Patterns in Science and Nature*, eds., Tyler Millhouse, Steve Petersen, and Don Ross, chapter 3, MIT Press, Cambridge, MA, (2026).

[13] James Ladyman and Lorenzo Lorenzetti, 'Effective ontic structural realism', *The British Journal for the Philosophy of Science*, (2023). Online ahead of print; forthcoming in vol. 78.

[14] James Ladyman and Don Ross, *Every Thing Must Go: Metaphysics Naturalized*, Oxford University Press, Oxford, 2007.

[15] James Ladyman and Karoline Wiesner, *What Is a Complex System?*, Yale University Press, New Haven, CT, 2020.

[16] Sarah E. Marzen and Simon DeDeo, 'The evolution of lossy compression', *Journal of The Royal Society Interface*, **14**(130), 20170166, (2017).

[17] Patricia Palacios, *Emergence and Reduction in Physics*, Cambridge Elements in Physics, Cambridge University Press, Cambridge, 2022.

[18] Oriol Roca-Martín, 'A New Problem for Real Patterns Ontic Structural Realism?—Separating Local from Global Indispensability', *Manuscript in preparation for submission*, (2026).

[19] Oriol Roca-Martín and Manolo Martínez, 'What Can Multiple Realizability Possibly Be? A Complexity Approach', *Manuscript in preparation for submission*, (2026).

[20] Fernando E. Rosas, Bernhard C. Geiger, Andrea I. Luppi, Anil K. Seth, Daniel Polani, Michael Gastpar, and Pedro A. M. Mediano. Software in the natural world: A computational approach to hierarchical emergence. arXiv preprint, 2024.

[21] Gerhard Schurz, *Hume's Problem Solved: The Optimality of Meta-Induction*, MIT Press, Cambridge, MA, 2019.

[22] Cosma Rohilla Shalizi and James P. Crutchfield, 'Computational mechanics: Pattern and prediction, structure and simplicity', *Journal of Statistical Physics*, **104**(3–4), 817–879, (2001).

[23] Cosma Rohilla Shalizi and Cristopher Moore, 'What is a macrostate? subjective observations and objective dynamics', *Foundations of Physics*, **55**(1), 2, (2025).

[24] Matthew H. Slater, 'Less work for theories of natural kinds', *The Philosophical Quarterly*, (2024).

[25] Tom F. Sterkenburg and Peter D. Grünwald, 'The no-free-lunch theorems of supervised learning', *Synthese*, **199**(3–4), 9979–10015, (2021).

[26] Abel Suñé and Manolo Martínez, 'Real patterns and indispensability', *Synthese*, **198**(5), 4315–4330, (2021).

[27] Thomas F. Varley, 'Information theory for complex systems scientists: What, why, and how', *Physics Reports*, **1148**, 1–55, (2025).

[28] David S. Watson, 'On the philosophy of unsupervised learning', *Philosophy & Technology*, **36**(2), 28, (2023).

[29] William C. Wimsatt, 'Robustness, reliability, and overdetermination', in *Scientific Inquiry and the Social Sciences*, eds., Marilynn B. Brewer and Barry E. Collins, 124–163, Jossey-Bass, San Francisco, (1981).

[30] David H. Wolpert, 'The implications of the no-free-lunch theorems for meta-induction', *Journal for General Philosophy of Science*, **54**(3), 421–432, (2023).

[31] David H Wolpert, Joshua A Grochow, Eric Libby, and Simon DeDeo, 'Optimal high-level descriptions of dynamical systems', *arXiv e-prints*, arXiv–1409, (2014).

[32] Bing Yuan, Zhang Jiang, Aobo Lyu, Jiayun Wu, Zhipeng Wang, Mingzhe Yang, Kaiwei Liu, Muyun Mou, and Peng Cui. Emergence and causality in complex systems: A survey on causal emergence and related quantitative studies. arXiv preprint, 2024.

[33] J. Zhang, R. Tao, K. H. Leong, et al., 'Dynamical reversibility and a new theory of causal emergence based on svd', *npj Complexity*, **2**, 3, (2025).