Logic and Logical Philosophy
Online First Articles (2025)
DOI: 10.12775/LLP.2025.013

Simon McGregor

Embodied Sensorimotor (Hyper)intensionality

Abstract. This article aims to stimulate interdisciplinary exchange between logicians and cognitive scientists. In particular, I claim that conceptual analogues of hyperintensionality and intensionality can be found when we apply statistical tools to analyse sensorimotor processes in embodied cognition. When considering the functional correlation between the internal state X of an agent, and the external state Y of its environment, I propose that the precise functional form of the correlation has a hyperintensional flavour, while the abstract information carried by the correlation has a purely intensional flavour.

Recent work by Kolchinsky and Wolpert attempts to bring 'semantics' to physical correlations by analysing the effects of those correlations on task performance. I argue that this 'semantic information' framework currently provides a model for intensional, but not hyperintensional, aspects of belief in a hypothetical mental arithmetic scenario.

In general, I suggest that cognitive scientists should be more familiar with the intensional/hyperintensional distinction (for instance, I argue that the 'Bayesian brain' approach cannot account for hyperintensional aspects of cognition), and that logicians should be aware of analogues of hyperintension in embodied cognition (for instance, I claim that hyperintensional-like phenomena occur as much in bacteria as in humans).

Keywords: hyperintensionality; sensorimotor cognition; embodied cognition

1. Introduction

The primary focus of this article is the concepts of *intension* and *hyperintension* from formal semantics, and how they relate to issues in the study of embodied cognition *qua* physical dynamics. Effectively, I propose to understand something closely resembling (hyper)intensionality using the tools of information dynamics, rather than the tools of formal

Formal Methods in Philosophy. Guest Editors: K. Świętorzecka and K. Zdanowski
Received October 30, 2024. *Accepted* June 26, 2025. *Published online* August 27, 2025

semantics. Note that I do not propose any formal theory of embodied hyperintension in this article, instead offering only a conceptual account.

To this end, I apply the information agent formalism from (Ay and Polani, 2008) to a task involving mental arithmetic. The mental arithmetic scenario suggests informal concepts that form embodied sensorimotor (ESM) analogues of intension and hyperintension: I propose that

- different statistical physical ensembles $P(X, Y)$ and $P'(X, Y)$ over the internal state X of some agent and the external state Y of that agent's environment are 'intensionally equivalent' in the sensorimotor context iff X carries the same functionally meaningful information about Y under P as it does under P';
- two ensembles P and P' are 'hyperintensionally equivalent' in the sensorimotor context iff the correlation between X and Y plays (in some intuitive sense) the same functional role under P as it does under P'.

To put it another way (still using very informal language), intensionality (or its analogue in the sensorimotor context) relates to the abstract functional informational relationship between X and Y, while hyperintensionality (or its analogue) relates to the specific functional form of the relationship.

We will see that the *semantic information* framework introduced by Kolchinsky and Wolpert (2018) allows us to make a rigorously formal distinction between functionally meaningful correlation and brute statistical correlation; in doing so, it helps to provide a statistical semantics for embodied cognition.

However, I will argue that this semantics has a 'purely intensional' flavour, failing to capture important hyperintension-like aspects of embodied semantics. In particular, it does not distinguish between functionally different forms of meaningful correlation. My argument here is conceptual, and I do not pretend to offer a formal solution to this problem.

While ostensibly the article is about relations between physical ensembles, and how these relations resemble concepts from formal semantics, it is also meant to offer an opportunity for fruitful dialogue between two intellectual communities: scientists formally modelling the physical dynamics of embodied cognition, and philosophers studying formal semantics.

The consideration of this intersection may prove valuable to members of both communities. Scientists are generally unfamiliar with the formal semanticists' distinction between intension and hyperintension; contrarily, semanticists are generally unfamiliar with statistical tools intended to describe informational processes in embodied cognition.

Human linguistic and explicit reasoning behaviour is presumed by many scientists to form a special case of the more general category of embodied cognition, which is why general properties of information flow in embodied cognition might be relevant to semanticists. Indeed, I argue later that the embodied analogues of intension and hyperintension are as applicable to bacteria as they are to humans.

2. Extension, intension and hyperintension

The contents of this section will be familiar to philosophers, but since the article is meant to be accessible to a scientific audience as well, I will recap some basics regarding the notions of intensionality and hyperintensionality.

Formal semanticists distinguish between several different aspects of the concept of meaning. Consider the phrase 'the tallest dog in the world'. There is some specific dog that this phrase picks out right now (according to the Guinness Book of Records, in 2024 it is a Great Dane called Kevin). This is the *extension* of the phrase, which is one aspect of its meaning. But this dog (Kevin) cannot be the totality of what the phrase 'the tallest dog in the world' means, because, e.g., 'the tallest dog in the world in 2013' picks out a different dog (according to the Guinness Book of Records, another Great Dane called Zeus). If 'the tallest dog in the world' just meant Kevin, 'the tallest dog in the world in 2013' would mean 'Kevin in 2013'. According to the dominant formal model in semantics (the possible worlds model), the *intension* of the phrase is a function from things called 'possible worlds' (contexts in which the phrase picks out different things) to extensions, determining which dog (if any) the phrase picks out in every possible world.

While, for many purposes, we can think of the meaning of a phrase as being synonymous with its intension, there are some contexts in which we cannot do this. For instance, consider the phrases A "there are seven biscuits in the jar" and B "the number of biscuits in the jar squared, plus five, is fifty-four". These phrases have the same intension, because they map to the same extension (true or false) in every possible world. But it might be true that 'Alice believes there are seven biscuits in the jar' without it being true that 'Alice believes the number of biscuits in the jar squared, plus five, is fifty-four'. Formally speaking, 'Alice believes that' is a *hyperintensional* operator H, which means that $H(A)$ and $H(B)$ can differ even when A and B are intensionally equivalent.

The standard model of intensionality is the possible-worlds model that underpins modal logic (see, e.g., Nolan, 2013). There is no standard model of hyperintensionality, and there is much debate surrounding hyperintensionality in general (for instance, about whether hyperintensionality is a purely 'representational' phenomenon or a 'metaphysical' one). I will not attempt to engage with these issues; the interested reader may wish to consult Berto and Nolan (2023) for a broad introduction to the concept.

In the next section we will provide an example of why this distinction might be of interest to cognitive scientists, namely in identifying conceptual limitations of so-called 'Bayesian brain' approaches to cognition.

3. Bayesian brains and hyperintension

In 'Bayesian brain' approaches, cognition is taken to involve the manipulation of subjective probability distributions encoded in physical variables (Doya, 2006). This assumption has proved very useful in many areas of neuroscience. I will argue that this approach faces considerable challenges in accounting for intrinsically hyperintensional phenomena such as human logical reasoning[1].

