






Taking Cognition Seriously

A Generalised Physics of Cognition

Sophie Alyx Taylor¹(✉) , Son Cao Tran¹ , and Dan V. Nicolau Jr.^{1,2,3} 

¹ School of Mathematics, Queensland University of Technology, Brisbane, Australia

sophie.taylor@hdr.qut.edu.au

² Respiratory Medicine Unit, Nuffield Department of Medicine, University of Oxford, Oxford, UK

³ Centre for Clinical Research, University of Queensland, Brisbane, Australia

Abstract. The study of complex systems through the lens of category theory consistently proves to be a powerful approach. We propose that cognition deserves the same category-theoretic treatment. We show that by considering a highly-compact cognitive system, there are fundamental physical trade-offs resulting in a utility problem. We then examine how to do this systematically, and propose some requirements for “cognitive categories”, before investigating the phenomena of topological defects in gauge fields over conceptual spaces.

Keywords: Mathematics of cognition · Topological field theories · Applied category theory · Computational cognitive architecture · Computation in the Schwarzschild metric

1 Introduction

The study of thought has occupied academia for millennia, from philosophy, to medicine, and everything in between. We demonstrate a promising approach to take advantage of the powerful tools developed by other disciplines in our study of cognition, by emphasising the role that category theory – a foundation of mathematics – ought to play.

We first motivate our approach by considering the effects of general relativity on a highly compact cognitive agent, and noting the structural similarities between metric field theories in physics, and the dynamics of mental content in conceptual spaces. We feel that these similarities practically beg for the study of field theories in *conceptual spaces themselves*, and consider the topological defects that arise; interpreting them, as is done in physics, as *particles of thought*.

2 Background

We first recall some required background which we build upon.

2.1 Cognitive Architecture and Conceptual Spaces

To model a dynamical system, we require specification of two things: How the data is to be represented, and the system that operates on it. In our approach, we use conceptual spaces and cognitive architectures, respectively.

Conceptual Spaces. *Conceptual spaces* [7,22] provide a convenient middle-ground approach to knowledge representation, situated between the symbolic and associationist approaches in terms of abstraction and sparseness. A conceptual space is a tensor product of conceptual *domains*, which can be arbitrary topological spaces equipped with certain measures. They provide a low-dimensional alternative to vector space embeddings, which often have no clear interpretation of their topological properties [3]. Conceptual spaces represent concepts as regions in a topological space. Because (meta-)properties of the conceptual spaces depend on context, this suggests that a proper description involves the notion of a (pre)sheaf. Conveniently, they also provide a well-motivated, natural construction for prior probabilities based on the specific symmetries of the domain [5].

Cognitive Architecture. Cognitive architectures are an system-level integration of many different aspects of “intelligence”. It can be useful to think of cognitive architectures as the cognitive analogue of a circuit diagram; specifying the hardware/operating system (in the case of robotic agents) or the wetware (in the case of biological agents) that the processes deliberate thought, intentions and reasoned behaviour occur. They generally take a “mechanism, not policy” approach: just like a mathematician has the same brain structure as a carpenter, a given cognitive architecture can produce wide range of agents, depending on what has been learned.

Cognitive architecture has been studied in the context of Artificial Intelligence for at least fifty years, with several hundred architectures described in the literature. Laird [12] describes the SOAR cognitive architecture, one of the most well-established examples in the field. Several surveys have been presented, such as Goertzel, Pennachin, and Geisweiller [8] and Kotseruba and Tsotsos [11]. Huntsberger [9] describes one particular architecture designed for human-machine interaction in space exploration.

Generalised Problem Space Conceptual Model. As a concrete example of a cognitive architecture, we show a modified version of the Problem Space Conceptual Model [12] in Fig. 1. It is reminiscent of a multi-head, multi-tape Turing machine; but with transition rules stored in memory, too. The core idea of the PSCM is that operators are selected on the basis of the contents of working memory, which is essentially short-term memory. Production rules elaborate the contents of working memory until quiescence; operators are essentially in-band signals about which production rules should be fired. If progress can’t be made, an *impasse* is generated; from here, we can enter a new *problem space*. This is essentially a formalisation of a “point of view” to solve a problem.




