

Two kinds of information processing in cognition

Mark Sprevak
University of Edinburgh

5 July 2018

What is the relationship between information and representation? Dating back at least to Dretske (1981), an influential answer has been that information is a rung on a ladder that gets one to representation. Representation is information, or representation is information plus some other ingredient. In this paper, I argue that this approach oversimplifies the relationship between information and representation. If one takes current probabilistic models of cognition seriously, information is connected to representation in a new way. It enters as a property of the representational content as well as a property of the vehicles of that content. This offers a new, logically independent, way in which information and representation are intertwined in cognitive processing.

1 Introduction

There is a new way in which cognition could be information processing. Philosophers have tended to conceptualise cognition's relationship to Shannon information in just one way. This suited an approach that treated cognition as inference over representations of single outcomes (*there is a face here, there is a line there, there is a house here*). Recent work conceives of cognition differently. Cognition does not involve an inference over representations of single outcomes but an inference over *probabilistic representations* – representations whose content involves multiple outcomes and their estimated probabilities.

My claim in this paper is that these recent probabilistic models of cognition open up new conceptual and empirical territory for saying that cognition is information

processing. Empirical work is already exploring this territory and researchers are drawing tentative connections between the two kinds of information in the brain. In this paper, my goal is not to propose a specific relationship between these two quantities of information in the brain, although I sketch some possible relationships in Section 6. My goal is to convince you that there *are* two logically and conceptually distinct kinds of information whose relationship should be studied.

Before we proceed, some assumptions. My focus in this paper is only on Shannon information and its mathematical cognates. I do not consider other ways in which the brain could be said to process information.¹ Second, I assume a representationalist theory of cognition. I take this to mean that cognitive scientists find it *useful* to describe at least some aspects of cognition as processing representations. I focus on the role of Shannon information within two different kinds of representationalist model: traditional models and probabilistic models. My claim is that *if* one accepts a probabilistic model, then there are two ways in which cognition involves Shannon information processing. I do not attempt to defend representationalist theories of cognition in general.²

Here is a brief preview of my argument. Under probabilistic models of cognition there are two kinds of probability distribution associated with cognition. First, there is the ‘traditional’ kind: probability distributions associated with a specific neural state occurring in conjunction with an environmental state (for example, the probability of a specific neural state occurring when a subject is presented with a line at 45 degrees in a certain portion of her visual field). Second, there is the new, ‘probabilistic representation’ kind: probability distributions *represented by* those neural states. These probability distributions are the brain’s guesses about the probability of various environmental outcomes (say, that the line is at 0, 35, 45, or 90 degrees). The two kinds of probability distribution – one associated with *a neural/environmental state occurring* and the other associated with *a specific outcome represented by that neural state occurring* – are conceptually and logically distinct. They have different outcomes, different probability values, and different types of probability (objective and subjective) associated with them. They generate two separate measures of Shannon information. The algorithms that underlie cognitive processing can be described as processing *either* or *both* of these Shannon

¹See Floridi (2011).

²Note that I define representationalist theories in terms of their *utility* for describing cognitive processes, not in terms of their *truth*. Some deny truth but accept utility: they endorse some form of instrumentalism about representationalist models in cognitive science (for example, Egan, 2010; Colombo and Seriès, 2012; Sprevak, 2013). On my view, this still falls within the representationalist paradigm. To the extent it is legitimate, even if only on pragmatic grounds, to use a representationalist model of cognition, it is legitimate to say that cognition involves two kinds of information processing.

information quantities.

2 Shannon information

Before attributing two kinds of Shannon information to the brain, one first needs to know what justifies attribution of *any* kind of Shannon information. Below, I briefly review the definitions of Shannon information to identify conditions sufficient for a system to be ascribed Shannon information. The rest of the paper shows that these conditions are satisfied by the brain in two separate ways. Definitions in this section are taken from MacKay (2003), although similar points can be made with other formalisms.