Suppose I judge that the distance AB from Athens A to Berlin B is probably less than 1900km, and the distance BC from Berlin to Copenhagen C is probably less than 500km, and I also believe the triangle inequality, i.e. $AC \leq AB + BC$. I might still be momentarily uncertain as to whether the distance AC from Athens to Copenhagen is less than 2500km.

This is difficult to represent using a subjective probability distribution over states of the world. In a world that obeys the triangle inequality, $P(AC < 2500)$ can never be lower than $P(AB < 1900 \wedge BC < 500)$. Say I assign 90% credence to $AB < 1900$ and 80% credence to $BC < 500$. Then it turns out that the highest credence I can consistently assign to

[1] When I use the term 'logical reasoning' here, I mean the psychological process of correctly reasoning according to logical norms in ordinary humans; I'm not invoking an idealised 'logically omniscient' reasoner.

$AC \geq 2500$ is 15%. So I shouldn't be very uncertain about whether $AC < 2500$.

The key word here is *consistently*. Probability values that are assigned directly to states of the world have to be consistent, because the world itself has to be consistent. But my beliefs are not (in the parlance of logicians) closed under logical implication. They do not have to be consistent.

One could, of course, consider probability distributions over *descriptions* of states of the world, rather than over states of the world directly. For instance, one could ascribe probabilities to collections of sentences rather than states of the world, and permit $P(\text{“}AC < 2500\text{”})$ to be lower than $P(\text{“}AB < 1900 \wedge BC < 500\text{”})$, representing a logically inconsistent set of beliefs.

But a significant part of the appeal of Bayesian brain theories is that they account for an organism's behaviour, and they do so only because we can readily model how appropriate behaviour relates to states of the world. If you only have enough fuel for a 2500km flight from Athens, then Copenhagen is within range, while Dubai is not. There are consequences for actions that follow straightforwardly from these *facts about the world*. Conventional decision theory derives prescriptions for action from (intensional) facts about what the consequences of different actions are in different circumstances. Much more elaborate machinery is required to go from (potentially incomplete or inconsistent) sets of sentences about the world to prescriptions for motor actuations.

Moreover, sentence-based models may not be appropriate for describing practical non-linguistic reasoning. For instance, when I assess the distance between Athens and Copenhagen, I might rely on visualisation rather than inner speech, with my visualisation implicitly involving the triangle inequality (as I cannot simultaneously visualise three points where $AC > AB + BC$).

So Bayesian brain models are caught in a dilemma:

- they can operate at the intensional level, where the consequences for action are clearly defined, but hyperintensional phenomena like logical reasoning cannot be modelled; or
- they can operate at a hyperintensional level, allowing distinctions between different representations of the same thing, but obfuscating the relationship between mental state and behaviour.

The broader relevance of hyperintensionality to scientific models of cognition is an interesting topic (it is plausibly relevant to much-vexed discussions about the term ‘representation’), but it is not the primary focus of this article. I am interested in what scientific tools have to offer philosophy as well. In section 6 I will introduce the *information agent* formalism of Ay and Polani (2008), which I hope will have a profitable application to semantics. The next couple of sections will lay some groundwork by introducing *4E cognitive science* and sensorimotor cognition in particular.

4. 4E cognitive science and ‘reasoning’

Cognitive science is an interdisciplinary field which concerns itself with the study of cognition in the broadest sense, including intelligence in animals and other biological organisms, and adaptive behaviour in artificial systems such as robots or AIs, as well as questions about what the two have in common. Contributions to this endeavour have been made by philosophers, psychologists, linguists, computer scientists, roboticists, and animal behaviourists, among others.

I belong broadly to a school of thought known as 4E (embodied, embedded, enactive and extended) cognitive science (see, e.g., Clark, 1996; Newen et al., 2018). This perspective emphasises the conceptual importance of purposeful everyday interactions with the external world, such as navigating the environment, rather than taking ‘higher cognition’ to be the primary phenomenon of interest.

4E approaches to cognitive science reject any conception of cognition as a purely intellectual phenomenon. Instead they hold that cognition is a process of real-time, physically embodied, purposeful interaction between an agent and its environment. Under this view, bacterial chemotaxis (following a chemical gradient towards a nutrient) is a cognitive phenomenon; supposedly ‘higher’ forms of cognition (such as human logical reasoning) are simply more complex versions.

This means there will be important differences in how philosophers usually conceptualise reasoning, and how I will talk about it in this article. In particular, as a 4E cognitive scientist, when I consider mathematical or logical reasoning, I have in mind the complex and messy processes that take place in humans: biased and scaffolded by our own embodiment, potentially subject to everyday distractions from the outside (a doorbell ringing) or from the inside (a stomach growling), and potentially involving the external world (pen and paper, or discussion with a colleague). Moreover, the starting point of my analysis is a *task* that involves physical interaction with the world: the bodily process of naming a number out loud in a particular language.

This informal conception of logical reasoning does not invoke any abstract, disembodied, ideal mind (such notions are regarded with suspicion by 4E cognitive scientists) but rather seeks to better understand those processes we call 'reasoning' by situating them on a broader spectrum of processes that occur in real humans, other organisms, and artificial agents such as robots. In this way, reasoning is construed as a real-world process occurring over time, necessarily involving a succession of meaningfully different internal states.

Since this perspective treats reasoning as inherently diachronic, it may also have productive applications to discussions about *diachronic rationality* in philosophy (see, e.g., Staffel, 2019). However, I will not explore that topic in this article.

The main thesis of this article is that we can gain some useful insights by applying a particular set of formal tools (involving sensorimotor information flow) to a mental arithmetic task. This does not require the reader to endorse any of the tenets of 4E cognitive science (although it is certainly inspired by my own practice as a 4E cognitive scientist). In fact, in some sense I am defying a standard doctrine in 4E cognitive science, which holds that in order to understand 'higher' cognitive capacities such as abstract reasoning, we need to first understand simpler sorts of everyday bodily competence; instead, I am applying a sensorimotor information flow analysis directly to 'higher' cognition in the form of (embodied) mental arithmetic.

5. Sensorimotor cognition

One particular strand of 4E cognitive science focuses on *sensorimotor cognition*: this relates to the coupled dynamics of an agent and its environment through sensors, actuators, and the agent's bodily morphology as well as its brain[2]. In particular, this process is understood as a complex 'sensorimotor loop' in which effective problem-solving looks very different than 'offline', less-interactive scenarios.