Fig. 1. A slight generalisation of the Problem Space Conceptual Model. Control lines are drawn in orange, while data flow is depicted in blue. Rounded corners on a black indicate that it while it has fixed functionality, it exposes control knobs to working memory, just like Direct Memory Access in a computer. (Color figure online)

Realisation of Cognitive Architecture. Cognitive architectures are *abstract models*, which need to be *concretely realised*. We almost always study cognition through a realisation via neuroscience, which itself is realised via biology, then chemistry, then physics. This is depicted in Fig. 2. Alternatively, we may realise it in the form of robotics and software simulations; like with neuroscience, these are also built upon a tower of realisations.

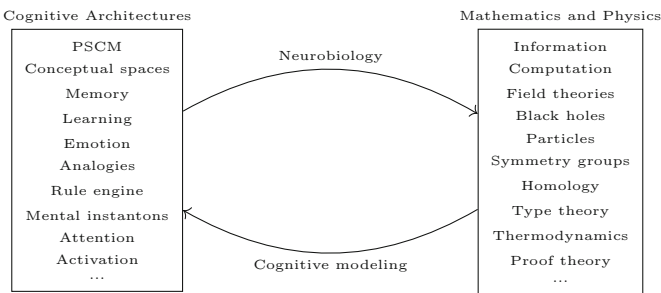


Fig. 2. The abstract models of cognitive architectures are typically concretely realised by way of neurobiology.

2.2 Limits on Computation

Computation has its limits; most famously, the halting problem. We can find limits of computation in at least two general classes: limits imposed by computation in-and-of-itself, and limits imposed by physics on the embodiment of a computation.

Busy-Beaver Machines. Limits on maximum output of a Turing machine for a given number of states, and number of transitions, $BB(n)$. The Busy Beaver game is to construct an n -state Turing machine which outputs the largest amount of “1”s possible, *then halts*. A variation on this is to count the number of steps it takes for the said Busy Beaver machine to halt. If you look at this as a function in n , whether counting the size of the output or the transitions, you have an extraordinarily fast-growing function, quickly outpacing Ackermann’s function. You can use this to decide whether n -state Turing machines halt, simply by waiting $BB(n)$ steps. If it takes longer than that, then you know it will never halt, by definition, because it would be the new $BB(n)$. Of course, $BB(n)$ itself is uncomputable in general: We have no free lunch.

The Bekenstein Bound. The Bekenstein bound is a limit on information density:

$$I \leq \frac{2\pi RE}{\hbar c \ln 2} \quad \text{bit} \quad (1)$$

where R is the radius of the system, E is the energy of the system, \hbar is the reduced Planck’s constant, and c the speed of light [21]. If one assumes the laws of thermodynamics, then all of general relativity is derivable from Eq. 1 [10]. If you fix the radius R , then as you try to fit more and more information in the given volume, you must eventually start to increase the energy in that volume. Eventually, the Schwarzschild radius of that energy equals R , and you have a black hole. This occurs when Eq. 1 is an equality.

The Margolus-Levitin Bound. The Margolus-Levitin bound is a bound on the processing rate of a system. A (classical) computation must take at least $\frac{\hbar}{4E}$. To decrease the minimum time to take a computational step, you must increase the energy involved. If this energy is contained in a finite volume, then you suffer the same problem arising from the Bekenstein bound; eventually the Schwarzschild radius catches up to you. It’s important to note, however, that this assumes no access to quantum memory.

2.3 Gauge Fields and Their Symmetries

A *gauge field* is an approach to formalising what we mean by a *physical field*. We think of a gauge field as data “sitting above” the points in a space (a *principle bundle*), as shown in Fig. 3, combined with a transformation rule (a *connection*)

to ensure the data transforms sensibly with a change of test point.[2] Gauge fields typically have a very useful property: We can reparameterise them to be more convenient to work with. For example, we can change the electric and magnetic potentials in an electromagnetic field by a consistent transformation derived from arbitrary data, yet the exact same physical effects will occur. This is called *gauge fixing*.




Fig. 3. A gauge field F over a spacetime Σ .