In order to define Shannon information, one first needs the notion of a *probabilistic ensemble*:

- Probabilistic ensemble X is a triple (x, A_X, P_X) , where the outcome x is the value of a random variable, which takes on one of a set of possible values, $A_X = \{a_1, a_2, \dots, a_i, \dots, a_I\}$, having probabilities $P_X = \{p_1, p_2, \dots, p_I\}$, with $P(x = a_i) = p_i$, $p_i \geq 0$, and $\sum_{a_i \in A_X} P(x = a_i) = 1$

The conditions sufficient for the existence of a probabilistic ensemble include the existence of a random variable with *multiple possible outcomes* and an associated *probability distribution*. If the random variable has a finite number of outcomes, the probability distribution takes the form of a mass function, assigning a probability value, p_i , to each outcome. If the random variable has an infinite number of outcomes, the probability distribution takes the form a density function, assigning a probability value, p_i , to the outcome falling within a certain range. In either case, multiple outcomes and a probability distribution over those outcomes is a sufficient condition for a probabilistic ensemble to exist.³

If a physical system has multiple outcomes and a probability distribution associated with those outcomes, then that system can be described as a probabilistic ensemble. If a neuron has multiple possible outcomes (e.g. firing or not), and there is a probability distribution over those outcomes (the chance of the neuron firing), then the neuron's activity can be described as a probabilistic ensemble.

Shannon information is a continuous scalar quantity measured in bits. Shannon information is predicated of at least three different types of entity: *ensembles*, *outcomes*,

³In principle, an ensemble might have only one outcome (necessarily, with probability of 1). As we will see this corresponds to the ensemble and its single outcome having 0 bits of associated Shannon information. For information processing to be non-trivial, we need ensembles with more than one outcome.

and *ordered pairs of ensembles*. Let us consider each in turn.

The Shannon information, $H(X)$, of an *ensemble* is:

- $H(X) = \sum_i p_i \log_2 \frac{1}{p_i}$

The independent variables of $H(X)$ are the possible outcomes (*is*) and their probabilities (p_i s). The Shannon information of an ensemble is thus a mathematical function of, and only of, the essential properties of a probabilistic ensemble. Merely being an ensemble – having multiple outcomes and a probability distribution over those outcomes – is enough to define a $H(X)$ measure and bestow a quantity of Shannon information on the ensemble. Therefore, any physical system that is described as a probabilistic ensemble *ipso facto* has a measure of Shannon information. If a neuron is described as an ensemble (because it has multiple outcomes and a probability distribution over those outcomes), then it has a quantity of Shannon information attached.

The Shannon information, $h(x)$, of a *single outcome* is:

- $h(x = a_i) = \log_2 \frac{1}{p_i}$

In the definition of $H(X)$, an expected value is taken across all possible outcomes for $h(x)$. A sufficient condition for satisfying the definition of $h(x)$ is again the presence of an ensemble. If some ensemble exists, then each outcome of the ensemble automatically has a measure of Shannon information attached based on the definition above. No further conditions need to be met. If a neuron is described as an ensemble, then under the definition above, each of its outcomes (firing or not firing) has an associated quantity of Shannon information.

Many different Shannon measures of information are defined for a *pair of ensembles*.⁴ Common ones include:

Joint information

- $H(X, Y) = \sum_{x,y \in A_X A_Y} P(x, y) \log_2 \frac{1}{P(x, y)}$

Conditional information

- $H(X | Y) = \sum_{y \in A_Y} P(y) [\sum_{x \in A_X} P(x | y) \log_2 \frac{1}{P(x | y)}]$

Mutual information

- $I(X; Y) = H(X) - H(X | Y)$

Although these measures differ from each other, a sufficient condition for satisfying any one of them is the existence of a sufficient number of ensembles. If two ensembles

⁴One member of of an ordered pair is usually called the ‘sender’ and the other the ‘receiver’.

exist, then the pairwise measures above are defined. Individual ensembles, X and Y , may have some degree (or no degree at all) of probabilistic dependence on each other. Some joint probability distribution, $P(x, y)$, quantifies their relationship along with any conditional probability measures, such as $P(x | y)$. No more than this is required to define the Shannon measures above. If X and Y exist, and if their outcomes have a probabilistically quantifiable relationship to each other (even if that is probabilistic independence), then the Shannon measures of joint information, conditional information, and mutual information are defined. If two neurons are described as two ensembles, then their ordered pair has associated measures of joint information, conditional information, and mutual information.