For instance, an influential model of how outfielders in baseball catch fly balls is the *optical acceleration strategy* (Chapman, 1968), in which

[2] If it has a brain.

the outfielder moves to keep the ball's image moving at a constant rate in their visual field. This strategy is robustly successful, without requiring any detailed internal 'model' of how balls move, or any complicated planning process.

Sensorimotor dynamics is also the focus of O'Regan and Noë's (2001) *sensorimotor contingency theory*, which emphasizes that perception is not a passive reception of visual data, but an active process involving the exploration of the environment. In vision, for example, perceiving the 3D shape of an object involves understanding how sensory input changes with movement — a dynamic sensorimotor interaction.

There is much to be said about how supposedly 'higher' forms of cognition (such as human mathematical reasoning) relate to more basic sensorimotor processing (see, e.g., Lakoff and Johnson, 2008; Lakoff and Núñez, 2000): according to 4E cognitive scientists, conscious human reasoning is not separate from everyday pragmatic bodily competences, but builds on (and is biased by) those bodily competences in complex ways. However, we will not consider those issues in this article; instead, we will see how a sensorimotor approach can be applied directly to a mental arithmetic task.

6. Information agents

One approach to sensorimotor cognition is the information agent formalism (Ay and Polani, 2008), which models sensorimotor processes using the tools of statistical information theory, by considering the dynamics of a probability distribution over the state variables of the agent-environment system. This distribution can be interpreted as a statistical ensemble in the sense of statistical mechanics. Let's call the agent's internal state X_t and the environment's state Y_t; these are random variables in some statistical ensemble $P(X_t, Y_t)$. The agent is assumed to interact with its environment through 'sensor' and 'actuator' channels S_t and A_t. Figure 1 shows the Bayesian graph for the resulting stochastic process. (A Bayesian graph, also known as a Bayesian network, encapsulates conditional independences between random variables in a probability space. Each variable is conditionally independent of its non-descendants, given its parents.)

Intuitively, we can think of adaptive behaviour in terms of information entering via the sensors and being emitted via the actuators; this

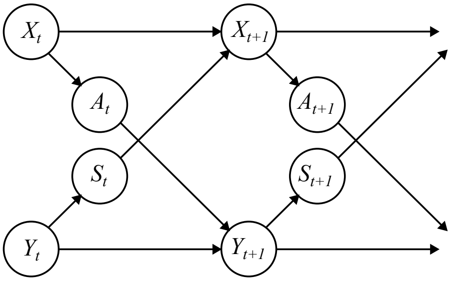

Figure 1. Partial Bayesian graph showing the conditional independence relations between successive time steps for internal and external variables X and Y and 'sensorimotor' variables S and A. In this paper, we assume that S is a deterministic function of Y and that A is a deterministic function of X.

intuition finds formal expression in the observation that (in such cases) there is non-zero mutual information between the agent's actions A_t and its previous sensory history $S_{<t}$; in other words $I(A_t; S_{<t}) > 0$.

This quantity $I(A_t; S_{<t})$ comes from Shannon information theory (see, e.g., Cover, 1999), and the information agent framework considers *information-theoretic* relationships between the variables in figure 1. Although Shannon originally framed his equations in terms of a 'theory of communication', his approach can be used more broadly, to measure degrees of abstract variation and covariation in random variables.

A key concept in Shannon information theory is the (Shannon) *entropy* $H(X)$ of a random variable X, defined for a discrete[3] probability measure p as

$$H(X) = -\sum_{x \in X} p(x) \ln p(x)$$

Entropy is usually described as a measure of *uncertainty*, but we can also think of it as simply an abstract measure of the variability of a distribution. Notice that discrete entropy depends only on the probability of outcomes x, and not on the value that x takes; this means that it is invariant to relabellings of the random variable X, unlike measures such as standard deviation which can change significantly if the measurement scale changes in a non-linear manner. Discrete entropy is strictly non-negative; it is minimised (equal to zero) when the probability distribution

[3] Shannon information theory can readily be extended from discrete distributions to continuous ones. In the continuous case, entropy is a bit slippery (due to how integration works, it is scale-dependent, and can be negative), but continuous mutual information has most of the same properties that discrete mutual information has.

is a delta function (i.e., all the probability mass is concentrated in a single value), and maximised by a uniform probability distribution.

Given a joint distribution over two variables $P(X, Y)$ the *mutual information* $I_P(X; Y)$ between X and Y can be defined as $H(X)+H(Y)-H(X, Y)$, where $H(X)$ and $H(Y)$ are the individual entropies of X and Y, and $H(X, Y)$ can be seen as the joint entropy of the random variable (X, Y) whose values are (x, y) pairs. Like discrete entropy, mutual information is non-negative and labelling-insensitive. When X and Y are independent, i.e. $P(X, Y) = P(X)P(Y)$, giving $I_P(X; Y) = 0$; when X and Y are related by some deterministic bijective function, we have $I_P(X; Y) = H(X) = H(Y)$.

Usually, for reasons I will omit here, this quantity $I_P(X; Y)$ is interpreted as the amount that one's uncertainty about X would decrease, on average, if one learned Y; it is symmetric, hence the term 'mutual' information. Alternatively, we can see it simply as a very general measure of how much X and Y are systematically related.

Since statistical (Shannon) information theory offers general measures of how non-uniform a probability distribution is, and how systematically related two variables are, it has extremely broad application. It has close relationships with algorithmic (Kolmogorov-Chaitin) information theory in computer science, and with thermodynamic quantities in physics. It is widely used in communications engineering, machine learning, theoretical neuroscience and bioinformatics, amongst other disciplines.

One might reasonably ask: what does any of this have to do with semantics? In brief, the answer is: the information agent framework applies as readily to (embodied applications of) conscious human reasoning as it does to bacterial cognition. We will see how this can be related to intension and hyperintension in the next section.

7. Information agents and (hyper)intension

Suppose Alice is performing a mental arithmetic task: she must compute the value $317 \times N$, for some positive integer N that she is given shortly after the task begins. Intuitively, we can think of Alice's beliefs changing as follows:

1. Alice begins at time t_1 not knowing either N or $317 \times N$. This is when she is told aloud in English that N has some specific value (say, "two hundred and forty five").

2. By time t_2, Alice has processed the auditory information specifying N, and now has a definite belief about the value of N (in this case, "N=245"), but has not performed the computation, so she is uncertain as to the value of $317 \times N$.
3. By time t_3, Alice has performed the computation and acquired the belief that $317 \times N$ has a particular value (in this case, "$317 \times N = 77665$").