2.4 Topological Defects

Topological defects are an important notion in a number of scientific and engineering fields. Intuitively speaking, they are *imperfections* in what we call an *ordered medium*; roughly, something which has a consistent rule for assigning values to a point in space, like a gauge field [15]. Some familiar examples include:

- Grain boundaries in metals, consisting of disruptions in the regular lattice placement of the constituent atoms,
- Singularities and vortices in vector fields, such as black holes and tornadoes, respectively, and
- Gimbal lock is a topological defect in the Lie algebra bundle $\mathfrak{spin}(3) \rightarrow T^3$ of rotations over the Euler angles.

In general, defects arise when there are non-trivial homotopy groups in the principle bundle of the gauge field; that is, when you can make a loop in the bundle space, but you can't shrink it down to a single point.

These are non-perturbative phenomena, arising due to the very *nature* of these systems. These manifest themselves as (*psuedo-*)*particles* or higher branes, such as membranes. However, due to the risk of confusion between the homonyms “brain” and “brane”, we will call them all particles, regardless of their dimension.

Dynamics. What makes defects particularly interesting is that they can evolve over time. Dynamical behaviour in the base space can lead to dynamical behaviour of defects; for example, the movement of holes in a semiconductor, or the orbit of a small black hole around a larger one.

Temperature and Spontaneous Symmetry Breaking. Another common phenomenon is that there is a *temperature* associated with the defect, which controls both its dynamics and its very existence. The obvious example of a temperature parameter is thermal temperature; if you melt a bar of iron, you won't have any more crystal structure to be defective. A more challenging one is temperature in superconductors or superfluids; raise it too high and Cooper pairing no longer works, so you lose superconductivity and superfluidity, respectively. These drastic changes in behaviour are *phase changes*.

Consider the opposite situation: Freezing a liquid to a solid. If you lower the temperature to a point where defects start appearing, then those defects have to actually appear *somewhere*. The configuration of the defects will be just one possible configuration out of a potentially astronomically large space of possibilities (such as the precise atoms involved in a grain boundary), chosen at random. This phenomena is called *spontaneous symmetry breaking*, and is responsible for the masses of elementary particles, via the Higgs mechanism.

2.5 Category Theory and Its Slogans

Category theory is a foundation of mathematics, as an alternative to set theories like ZFC. Instead of focusing on membership, category theory focuses on *transformation* and *compositionality*. Its methodology for studying an object is to look at how that object relates to other objects. There are many excellent introductions to category theory aimed at different audiences (such as [6, 14, 16, 20]). These introduce the basic notions and show the versatility of category theory in a wide range of fields.

Categories and Functors. A category is a collection of objects, and a collection of morphisms between those objects. For example, the category of sets has sets as objects, and functions as morphisms. Different classes of category allow different constructions; for example, monoidal categories allow *pairing* of objects. You can translate between categories with *functors*, which are themselves just morphisms in the category of categories.

Adjunctions. A common situation in mathematics arises when you want to translate back and forth between two types of structures via a pair of functors, but the structures aren't isomorphic. The best we can do is approximate an isomorphism; this is called an *adjunction* [16].

Enrichment. *Enriched categories* are a generalisation of plain categories. Whereas in standard category theory, the morphisms between objects are described by their homset, a category enriched in a monoidal category \mathcal{V} has hom-objects from $\mathbf{Ob}(\mathcal{V})$ [6, pp. 56,139]. Ordinary categories are just categories enriched over the category of sets [6, p. 87].

An Approach to Science and Mathematics. Let us take a step back for a moment, and consider *how* science should be done. Bill Lawvere offers this:

When the main contradictions of a thing have been found, the scientific procedure is to summarize them in slogans which one then constantly uses as an ideological weapon for the further development and transformation of the thing [13].

Let us take this approach. What are some slogans for category theory?

- Better to have a nice category of objects than a category of nice objects [4]
- Dialectical contradictions are adjunctions [13]
- Find and characterise the principal adjunction of any problem

We shall now provide some motivation for our work.

3 Motivation

We motivate our approach by considering arbitrary cognitive systems in terms of physics. First, we consider the consequences of a physically dense cognitive agent; second, we consider the role self-parameterisation in conceptual spaces.