A sufficient condition for a physical system to be attributed Shannon information (under any definition above) is that it *has multiple outcomes and a probability distribution over those outcomes*. Shannon information of an ensemble, a single outcome, or a pair of ensembles, is a function of, and only of, the outcomes and probability distribution associated with an ensemble, single outcome, or pair. If a physical system is described as an ensemble (or pair of ensembles), then *ipso facto* it has associated Shannon information.

If a physical system changes its probabilities over time, its associated Shannon measures are likely to change too. Such a system may be described as ‘processing’ Shannon information. This change could happen in at least two ways. If a physical system like a neuron modifies the probabilities associated with its *physical states* occurring (e.g. making certain physical states such as firing more or less likely), it can be described as processing Shannon information. Alternatively, if the firing activity of the neuron *represents* a probability distribution, and that represented probability distribution changes over time, perhaps in inference or learning, then its associated Shannon measures will change too. In both cases, probability distributions change. In both cases, Shannon information changes and is ‘processed’. But distinct probability distributions and distinct measures of Shannon information are changed in each case. The remainder of this paper will unpack the distinction between them.

3 The traditional kind of Shannon information

An important role for Shannon information has been as a building block in the project of *naturalising representation*. Many versions of information-theoretic semantics try to explain semantic content in terms of Shannon information. These accounts aim to explain how representation, and in particular mental representation, arises from Shannon information-theoretic relations. Such theories often

claim that Shannon information is a source of naturalistic, objective facts about representation. Dretske formulated one of the earliest such theories.⁵ Dretske's (1981) theory aimed to reduce facts about representation entirely to facts about Shannon information. More recent accounts, including Dretske's (1988; 1995), view an information-theoretic condition as only one part of a larger condition for naturalising representation. Additional requirements that need to be satisfied for representation might include conditions regarding teleology, instrumental (reward-guided) learning, structural isomorphism and/or appropriate use.⁶ In what follows, I focus solely on the information-theoretic part of a semantic theory.

Information-theoretic semantics explains representation in terms of some physical state 'carrying information' about another state. The relation of 'carrying information' is assumed to be a precursor to, or a precondition for, certain kinds of representation. In the context of the brain, such a theory says:

(R) Neural state, n (from N), represents an environmental state, s (from S) only if n 'carries information' about s .

Implicit in R is the idea that neural state, n , and environmental state, s , come from a set of possible alternatives. Neural state n represents s because n bears a 'carrying information' relationship to this s and not to others. Different neural states could occur in the brain (e.g. different neurons in a population might fire). Different states could occur in the environment (e.g. a face or a house could be present). The reason why certain neural firings represent a face and not a house is that those firings, and only those firings, bear the 'carrying information' relation to *face* outcomes, and the relevant firings do not bear this relation to *house* outcomes. R assumes that we are dealing with multiple possible outcomes: multiple possible representational vehicles (N) and multiple possible environmental states (S). It tries to identify a special relation between individual outcomes within that multiplicity. Representation occurs when n from N bears an 'carrying information' relationship to s from S .

The primary task for an information-theoretic semantics is to explain what 'carrying information' is. Different versions of information-theoretic semantics do this differently.⁷ Theories can roughly be divided into two camps: those that are 'correlational' and those that invoke 'mutual information'.

⁵Prior to Dretske's work, Shannon information had been linked to semantic content, although not always in reductive fashion (Wiener, 1961; Bar-Hillel and Carnap, 1964).

⁶See Millikan (1984); Papineau (1987); Dretske (1988); Shea (2007); Shea (2014); Skyrms (2010); W. M. Ramsey (2016) for a variety of such proposals.

⁷The relation of 'carrying information' is also sometimes described as one physical state 'having natural information' about another; see Stegmann (2015).