The mental transition from t_2 to t_3 is a classic example of hyperintensional change: "N = 245" has the same intension as "$317 \times N = 77665$", because there is no possible world in which one of these holds and not the other. But since Alice's beliefs are not closed under logical entailment (in the formal semanticists' parlance), she can believe "N = 245" without being sure whether "$317 \times N = 77665$". This is why "Alice believes that" is a hyperintensional operator.

This is a standard philosophical account of the semantics involved in a mental arithmetic task. But if we shift our attention from Alice's beliefs to Alice's internal (physical) state, and apply statistical tools in the spirit of the information agent framework, we will see something rather interesting.

The fact that the function $f(n) = 317n$ is a bijection is what makes the belief "$N = \langle n \rangle$" possess the same intension as the belief "$317 \times N = \langle 317n \rangle$". A basic result in information theory establishes that mutual information is invariant under relabelling of variables[4], i.e. $I(X;A) = I(X;B)$ when there is a bijection between A and B. In our scenario: suppose some arbitrary random variable X is correlated with the random variable N; then X must be correlated to an identical degree with the random variable $317 \times N$.

Let's model the arithmetic task, across a range of values of N, using the Bayesian graph in Figure 2. At time 0, Alice's state X_0 is not correlated with the value of N. Sensory information about N enters through S_1 and is appropriately emitted some time later in Alice's answer A_3.

In this model, I have assumed that Alice does the calculation entirely "in her head", without storing any relevant information externally (e.g., by using pen and paper). This guarantees that correlation between A_3 and N is mediated entirely through Alice's intermediate state X_2, i.e.

[4] In the discrete case, invariant under any relabelling; in the continuous case, invariant under continuous relabelling transformations.

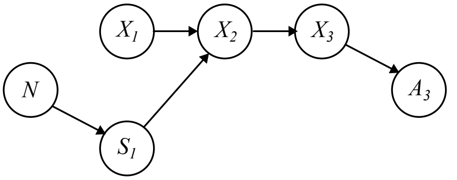

Figure 2. Bayesian graph for a mental arithmetic task. Irrelevant sensorimotor variables are not shown. N is a random number; X_t is Alice's internal state at time t; S_1 is a sensory variable at time $t = 1$; A_3 is an actuator value at time $t = 3$.

that the conditional mutual information $I(X_3; N \mid X_2)$ between N and X_3 given X_2 is equal to zero.

I'll also assume that Alice doesn't forget any information about N between time t_2 and t_3, which implies that the conditional mutual information $I(X_2; N \mid X_3)$ is also zero; putting $I(X_3; N \mid X_2) = 0$ and $I(X_2; N \mid X_3) = 0$ together, we obtain that $I(X_2; N) = I(X_3; N)$. In other words, X_2 correlates with N to the same degree that X_3 does.

Informally, we can think of Alice's dynamics as pumping statistical information about N that is contained in the physical sensory stimuli S_1 into her physical actuators A_3. Remember, the statistical information about N is identically information about $317 \times N$, because of the bijection. This same information (about N / $317 \times N$) is present in Alice's state at time $t = 1$, but it has not yet made it out into her actuators.

The statistical model here is a purely physicalist one (the variables are meant to be understood as physical variables), so in the absence of a theory specifying how to map physical states to mental states, we need to be careful about relating it to belief-talk and knowledge-talk. But it seems reasonable to suppose that knowing "$317 \times N = \langle 317n \rangle$" poises an English speaker to name $317 \times N$ in English more quickly than knowing "$N = \langle n \rangle$". This allows us to relate the semantics to the physics: in state X_3, the Alice-ensemble is poised to perform, without meaningful delay, the sequence of muscle actions that correspond to naming $317 \times N$ out loud in English; state X_2 also produces this sequence of muscle actions, but only after a delay (since they don't occur until time t_3).

In this sense, despite the fact that X_2 and X_3 correlate to the same extent with N, the correlations are not *functionally equivalent*. It matters whether the Alice-environment-ensemble is distributed according to $P(X_2, N)$ or $P(X_3, N)$: in the second ensemble, the information has been 'processed' into a form that can be fed directly[5] into her actuators; in the first ensemble, that has yet to happen.

At a conceptual (rather than formal) level, there is a tantalising analogy between the semantics and the statistical dynamics:

- Alice's beliefs "$N = \langle n \rangle$" at t_2 and "$317 \times N = \langle 317n \rangle$" at t_3 are intensionally equivalent, but hyperintensionally different. They can be said to represent the same (intensional) information about N, in the sense that they pick out the same set of possible values for N, but in a different form (in that their internal structure is different).
- The states X_2 and X_3 of the Alice-ensemble at times t_2 and t_3 carry the same (statistical) information about N, in that $I(X_3; N \mid X_2) = I(X_3; N \mid X_3) = 0$, but in a different form, since the distribution $P(X_3, N)$ differs from the distribution $P(X_2, N)$.

Moreover, it's reasonable to suppose that differences in the internal structure of Alice's beliefs correspond to differences in her internal state and sensorimotor dynamics. Apart from anything else, if she can respond more quickly to questions of the form "What is $317 \times N$" when she believes that "$317 \times N = \langle 317n \rangle$" than when she lacks such a belief, then this must correspond to some relevant functional difference in her internal state.

These parallels may naively suggest a straightforward model of sensorimotor intension: consider some joint distribution $P(X, X', Y)$. Then we might imagine saying that X and X' are 'intensionally equivalent' iff X and X' carry the same information about Y, relative to Y, i.e. $I(X; Y \mid X') = I(X'; Y \mid X) = 0$; and we might find some similar story to tell about 'hyperintensional equivalence'.

Unfortunately, even as a model of merely intensional equivalence, this naive formulation would overlook a well-known problem with statistical information theory: correlations are not always meaningful. I discuss this issue further in the next section.

8. 'Syntactic' and 'semantic' statistical information

Shannon information deals only with brute statistical correlation, sometimes called 'syntactic information' (Kolchinsky and Wolpert, 2018; Lombardi, 2004; Zhong, 2017). This complicates the study of information dy-

[5] Well, relatively directly. In reality, vocally articulating words is a complex real-time process involving muscular and auditory feedback.

namics in embodied cognition, where some correlations are meaningful and some are not.

For instance, C3 plants such as rice and wheat have less 13C in their tissues than C4 plants such as maize and sorghum; these carbon isotopes find their way into the tissues of animals which eat the plants (and animals that prey on those animals). Consequently, there is a correlation between the 12C / 13C isotope ratio in an animal's tissues, and the prevalence of C3 and C4 plants in its environment. This correlation seems of a different sort than the correlation between, e.g., activity in the optic nerve and visual stimuli presented to the eye; the latter clearly carries functionally relevant biological information, while (as far as we know) the information carried by isotopes in an animal's tissues about the plants in its environment plays no meaningful role.