3.1 The High-Density Regime of Cognition and Its Consequences

Let us consider the high-density regime of cognition; that is, where the cognitive agent's radius is close to its Schwarzschild radius. We shall assume that the agent can learn from its experiences. As the agent experiences the world, it may learn new information,¹ which requires storage in the agent's memory. Consider this increase in information in the context of Eq. 1: more and more energy will be required to store it in a given volume. This, in turn, results in an increase in the Schwarzschild radius. As we are already close to our Schwarzschild radius, and we presumably don't want to collapse into a black hole,² we must increase our containment radius; this results in both the information storage *and* processing gadgets spreading out. This results in the average distance between two pieces of information in the containment volume increasing, which in turn results in increased average propagation delay to shunt information around. Thus, the average serial processing rate will decrease—in other words, *learning can make you slower*. This isn't even taking into account gravitational time dilation, which will only further enhance the slowdown relative to an outside observer, as well as the energy required to interact with the external environment.

¹ These experiences may allow generalisation of existing information and thus discarding of the individual cases; but over time, there will be a point where generalisations can't be reasonably made.

² It is ill-advised to become a black hole, if only to be able to affect the external environment in a reasonable manner (but see [1]).

Setting. We consider an agent comprised of an incompressible fluid which is spherically symmetric and non-rotating, and electrically neutral, resulting in the *interior Schwarzschild metric* [19]. In Schwarzschild coordinates, the line element is given by

$$c^2 d\tau^2 = \frac{1}{4} \left(3\sqrt{1 - \frac{R_S}{R_g}} - \sqrt{1 - \frac{r^2 R_S}{R_g^3}} \right)^2 c^2 dt^2 - \left(1 - \frac{r^2 R_S}{R_g^3} \right)^{-1} dr^2 - r^2 (d\theta^2 + \sin^2 \theta d\varphi^2), \quad (2)$$

where

- (t, r, θ, ϕ) = The spacetime coordinates in $(+ - - -)$ convention
- R_S = The Schwarzschild radius of the agent
- R_g = The value of r at the boundary of the agent, measured from the interior.

Slowdown Due to Propagation Delay. If we consider an instantaneous path through our agent – that is, one where $dt = 0$ – we can see from Eq. 2 that when R_g is close to R_S , the radial term blows up as the path approaches the surface; indeed, we can see that the distance required to be travelled to travel to the surface when $R_g = R_S$ is infinite.

Slowdown Due to Gravitational Time Dilation. The time dilation experienced by the agent compared to an observer at infinity increases with density. The scaling compared to the outside observer is given by

$$\sqrt{1 - \frac{R_S}{r}} \quad (3)$$

valid for $r > R_S$.

Extra Energy Required to Interact with the External Environment.

Not only will the extreme gravitational field reduce rate of processing, it will also increase the energy required to interact with the external environment. Consider the case of our agent trying to communicate with a distant observer via some protocol based on the exchange of radio signals. If there is any specification of frequency of the carrier signal, then we must account for the gravitational redshifting of our agent’s emissions. This redshifting results in longer wavelength signals; equivalently, lower frequency signals. Because the energy of the photons comprising the signal is directly proportional to its frequency, this can be stated in terms of *lower energy signals*. Thus, in order to ensure outbound transmissions meet the frequency specifications of the protocol, the photons must be emitted with extra energy to overcome the extreme curvature of the gravitational field.

In the simplest case, assume the exterior Schwarzschild metric [18]. To signal an outside observer at infinity with a photon of a given energy E emitted at a

radius R_E , the photon will have to be emitted with an energy increased by a factor of

$$\frac{R_E}{\sqrt{R_E(R_E - R_S)}} - 1. \quad (4)$$

Some Tasks Are Too Complex to Be Solved in a Given Volume. We have thus seen that there is a *fundamental physical trade-off* between learning, and the time it takes to solve an arbitrary task. This immediately implies a number of things:

- There is a trade-off between the general capability of an agent, and its average reactivity,
- Some complex tasks with a time requirement are simply too complicated for any agent to solve the task in a given volume of space,
- As you increase the average capability of a group of agents, less of them will be able to fit in a given volume.