The starting point for ‘correlational’ theories is that one physical state ‘carries information’ about another just in case there is a correlation between the two that satisfies some probabilistic condition. This still leaves plenty of questions unanswered: What kind of ‘correlation’ is relevant to information carrying (Pearson, Spearman, Kendall, mutual information, something else?)⁸ How should physical states be typed into outcomes between which a correlation holds? How high should the value of a correlation be to produce information carrying? Does it matter to information carrying if the correlation is ‘accidental’ or underwritten by a law or modal disposition? Rival information-theoretic semantics spell out ‘correlation’ differently.

1. $P(S = s \mid N = n) = 1$
2. $P(S = s \mid N = n)$ is ‘high’
3. $P(S = s \mid N = n) > P(S = s \mid N \neq n)$

Dretske (1981) endorses (1): a neural state carries information about an environmental state just in case an agent, given the neural state, could infer with certainty that the environmental state occurs (and this could not have been inferred using the agent’s background knowledge alone). Millikan (2000; 2004) endorses (2): the conditional probability of the environmental state need only be ‘high’, where what counts as ‘high’ is a complex matter involving the probabilistic relation having influenced past use via genetic selection or learning.⁹ Shea (2007) and Scarantino and Piccinini (2010) propose that correlation should be understood in terms of probability raising as in (3): the neural state should make the occurrence of the environmental state more probable than it would have been otherwise.

At first glance, there may seem nothing particularly Shannon-like about proposals (1)–(3). Probability theory is sufficient to express the conditions sufficient for representation without any enrichment from Shannon’s concepts. The theories are perhaps better termed ‘probabilistic’ semantics than ‘information-theoretic’ semantics. Nevertheless, there is a legitimate way in which these accounts show that cognition is Shannon information processing. According to (1)–(3), ‘carrying information’ is a relation between particular outcomes and those outcomes come from ensembles which have probability distributions. A sufficient condition for a system to have Shannon information is that it *has multiple outcomes and a*

⁸Millikan (2001) suggests that one should look at the probabilistic relations that are ‘learnable’ for an agent: A is correlated with B, and hence carries information about B, if B is learnable (or inferable) from A. However, any degree of probabilistic dependence between A and B (no matter how slight) could, in principle, allow an agent to learn, or infer, one from the other. With suitable rewards on offer, even the mildest degree of probabilistic dependence could be a target of learning as an agent could reap arbitrarily large rewards from doing so. The notion of a ‘learnable’ relation – if it is not merely a synonym for *not probabilistically independent* – is as much need of explication as the notion of ‘correlation’.

⁹See Stegmann (2015), pp. 873–874 for helpful analysis of Millikan’s view.

probability distribution over those outcomes. (1)–(3) assure us that this condition is satisfied by a cognitive system that processes representations. According to (1)–(3), the representational content of a neural state arises because that neural state is an outcome from an ensemble with other possible outcomes (other possible neural states) that could occur with certain probabilities (and probabilities conditional on various environmental outcomes). If cognition involves representation, and representations gain their content via any of (1)–(3), then cognition *ipso facto* involves Shannon information. Shannon information attaches to representations because of the probabilistic nature of their vehicles. That probabilistic nature is essential their representational status. Consequently, to the extent that cognition can be described as processing representations, and to the extent that we accept one of these versions of information-theoretic semantics, cognition can be described as processing states with a probabilistic nature, and *ipso facto* can be described as processing states with Shannon information.

‘Mutual information’ versions of information-theoretic semantics unpack the ‘carrying information’ relation differently. They invoke the Shannon concept of mutual information – or rather, they invoke *pointwise* mutual information, the analogue of mutual information for pairs of single outcomes from two ensembles. Mutual information is the expected value of pointwise mutual information over all possible outcomes of a pair of ensembles.¹⁰ Pointwise mutual information for a pair of single outcomes, x, y , from two ensembles, X, Y , is defined as:

- $PMI(x, y) = \log_2 \frac{P(x, y)}{P(x)P(y)} = \log_2 \frac{P(x|y)}{P(x)} = \log_2 \frac{P(y|x)}{P(y)}$

Skyrms (2010) and Isaac (forthcoming) define the ‘informational content’ of a physical state, n , as a vector consisting of the $PMI(n, s)$ value for every possible environmental state, s , from S , $\langle PMI(n, s_1), \dots, PMI(n, s_n) \rangle$. Isaac identifies the meaning and semantic content of n with this PMI -vector. Skyrms says that the meaning/semantic content is likely to be a more traditional semantic object, such as a set of possible worlds. He says that this object can be derived from the PMI -vector by examining the environmental states responsible for high-value elements of the PMI -vector. The set of possible worlds in which these s_i s occur is the semantic content of n . This specifies the conditions under which n is ‘true’ or ‘accurate’.