Another way to think about this distinction is to observe that in the information-agent framework, a mutual information quantity such as $I(X;Y)$ represents the amount of information that X provides *to an external theorist* about Y, which need not correspond to any meaningful information *for the agent itself* (Beaton and Aleksander, 2012).

We can make this clear in the mental arithmetic scenario as follows. Suppose that, prior to t_1, Alice has swallowed a hardy capsule with N written inside it on a piece of paper. Although this is not cognitively relevant, it seems fair to describe the capsule as part of her internal physical state, since it is literally inside her.

This version of the task has a different Bayesian graph (see Figure 3). In the information agent framework, the environment can only affect the internal state of the organism via "sensory" variables, so we need to introduce a new variable S_0 which includes the state of the pill as it is swallowed. This variable feeds into X_1, correlating Alice's internal state with N even before she hears the value of N via S_1.

Remember that the statistics of this scenario represent information for the *theorist*. If an external daemon-like theorist, understanding the setup, wanted to know what N was, they could infer the same information about N's value from knowing the state of Alice's stomach (containing the pill) as they could from knowing the state of Alice's brain (having processed the name of N in English).

In the pill-swallowing scenario, both X_1 and X_2 carry the same information about N, albeit in different ways. According to the naive brute-statistical notion of embodied semantics suggested in the last sec-

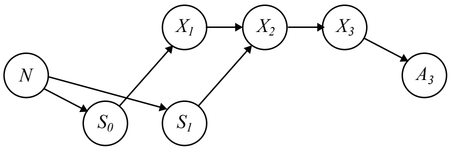

Figure 3. Bayesian graph for a mental arithmetic task with a functionally irrelevant variable introduced via a pill. N is a random number; X_t is Alice's internal state at time t; S_0 is a functionally irrelevant causal variable that correlates X_1 with n; S_1 is a 'true' sensory variable at time $t = 1$; A_3 is an actuator value at time $t = 3$.

tion, Alice's state X_1, before she has been told N in English, must thus be intensionally equivalent (relative to N) to her state X_2.

But this is nonsense. The reason that Alice gives the correct answer A_3 at time t_3 is because she heard N named via S_1; it is certainly not that she swallowed a pill with N written on it, since there is no functional pathway that allows this information to flow appropriately into Alice's vocal muscles.

In our model, this fact is a causal one, not a statistical one, since I have supposed that S_1 and S_0 are equally correlated with A_3. In order to see that S_1 is relevant while S_0 is not, we will need to consider counterfactual interventions à la Pearl (2000). Kolchinsky and Wolpert (2018) introduce a useful trick here (there is a similar approach in (McGregor and Mediano, 2018)), which we will encounter in the next section.

9. Intervening to decorrelate

In the pill-swallowing scenario, Alice's states X_1, X_2 and X_3 are all equally correlated with N. But we are interested in picking out those correlations that serve a purpose. We will pursue an approach similar in spirit to Pearlian intervention (Pearl, 2000) but which modifies entire distributions instead of individual variables, typically changing the Bayesian graph (unlike Pearl's framework). Suppose that we magically 'intervened' at time t to decorrelate X_t from N. In the original distribution P, we can factor the joint distribution $P(A_3, X_t, N)$ as follows:

$$P(A_3, X_t, N) = P(A_3 \mid X_t, N)P(X_t, N) \tag{1}$$

Let's define a 'scrambled' distribution $\hat{P}_t$, defined as

$$\hat{P}_t(A_3, X_t, N) = P(A_3 \mid X_t, N)P(X_t)P(N) \tag{2}$$

We have decorrelated X_t and N in $\hat{P}_t$, but retained the conditional dependence $P(A_3 \mid X_t, N)$.

In the scrambled ensemble $\hat{P}_1(A_3, X_1, N) = P(A_3 \mid X_1, N)P(X_1)$ $P(N)$, there is still a pill in Alice's stomach at time $t = 1$ with a number written on it, but this writing is no longer correlated with N. Now we can ask the question: how would the scrambled Alice-ensemble perform on the mental arithmetic task? It seems overwhelmingly likely that her performance would be unchanged by this scrambling. The important correlation, which is between her relevant brain state and the value of N, is preserved under $\hat{P}_1$ because the term $P(A_3 \mid X_1, N)$ involves the crucial pathway through S_1 and X_2.

Now let's consider the implications for Alice's performance if we scrambled at time $t = 2$. We have defined $\hat{P}_2(A_3, X_2, N) = P(A_3 \mid X_2, N)P(X_2)P(N)$. This corresponds to the behaviour of an ensemble in which Alice performs the mental arithmetic task, but for a number that is uncorrelated with N. Clearly, her performance will drop significantly in $\hat{P}_2$ compared to P.

Similarly, we can perform the scrambling at time $t = 3$, giving $\hat{P}_3(A_3, X_3, N) = P(A_3 \mid X_3, N)P(X_3)P(N)$. In this ensemble, Alice has completed the mental arithmetic task, but for a number uncorrelated with N, and gives her answer A_3. Again, Alice's performance will drop significantly.

Kolchinsky and Wolpert present a formal version of this idea. They begin with an ensemble $P(X, Y)$ and a 'viability function'[6] $V : \mathrm{Prob}(\mathcal{X} \times \mathcal{Y}) \to \mathbb{R}$. They then define what they call a 'viability-optimal intervention' distribution $\hat{P}^{\mathrm{opt}}(X, Y)$ (I will call it a 'semantic ensemble'), which is a 'scrambled' version of $P(X, Y)$ in which, intuitively speaking, X and Y have been decorrelated as much as possible without impacting performance on the viability function V.

The details of this optimal scrambling involve a 'coarse-graining' function ϕ of the sample space $\mathcal{Y}$, which allows to perform partial decorrelations as well as complete decorrelations (Kolchinsky and Wolpert,

[6] Kolchinsky and Wolpert specifically define V as the entropy of X after a fixed time t (under normal dynamics); this is a special case of the version I give here.

2018). Partial decorrelations turn out not to arise in my mental arithmetic scenario, so I will present a highly abbreviated version of the coarse-graining idea. Given a coarse-graining $\phi\colon \mathcal{Y} \to Z$, where Z is an arbitrary set of labels, they define the ϕ-scrambled distribution $\hat{P}^{\phi}$ as

$$\hat{P}^{\phi}(X = x, Y = y) = P(X = x \mid \phi(Y) = \phi(y))P(Y = y) \tag{3}$$

In the case where ϕ assigns the same label to every $y \in \mathcal{Y}$, this corresponds to a complete decorrelation $\hat{P}^{\phi}(X, Y) = P(X)P(Y)$; in the case where ϕ is the identity function, we do not decorrelate at all: $\hat{P}^{\phi}(X, Y) = P(X, Y)$. Other functions can correspond to partial decorrelation.