Lattice of Cognitive Skills. We can use these limits to create various partial orders:

- We can rank tasks by the minimum and maximum volume of agents capable of solving it,
- We can rank agents by their capabilities with respect to physically-embodied Busy-Beaver machines,
- We can rank space-time volumes by combinations of the above

3.2 Self-parameterisation of Behaviour

Let us consider our second motivational example: the self-parameterisation of behaviour. The behaviour of an agent depends on its accumulated knowledge up until that point, in particular, its procedural knowledge. Its knowledge, in turn, depends on its past behaviour acquiring the knowledge. This self-referential quality appears in physics, in the form of *metric field theories*. Consider, again, general relativity: The metric tensor controls the dynamics of a system, and the distribution of that system determines the metric tensor. The previous example considered the *effects* of general relativity on a cognitive agent, but the agent didn't particularly play an *active* role in our considerations.

But, why consider only general relativity? Can we consider other metric field theories? For that matter, do we even have to consider *physical* field theories?

4 Goal

Consider the position our motivational examples has put us in: Instead of considering cognition as an embodied agent comprised of wetware or hardware, we just re-enacted the old “assume a spherical cow in a vacuum” joke. We didn't bother

trying to determine *how* a dense sphere of incompressible fluid could embody a cognitive agent; it turned out we didn't even need to! We just assumed that the translation from an abstract model of cognition, to the concrete sphere of cosmic horror, and then back again to an abstract model preserved the semantics of cognition, whatever they may be.

4.1 Taking Advantage of Tools from Other Disciplines

Let us now state what we wish to do, in order to find a solution. We wish to use the mathematical tools from other disciplines in order to study cognition. How is this typically done?

The Physics of Cognition, by Way of Biology. As shown in Fig. 2, we translate through a chain of “domains” of science before reaching physics. Each translation introduces an extraordinary amount of complexities that exist solely due to choosing a specific concrete realisation *at each stage*. While, of course, neurology and biology are *incredibly* important subjects of study in humans – if only for their phenomenological observations, not to mention the pathophysiological importance – they are complicating factors in the study of the phenomena of cognition in-and-of-itself.

We Only Need to Find Nicer Realisation Morphisms Which Preserve Behaviour. Instead of considering the familiar realisation via neurobiology, the realisation transformation only needs to preserve the structures and behaviours of interest; that is, *we only need to find toy models*. To do this, we need to find a nice class of “cognitive categories”, and adjunctions out of them, as shown in Fig. 4. A slightly more category-theoretic diagram is shown in Fig. 5.

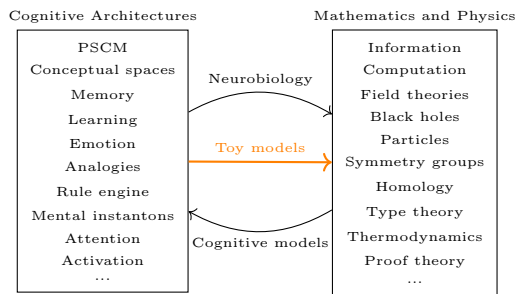


Fig. 4. All we need to do is to preserve the behaviours, not any particular concrete realisation strategy.




Fig. 5. Transformation expressed in a more traditional categorical diagram

4.2 Cognitive Categories

If we are going to use category theory to study cognition, then we ought to specify what cognitive categories actually are. This is an open problem, but some plausible requirements include:

- Subcategories of cognitive categories ought to include Turing categories, in order to capture computational behaviour.
- There should be an opportunity for enrichment in a category of conceptual spaces.

5 Example: Topological Defects in Conceptual Spaces

Let us consider the metric field analogy a little deeper. Since many conceptual spaces of interest have a genuine geometric structure, and we can form fibre bundles over them which allow parallel transport, then we can at least *consider* gauge fields.

5.1 Cognitive Gauge Fields

There are a lot of potential options for cognitive gauge fields. Some are generic, which might apply to any conceptual space; whereas some might only apply to certain classes of conceptual space. We assume some mechanism for smooth interpolation in cases of noisy discrete data.