Usher (2001) and Eliasmith (2005), like Skyrms and Isaac, appeal to pointwise mutual information to define ‘carrying information’. Unlike Skyrms and Isaac however, they define ‘carrying information’ as a relation that holds between a single neural state, n , and a single environmental state, s . They say that n carries information about s just in case s is the state for which $PMI(n, s)$ has its maximum value for that n . A neural state, n , carries information about the s that produces the peak value

¹⁰ $I(X; Y) = \sum_{x \in A_X A_Y} P(x, y) PMI(x, y)$

in its *PMI*-vector. Usher and Eliasmith connect this to ‘encoding’ experiments in neuroscience. In an encoding experiment, many environmental states are presented to a brain and researchers look for the environmental outcome that best predicts a specific neural response – that maximises $PMI(n, s)$ as one varies s for a given n . Usher and Eliasmith offer a second, rival definition of ‘carrying information.’ This is based around ‘decoding’ experiments. In a decoding experiment, researchers examine many neural signals and classify them based on which one best predicts an environmental state – which maximises $PMI(n, s)$ for a given s . Here, one looks for the peak value of $PMI(n, s)$ as one varies n and keeps s fixed. There is no reason why the encoding and decoding accounts of ‘carrying information’ should coincide: they might pick out two entirely different sets of informational relations between brain and world. Usher and Eliasmith argue that they provide two different, equally valid, information-theoretic accounts of neural representation.

On each of these semantic theories, Shannon information arises in a cognitive system via the properties of neural states *qua vehicles*. It is because a given neural state is an outcome from a range of possible alternatives combined with the probability of various environmental outcomes, that the cognitive system has the Shannon information-theoretic properties relevant to representation and hence to cognition. In the next section, I describe a different way in which Shannon information enters cognition. Here, the Shannon quantity arises, not from the probabilistic nature of the physical vehicles and environmental states, but from the *represented content*. Probabilistic models of cognition claim that the represented content of a neural state is probabilistic. This means that Shannon information attaches to a cognitive system via its content rather than via the probabilistic occurrence of its neural vehicles.

4 The new kind of Shannon information

Probabilistic models of cognition, like the representationalist accounts discussed in the previous section, attribute representations to the brain. Unlike those accounts however, probabilistic models do not aim to naturalise representation. They *help themselves* to the existence of representations. Their distinctive claim is that these representations have content of a particular kind. They are largely silent about *how* these representations get this content. In principle, a probabilistic model of cognition is compatible with a variety of semantic theories, including certain versions of information-theoretic semantics.¹¹

The central claim of any probabilistic model of cognition is that neural represent-

¹¹I discuss how information-theoretic semantics might interact with probabilistic models of cognition in Section 6.

ations have probabilistic content and they enter into some kind of probabilistic inference. Traditional approaches, including those described in Section 3, assume that *single outcomes* are the represented content of representational vehicles. A neural state, n , represents a single environmental outcome (or in the case of Skyrms, a single set of possible worlds). The job of a semantic theory is to explain how and why this representation relation occurs. The relationship between a representational vehicle and its content is assumed to be a relatively determinate matter. A common mark against information-theoretic proposals is that they sometimes fail to provide sufficiently determinate content in terms of a pairing to a single represented outcome (Dretske, 2009; Shea, 2013). Thinking about neural representations in these terms has prompted description of neural states early in V1 as *edge detectors*: their activity represents the presence (or absence) of an edge at a particular angle in a portion of the visual field. The represented content is a particular outcome (*edge at ~45 degrees*). Similarly, neurons in the inferior temporal (IT) cortex are described as *hand detectors*: their activity represents the presence (or absence) of a hand. The represented content is a particular outcome (*hand present*). Similarly, neurons in the fusiform face area (FFA) are described as *face detectors*: their activity represents the presence (or absence) of a face. The represented content is a particular outcome (*face present*) (for example, see Gross, 2007; Logothetis and Sheinberg, 1996; Kanwisher, McDermott and Chun, 1997).