Kolchinsky and Wolpert define a semantic ensemble (*viability-optimal intervention*) $\hat{P}^{\text{opt}}$ as a scrambled distribution $\hat{P}^{\phi^*}$ such that

$$\phi^* = \underset{\phi : V(\hat{P}^{\phi}) = V(P)}{\arg\min} \hat{I}^{\phi}(X; Y) \tag{4}$$

where $\hat{I}^{\phi}(X; Y)$ is the mutual information between X and Y under the intervened distribution $\hat{P}^{\phi}$. Intuitively speaking, ϕ^* coarse-grains Y as much as possible without changing the performance V.

To incorporate coarse-graining into our scenario, we'll consider different coarse-grainings $\phi_t : t \in \{1, 2, 3\}$ that we will use to induce scrambled distributions as follows:

$$\hat{P_t}^{\phi_t}(A_3, X_t, N) = P(A_3 \mid X_t, N)P(X_t \mid \phi_t(N))P(N)$$

The above equation is a version of equation (2) that introduces partial dependence of X_t on N via a coarse-graining function ϕ_t, as per equation (3). The optimal scramblings (for Alice's mental arithmetic task) are straightforward; in the interests of brevity, I present only informal derivations here.

Where X_1 is concerned (Alice's state after ingesting the pill, but before hearing the value of N), a viability-optimal ϕ_1 assigns the same label to every n in the support of $P(N)$, e.g., $\phi_1(n) = \emptyset$. This means that

$$\hat{P_1}^{\phi_1}(A_3, X_1, N) = P(A_3 \mid X_1, N)P(X_1)P(N) \tag{5}$$

which is the same fully-decorrelated distribution $\hat{P}_1$ we saw in equation (2), with $t = 1$. As we discussed earlier, Alice's performance is the same under such a $\hat{P_1}^{\phi_1}$, and the scrambled mutual information $\hat{I}^{\phi_1}(X_1; N)$ is zero, which must be minimal (since mutual information is non-negative). Hence this ϕ_1 is viability-optimal.

By contrast, where X_2 and X_3 are concerned, a viability-optimal ϕ_2 or ϕ_3 assigns a distinct label to every n in the support of $P(N)$, e.g., $\phi_2(n) = n$. This means that

$$\hat{P}_2^{\phi_2}(A_3, X_2, N) = P(A_3 \mid X_2, N)P(X_2, N) \tag{6}$$

$$\hat{P}_3^{\phi_3}(A_3, X_3, N) = P(A_3 \mid X_3, N)P(X_3, N) \tag{7}$$

which are the fully-correlated distributions from equation (1) for $t = 2$ and $t = 3$. For both $t = 2$ and $t = 3$, any less-informative coarse-graining than this ϕ_t will impact Alice's performance, because it must lump together two distinct values n, n' which are meant to have different answers, but for which Alice's answer distribution will be the same. Hence these ϕ_t are viability-optimal.

Kolchinsky and Wolpert (2018) use their notion of an 'optimal intervention' (what I am calling a semantic ensemble) to define something called a 'semantic content' distribution: a conditional scrambled distribution $\hat{P}^{\phi}(Y \mid X = x)$ where ϕ is a viability-optimal coarse-graining of P under V.

I have some reservations about how to interpret this semantic content distribution. For instance, while Kolchinsky and Wolpert talks about 'the' semantic content distribution $\hat{P}^{\phi}(Y \mid x)$, in general there can be more than one distinct viability-optimal scrambling $\hat{P}^{\phi}$, since the arg min operator in (4) may not have a unique value. There are clearly nuances to be teased apart here, and work still to be done.

Nevertheless, whatever the correct interpretation of a semantic content distribution $\hat{P}^{\phi}(Y \mid x)$ may be, the distribution expresses something akin to a task-oriented concept of representational content. I will argue in the next section that this distribution is 'purely intensional' in its semantic flavour, failing to capture hyperintensional-like features, thus illustrating an embodied analogue of the distinction between intension and hyperintension.

10. 'Semantic content' distributions in the mental arithmetic task

As explained in the previous section, Kolchinsky and Wolpert use their notion of an 'optimal intervention' $\hat{P}^{\text{opt}}(X, Y)$ (what I am calling a semantic ensemble) to define something called a ‘semantic content’ distribution $\hat{P}^{\text{opt}}(Y \mid x)$ for a particular state x. By analogy, let’s call the conditional distribution $P(Y \mid x)$ under the syntactic ensemble P the ‘syntactic content’.

Consider again the graph in figure 3. I’ll suppose that N is a deterministic function $N = \lambda_1(X_1)$ of X_1 (Alice’s state after swallowing the inscribed capsule), another deterministic function $N = \lambda_2(X_2)$ of X_2 (Alice’s state after hearing the value of N in English, but before performing the computation), and yet another deterministic function $N = \lambda_3(X_3)$ of X_3 (Alice’s state after performing the computation), so that in general

$$N = \lambda_t(X_t) \tag{8}$$

Let us compare the syntactic content $P(N \mid X_t = x_t)$ and the semantic content $\hat{P}_t^{\text{opt}}(N \mid X_t = x_t)$ of states x_t at different times. Because we assumed that $N = \lambda_t(X_t)$, the *syntactic* content for every t is a Kronecker delta distribution

$$P(N = n \mid X_t = x_t) = \delta_{n\,\lambda(x_t)} \tag{9}$$

with all the probability mass concentrated on the particular value $n = \lambda_t(x_t)$. For *semantic* content the story is a little different:

$$\hat{P}_t^{\text{opt}}(N = n \mid X_t = x_t) = \begin{cases} P(N = n) \text{ if } t = 1 \\ \delta_{n\,\lambda_t(x_t)} \text{ otherwise} \end{cases}$$

This is because $\hat{P}_1^{\text{opt}}(N, X_t) = P(N)P(X_t)$ from equation (5), while for $t \in \{2, 3\}$ we have $\hat{P}_t^{\text{opt}}(N, X_t) = P(N, X_t)$ from equations (6) and (7), and hence $\hat{P}_t^{\text{opt}}(n \mid x_t) = P(n \mid x_t)$, which is $\delta_{n\,\lambda(x_t)}$ as in equation (9).

How does this correspond to an informal account of the correlations between Alice’s state X_t and N? The syntactic content distribution $P(N \mid x_t)$ of an internal state x_t at time t is essentially the epistemic state of an external daemon-like theorist who learns that $X_t = x_t$; in this case, this is the certain knowledge that $n = \lambda_t(x_t)$. This makes sense for all $t \in \{1, 2, 3\}$.