Some example plausible generic gauge fields include:

- The activation value of a specific memory is related to its probability of being recalled due to a query. The higher the activation value, the more it will be recalled. These values can spread to adjacent memories. Many “forgetting” mechanisms are based on forgetting memories with low activation value.
- We can consider the emotional state of an agent when storing the memory, decomposed into a set of *valuation* and *valence* affect dimensions.
- We can also consider the subjective importance of a memory, which is separate to its activation; an example in procedural memory is a rule saying to not completely flood a room with water if there are people in it. It’s not very *likely* to be subject to recall, but it’s certainly an important rule!

Some plausible non-general gauge fields might be:

- Trustworthiness of data gathered socially.
- Difficulty of tasks and behaviours. This can evolve based on experience and better understanding of the situation.

5.2 Defects and Their Dynamics: Particles of Thought

Some cognitive gauge fields will have non-trivial topology, resulting in topological defects. As the underlying conceptual spaces evolve, whether contents or the topology itself, we might observe dynamical evolution of these defects. Thus, we might (not-so-metaphorically) call these “particles of thought”.

Leaving aside whether such a thing has a meaningful interpretation,³ we can at least ask more about the nature of such things.

5.3 Production Rules as Potentials

Can we encode production rules as a something akin to a potential? If so, what actually generates those potentials?

5.4 Transmission of Influence and Cognitive Gauge Bosons

How is the influence of a gauge field propagated? Is it wave-like, as in classical physics, or are there ‘force-carrying particles’, like gauge bosons in particle physics?

5.5 Phase Changes

Are there phase changes? That is, is there some order parameter where the particles are only manifest in a given range of parameter values? Does this relate to switching problem spaces?

5.6 Cognitive Event Horizons

Are there “cognitive event horizons”, where there are boundaries from which the effects of a particle can never escape, not just effects on its surrounding particles or memories, but on behaviour too? If so, how do they form? How do they evolve? Are there mechanisms to affect the underlying topological structure of a conceptual space? Is there a “maximum resolution” to some conceptual spaces as a result, analogous to the Schwarzschild radius in general relativity?

³ This almost certainly depends on the conceptual space and gauge field involved.

6 Discussion

There is an important consideration to be made when talking about theoretical modelling: we must stress when we are only talking about *effective theories*; that is, theories which model the *effects* of cognition, but do not make any claims as to whether there is any actual *causal* connection with the reality of cognition. The ontological status of particles of thought certainly rests on the status of conceptual spaces, and gauge fields over them. Further, assuming they *are* ontologically valid, whether they actually have any causal role requires considerable further study.

6.1 Future Research

We have a number of interesting questions prompted by our approach:

- How might we study the flow of information throughout an agent’s lifetime? Can this be linked to “particles of thought”?
- How can an agent perceive the Self? Does an agent who is aware of itself encounter, for example, the Barber’s paradox? How can it reason about things which are not true? What are the connections with paraconsistent logic?
- How can we more fruitfully take advantage of topological data analysis?
- How does analogy work in our approach? Does it rely on the homology structure of the relevant conceptual spaces having particular forms?
- What is a good model for various learning mechanisms; both of mental content itself, and of new conceptual spaces?
- What group structures over different conceptual spaces can we find to yield different field theories?
- Behaviour and cognition depends heavily on emotional state [17]. What is the most appropriate conceptual space to represent this, and do they have any psuedoparticles?

References

1. Andrews, G.: Black hole as a model of computation. *Results Phys.* **13**, 102–188 (2019). <https://doi.org/10.1016/j.rinp.2019.102188>
2. Baez, J.C., Muniain, J.P.: *Gauge Fields. Knots and Gravity*. World Scientific Publishing Company, Singapore (1994)
3. Bolt, J., Coecke, B., Genovese, F., Lewis, M., Marsden, D., Piedeleu, R.: *Interacting Conceptual Spaces I: Grammatical Composition of Concepts* (2017). [arXiv: 1703.08314](https://arxiv.org/abs/1703.08314), 29 September 2017
4. Corfield, D.: *Modal Homotopy Type Theory: The Prospect of a New Logic for Philosophy*. University Press, Oxford (2020)
5. Decock, L., Douven, I., Sznajder, M.: A geometric principle of indifference. *J. Appl. Logic* **19**, 54–70 (2016)
6. Fong, B., Spivak, D.I.: *An Invitation to Applied Category Theory: Seven Sketches in Compositionality*. Cambridge University Press, Cambridge (2019). <https://doi.org/10.1017/9781108668804>