There is increasing suspicion that neural representation is not like this. The represented content is rarely a single environmental state (e.g. *hand present*); rather, it is a probability distribution over many states. The brain represents multiple outcomes simultaneously to ‘hedge its bets’ during processing. This allows the brain to store, and make use of, information about multiple rival outcomes when it is uncertain which is the true outcome. Uncertainty may come from the unreliability of the perceptual hardware, or from the brain’s epistemic situation that even with perfectly functioning hardware it only has incomplete access to its environment. Ascribing probabilistic content to cognitive agents is not new (de Finetti, 1990; F. P. Ramsey, 1990). However, there is an important difference between past approaches and new probabilistic models. In the past, probabilistic representations were treated as *personal-level* states of a thinking agent – personal probabilities, credences, or degrees of belief. In the new probabilistic models, representations with probabilistic content are attributed to subpersonal *parts* of the agent – to neural populations, even single neurons. Regardless of which personal-level states are attributed to the agent, various parts of the agent token diverse (perhaps even conflicting) probabilistic representations. Thinking in these terms has prompted description of neural states early in V1 as probabilistically nuanced ‘hypotheses’, ‘guesses’, or ‘expectations’ about edges. Their neural activity does not represent one state of affairs (*edge at ~45 degrees*) but a probability distribution over multiple possible edge orientations (Alink

et al., 2010). Their representational content is a probability distribution over how the environment stands with respect to edges. Similarly, neural activity in the IT cortex does not represent one state of affairs (*hand present*) but a probability distribution over multiple possible outcomes regarding hands. The represented content is a probability distribution over how the environment stands with respect to hands. Similarly, neural activity in FFA does not represent one state of affairs (*face present*) but a probability distribution over multiple possible outcomes regarding faces. The represented content is a probability distribution over how the environment stands with respect to faces (Egner, Monti and Summerfield, 2010).

Traditional models of cognitive processing tend to describe that processing as a computational inference over specific outcomes – *if there is an edge here, then that is an object boundary*. Probabilistic models of processing tend to describe the processing as a computational inference over probability distributions – *if the probability distribution of edge orientations is this, then probability distribution of possible object boundaries is that*. Cognitive processing is a series of computational steps that use one probability distribution to condition, or update, another probability distribution. Representation in this inference may remain probabilistic right until the moment when the brain is forced to plump for a specific environmental outcome in action. At this point, it may select the most probable outcome from its current represented probability distribution conditioned on all its available evidence (or some other point estimate that is easier to compute).

Modelling cognition as a form of probabilistic inference does not mean modelling cognition as non-deterministic or chancy. The physical hardware and algorithms underlying the probabilistic inference may be entirely deterministic. Consider that when your electronic PC filters spam messages from incoming emails, it performs a probabilistic inference, but both the PC's physical hardware and the way in which it follows steps in its abstract algorithm are entirely deterministic. A probabilistic inference takes representations of probability distributions as input, yields representations of probability distributions as output, and transforms input into output based on rules of valid (or pragmatically efficacious) probabilistic inference. The physical mechanism and the algorithmic rules for processing the representations in this can be entirely deterministic. What makes the processing probabilistic is not the chancy nature of the physical vehicles or abstract rules but that probabilities feature in the represented content that is being manipulated.