The correlation between N and Alice’s state X_1 (after swallowing the pill, but before hearing the English name of N) is not functionally meaningful, so the semantic distribution $\hat{P}_1^{\text{opt}}(N \mid x_1)$ for every x_1 is the uninformed prior distribution $P(N)$. Although the information about N

is present in Alice's state "for the theorist", it is not present in a way that can appropriately influence Alice's actions.

After hearing the value of N, Alice's state is (perfectly) functionally correlated with N, and the semantic content distribution $\hat{P}_t^{\mathrm{opt}}(N \mid x_t)$ is the Kronecker delta function $\delta_{n\,\lambda_t(x_t)}$. This is the same as the syntactic content, because all the information available to the theorist about N in Alice's state actually finds its way out into appropriate action.

Imagine that we sample values (x_2, x_3, n) from $P(X_2, X_3, N)$. Then, under the assumptions we have made, it will hold with probability 1 that $\lambda_2(x_2) = \lambda_3(x_3) = n$. Let's suppose that $\lambda_2(x_2)$ has the particular value 143. We might reasonably describe the state x_2 as encoding the fact that $N = 143$. But it's important to note that (intensionally speaking) this is equivalent to the fact that $(317 \times N) = 45331$.

So we cannot make the claim that Alice's state X_3 carries (meaningful) statistical information about the correct answer to the question 'What is $317 \times N$?' which was not present in X_2. But something important does change between X_2, before Alice has performed the computation, and X_3, after she has performed it. The thing that changes is *how* information about N is (meaningfully) encoded in Alice's state.

In this article, I do not offer a formal theory of how information is functionally encoded in internal state, but the next section discusses the concept further at an informal level.

11. The changing form of functional information

In this section I will consider why it might make intuitive sense to distinguish between the *form* and the *content* of (statistical) information carried by functionally meaningful correlations between Alice's state and her environment.

Alice's physical dynamics are such that, in order for the information about N in S_1 to feed into her vocal muscles so that she names the value of $(317 \times N)$ in English, a series of internal changes must occur. These changes do not involve the acquisition of new (Shannon) information about $(317 \times N)$: the changes do not make the overall state of the Alice-ensemble any more correlated with N. Instead, as it were, they are part of the process by which Alice pumps relevant information, over time, from her sensory stimuli into her actuators.

During this process of (meaningful) information working its way from sensors to actuators, the manner in which Alice's internal state encodes the relevant information changes. For instance, when $(317 \times N) = 45331$, this will be encoded at $t = 2$ by some x_2 such that $\lambda_2(x_2) = 143$ according to equation (8). But in this state, Alice is not yet ready to give the answer. Her state has to change into some x_3, where $\lambda_3(x_3) = 143$, before her muscles will twitch in the right way.

Philosophers often abstract away the sensorimotor processes involved in human cognition, talking directly about mental states such as belief, but the physics is worth paying attention to here. Very few physical systems can convert a series of pressure waves in air, carrying sounds that correspond to the name of an integer N in English, into a series of pressure waves that name the integer $(317 \times N)$ in human-audible English, over a wide range of different values of N. Rocks certainly do not perform this conversion. While humans like Alice are not the only system that can convert one to the other (e.g., Amazon's digital assistant Alexa can also do so), the capability requires some complicated physical mechanisms however it is implemented.

We saw that Alice's state X_2 at $t = 2$ encodes the value of N in a different form than her state X_3 at $t = 3$. This is not at all surprising, if we see these states as intermediate steps in the state of a machine that is implementing an intricately complex map from pressure waves to pressure waves.

In many real-world cases, there is complex intermediate work to be done when channelling sensory information into appropriate actuator behaviour (i.e. producing actuator responses that are appropriately matched to sensory stimuli). A succession of different internal states is required to convert sensory stimuli into actuator signals, and these successive states are functionally different from one another precisely because they constitute different steps in the process.

If I am right, then (an embodied analogue) of hyperintensionality can, in part, be accounted for by the sheer physical complexity of the stimuli-to-actuator conversion processes that underpin sensorimotor cognition. In particular, the 'pipeline' between the stream of sensory stimuli and the stream of actuator behaviour requires changes in the functional form of organism-environment correlation, even when no new relevant sensory information is being received. This is a fundamentally 4E explanation, to do with how bodies physically relate to, and interact with, their environment; it does not involve abstract properties of ideal thought.

12. Hyperintensionality in bacteria

I've drawn attention to the processes in Alice, during the mental arithmetic task, that involved transformations in how her internal state was correlated with her sensory stimuli, in a way that allowed her actuator behaviour to match the sensory stimuli appropriately. I have argued that we can relate these processes to the formal semanticists' distinction between intension and hyperintension.

Roughly speaking, I have contrasted what we might call the 'form' of a (functional) correlation with its 'content'; the 'form' is determined by the details of the joint distribution of the variables, while the 'content' is determined by more abstract information-theoretic properties. 'Form', I claim, is an analogue of hyperintension, while 'content', which is labelling-invariant, is an analogue of intension.

This narrative is very general: we can apply it to processes in bacteria just as readily as we can apply it to mental arithmetic in humans. It does not presuppose that the internal states of the system should be interpreted in terms of sentence-like 'beliefs' or even as 'representational' (except inasmuch as they functionally correlate with the external environment).

For instance, consider a bacterium trying to move towards food. We start at time t_1 with an ensemble of bacteria whose internal states X_1 are determined by the distribution $P(X_1)$. At time t_2 the bacterium detects a burst of chemicals S_1 which might be a toxin or a nutrient, resulting in a new state X_2. After a short processing delay, the bacterium enters state X_3 and emits an action A_3 which corresponds to the rotation of its flagella: if S_1 was toxic, the flagella rotate clockwise, thereby rotating the bacterium on the spot; if S_1 was nutritious, the flagella rotate anticlockwise, thereby moving the bacterium forward.

The Bayesian graph for this process is shown in figure 4. Note that this is identical to the graph for the mental arithmetic task in figure 2. In both cases,

- information enters the system's state X_2 through sensory stimuli S_1 at t_1 (sounds naming N in Alice's case; chemicals indicating N in the bacterium's case)
- and is emitted after a delay via state X_3 which is a proximate cause of the appropriate 'action' A_3 (sounds naming $317 \times N$ in Alice's case; direction of flagellar rotation in the bacterium's case).