7. Gärdenfors, P.: *Conceptual Spaces: The Geometry of Thought*. MIT Press, Cambridge (2004)
8. Goertzel, B., Pennachin, C., Geisweiller, N.: Brief survey of cognitive architectures. In: *Engineering General Intelligence, Part 1. ATM*, vol. 5, pp. 101–142. Atlantis Press, Paris (2014). https://doi.org/10.2991/978-94-6239-027-0_6
9. Huntsberger, T.: Cognitive architecture for mixed human-machine team interactions for space exploration. In: *2011 Aerospace Conference*, pp. 1–11 (2011). [10/cq4m28](https://doi.org/10.1109/AERO.2011.5945288)
10. Jacobson, T.: Thermodynamics of spacetime: the Einstein equation of state. *Phys. Rev. Lett.* **75**(7), 1260–1263 (1995). <https://doi.org/10.1103/PhysRevLett.75.1260>. [arXiv: gr-qc/9504004](https://arxiv.org/abs/gr-qc/9504004)
11. Kotseruba, I., Tsotsos, J.K.: *A Review of 40 Years of Cognitive Architecture Research: Core Cognitive Abilities and Practical Applications* (2016). [arXiv: 1610.08602](https://arxiv.org/abs/1610.08602), 26 October 2016
12. Laird, J.E.: *The Soar Cognitive Architecture*. MIT Press, Cambridge (2012)
13. Lawvere, F.W.: Categories of space and of quantity. In: Echeverria, J., Ibarra, A., Mormann de Gruyter, T., (eds.) *The Space of Mathematics*. Boston, Berlin (1992). <https://doi.org/10.1515/9783110870299.14>
14. Lawvere, F.W., Schanuel, S.H.: *Conceptual Mathematics: A First Introduction to Categories*. Cambridge University Press, Cambridge (2012)
15. Mermin, N.D.: The topological theory of defects in ordered media. *Rev. Mod. Phys.* **51**, 591–648 (1979). <https://doi.org/10.1103/RevModPhys.51.591>
16. Riehl, E.: *Category Theory in Context*. Dover, New York (2016)
17. Rosales, J.-H., Rodríguez, L.-F., Ramos, F.: A general theoretical framework for the design of artificial emotion systems in autonomous agents. *Cogn. Syst. Res.* **58**, 324–341 (2019). [10/ggh7sr](https://doi.org/10.1016/j.cogsr.2019.07.007)
18. Schwarzschild, K.: On the gravitational field of a mass point according to Einstein – slashhcs theory. *Sitzungsber. Preuss. Akad. Wiss. Berlin (Math. Phys.)* **1916**, 189–196 (1916). [arXiv: physics/9905030](https://arxiv.org/abs/physics/9905030)
19. Schwarzschild, K.: On the gravitational field of a sphere of incompressible fluid according to Einstein’s theory. *Sitzungsber. Preuss. Akad. Wiss. Berlin (Math. Phys.)* **1916**, 424–434 (1916). [arXiv: physics/9912033](https://arxiv.org/abs/physics/9912033)
20. Spivak, D.I.: *Category Theory for the Sciences*. MIT Press, Cambridge (2014)
21. Tipler, F.J.: The structure of the world from pure numbers. *Rep. Prog. Phys.* **68**(4), 897–964 (2005). <https://doi.org/10.1088/0034-4885/68/4/R04>. [arXiv: 0704.3276](https://arxiv.org/abs/0704.3276)
22. Zenker, F., Gärdenfors, P. (eds.): *Applications of Conceptual Spaces*. SL, vol. 359. Springer, Cham (2015). <https://doi.org/10.1007/978-3-319-15021-5>