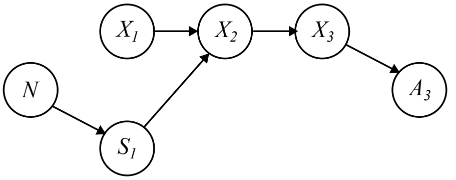

Figure 4. Bayesian graph for a bacterial locomotion task. Irrelevant sensorimotor variables are not shown. X_t is the bacterium's internal state at time t; S_1 is a sensory variable at time $t = 1$ corresponding to a chemical pulse; N is the source of the chemical pulse; A_3 is an actuator value at time $t = 3$ which is sensitive to S_1.

In Alice's case it was natural to attribute sentence-like 'beliefs' to her ("$N = \langle n \rangle$" at time t_2, and "$317 \times N = \langle 317n \rangle$" at time t_3). These sentence-like beliefs are not too far from sentences in the formal logics that historically underpin formal semantics.

But a sensorimotor information flow analysis is concerned with correlations between sensory stimuli history and appropriate actuator behaviour; these are physical variables, not sentence-like beliefs.

I told a dynamical narrative about Alice's performance on the mental arithmetic task: the correlations between her actuator behaviour and her sensory history are mediated by changes in internal state that themselves must be correlated with relevant features of the environment, and must somehow implement a complex transformation with a variety of intermediate 'forms' in which the same information can be manifested differently.

None of this dynamical narrative presupposes sentence-like beliefs. It is equally applicable to biological information-processing in single-celled organisms - a topic that is not usually the focus of formal semanticists.

13. Summary

I have claimed in this article that there are parallels between two different domains of research: the world of formal semantics, and the study of information dynamics in sensorimotor cognition. In logic, the notion of hyperintensionality arose because the possible-worlds model underpinning standard modal logics was inadequate for expressing certain operators. This concept is almost unknown amongst scientists, but I think it has relevance to the study of sensorimotor information dynamics.

For instance, Kolchinsky and Wolpert (2018)'s "semantic information" framework helps to bridge the gap between statistical 'syntax' and statistical 'semantics'. Unfortunately, this approach overlooks aspects of statistical 'semantics' which relate to functional changes that occur in the absence of new sensory information.

These functional changes are, so to speak, more to do with the particular internal physical *form* in which information is manifested, than the *content* of that information (in the sense of what it would tell an ideal external observer about the world); this informal distinction between content and form has conceptual parallels to the semanticists' distinction between intension and hyperintension.

Likewise, attention to the real-world dynamics of human cognition, and particularly to the functional pathway between sensors and actuators, may be of conceptual interest to formal semanticists. By abstracting such details away, they may be missing important parts of the picture. For instance, I've argued that 'sensorimotor hyperintensionality' occurs in bacteria as well as humans, and for the same reasons. Perhaps logicians could spare a few moments to contemplate the semantics of bacterial cognition.

I hope that the reader will forgive me for not making any very definite claims about exactly how we should formalise 'embodied sensorimotor hyperintensionality', how close these conceptual parallels are, or what insights they could offer to formal semanticists. The purpose of the article is to stimulate discussion, rather than to provide any answers.

Acknowledgments. I would like to thank Beatrice Buonaguidi and Artemy Kolchinksy for insightful and helpful comments on draft versions of this article, and in general to participants and organisers at "Formal Methods and Science in Philosophy V" who made an outsider feel welcome and offered encouragement for my research.

References

Ay, N., and D. Polani, 2008, "Information flows in causal networks" *Advances in Complex Systems*, 11(1): 17–41. DOI: 10.1142/S0219525908001465

Beaton, M., and I. Aleksander, 2012, "World-related integrated information: enactivist and phenomenal perspectives", *International Journal of Machine Consciousness*, 4(2): 439–455.

Berto, F., and D. Nolan, 2023. “Hyperintensionality”, in E. N. Zalta and U. Nodelman (eds.), *The Stanford Encyclopedia of Philosophy*, Metaphysics Research Lab, Stanford University, Winter 2023 edition. https://plato.stanford.edu/entries/hyperintensionality/

Chapman, S., 1968, “Catching a baseball”, *American Journal of Physics*, 36 (10): 868–870. DOI: 10.1119/1.1974297

Clark, A., 1996, *Being There: Putting Brain, Body, and World Together Again*, MIT Press. DOI: 10.7551/mitpress/1552.001.0001

Cover, T. M., 1999, *Elements of Information Theory*, John Wiley & Sons. DOI: 10.1002/0471200611

Doya, K., 2006, *Bayesian Brain: Probabilistic Approaches to Neural Coding*, MIT Press. DOI: 10.7551/mitpress/9780262042383.001.0001

Kolchinsky, A., and D. H. Wolpert, 2018, “Semantic information, autonomous agency and non-equilibrium statistical physics”, *Interface Focus*, 8(6): 20180041. DOI: 10.1098/rsfs.2018.0041

Lakoff, G., and M. Johnson, 2008, *Metaphors We Live By*, University of Chicago Press. DOI: 10.7208/chicago/9780226470993.001.0001

Lakoff, G., and R. Núñez, 2000, *Where Mathematics Comes From*, volume 6, New York: Basic Books.

Lombardi, O., 2004, “What is information?”, *Foundations of Science*, 9(2): 105–134. DOI: 10.1023/B:FODA.0000025034.53313.7c

McGregor, S., and P. A. M. Mediano, 2018, “Measuring fitness effects of agent-environment interactions”, *Artificial Life*, 24(3): 199–217. DOI: 10.1162/artl_a_00269

Newen, A., L. De Bruin, and S. Gallagher, 2018, *The Oxford Handbook of 4E Cognition*, Oxford University Press. DOI: 10.1093/oxfordhb/9780198735410.001.0001

Nolan, D., 2013, “6 possible worlds semantics”, pages 262–272 in G. Russell and D. Graff Fara (eds.), *Routledge Companion to Philosophy of Language*, Routledge. DOI: 10.4324/9780203206966.ch2_6

O’Regan, J. K., and A. Noë, 2001, “A sensorimotor account of vision and visual consciousness”, *Behavioral and Brain Sciences*, 24(5): 939–973. DOI: 10.1017/S0140525X01000115

Pearl, J., 2000, *Causality: Models, Reasoning and Inference*, Cambridge University Press, New York.

Staffel, J., 2019, "Credences and suspended judgments as transitional attitudes", *Philosophical Issues*, 29(1): 281–294. DOI: 10.1111/phis.12154

Zhong, Y., 2017, "A theory of semantic information", *China Communications*, 14(1): 1–17. DOI: 10.1109/CC.2017.7839754

SIMON MCGREGOR
University of Sussex
Brighton, United Kingdom
s.mcgregor@sussex.ac.uk
https://orcid.org/0000-0002-4090-0375