Perhaps the best known example of a probabilistic model is Bayesian inference. The Bayesian brain hypothesis claims that cognition is a probabilistic inference in which probabilistic representations are processed according to Bayesian or approximately Bayesian rules (Knill and Pouget, 2004). Predictive processing is one way in which this Bayesian inference could be implemented in the brain (Clark, 2013). It is worth

stressing that the motivation for probabilistic representations, and for probabilistic models of cognition in general, is much broader than that for the Bayesian brain hypothesis (or for predictive processing). The brain's inferential methods could, in principle, depart very far from Bayesianism and still produce adaptive behaviour under many circumstances. It remains an open question to what extent humans are best modelled as Bayesian (or approximately Bayesian) reasoners. Probabilistic techniques adapted from AI, such as deep learning, reinforcement learning, and generative adversarial models, can produce impressive behaviour despite having complex and qualified relationships to Bayesian inference. The idea that cognition is a form of probabilistic inference is a much more general idea than that cognition is Bayesian. A researcher in cognitive science may subscribe to probabilistic representation even if they take a dim view of the Bayesian brain hypothesis.

The essential difference between a traditional representation and a probabilistic representation lies in their content. Traditional representations aim to represent a single state of affairs. In Section 3, we saw that schema R treats representation as a relation between a neural state, n , and an environmental outcome, s . The represented content is specified by a truth or accuracy condition. Truth/accuracy is assumed to be an all-or-nothing matter. A traditional representation effectively 'bets all its money' that a certain outcome occurs. An edge detector says that *there is an edge*. Multiple states of affairs may sometimes feature in represented content (for example, *there is an edge between ~43–47 degrees*), but those states of affairs are grouped together into a single outcome that is represented as true. There is no probabilistic nuance, or apportioning of different degrees of uncertainty, to different outcomes.

In contrast, probabilistic representations aim to represent multiple outcomes and a probability distribution over those outcomes. The probability distribution is a graded estimation of how much the system 'expects' that the relevant outcome is true. Unlike traditional representations, represented content does not partition environmental states into two classes (*true* and *false*). Representation is not an all-or-nothing matter; it involves a weighted probability value associated with various outcomes. In contrast to R , the probabilistic notion of representation pairs a neural state, n , with multiple outcomes and a probability distribution over those outcomes. As we will see, these represented outcomes need not even coincide with the outcomes of S . Whereas the content of a traditional representation is specified by a truth condition, the content of a probabilistic representation is specified by a probability mass or density function over a set of possible outcomes.

Probabilistic representations could, in principle, use any physical or formal vehicle to represent their content. There is nothing about the physical make-up of a representational vehicle that determines whether that representation is traditional rather

than probabilistic. Either type of representation could also, in principle, use any number of different representational formats as the formal structure over which its algorithms operate. Possible formal formats include a setting of weights in a neural network, a symbolic expression, a directed graph, a ring, a tree, a region in continuous space, or an entry in a relational database (Griffiths, Chater et al., 2010; Tenenbaum et al., 2011). The choice of physical vehicle and formal format affect how easy it is to mechanise the inference with computation (Marr, 1982). Certain physical vehicles and certain formal formats are more apt than others to serve in certain computations. But, in principle, there is nothing about the physical make-up or formal format of a representation to determine whether it is traditional or probabilistic. That is determined solely by its represented content.

The preceding discussion should not be taken as suggesting that a model of cognition may only employ one type of representation (probabilistic or traditional). A model of cognition could mix the two types of representation in its inferences, provided it has the right rules to take the system between them. The preceding discussion should also not be taken as suggesting that there are no conceptual or empirical connections between the two types of representation. Non-accidental links are certainly possible. For example, a cognitive system might use suitably structured complexes of traditional representations to express the probability calculus and thereby express probabilistically nuanced content (maybe this is what we do with the public language of mathematical probability theory). Alternatively, a cognitive system might use suitably structured complexes of probabilistic representations to express all-or-nothing-like truth conditions. Feldman (2012) argues that traditional representations can be approximated with probabilistic representations that have modal (highly peaked) probability distributions.¹² Traditional and probabilistic representations may be mixed in cognition, and one could, under the right conditions, give rise to the other.

5 Two kinds of information processing

We assumed at the start that cognition is profitably described by saying it involves representation. Representations, both traditional or probabilistic, have multiple outcomes and probability distributions associated with them. In Section 2, we saw that having multiple outcomes and a probability distribution over those outcomes is enough to have an associated measure of Shannon information. Consequently, both types of representation – traditional and probabilistic – have an associated measure

¹²Feldman calls these ‘symbolic representations’, but his claim is about their content, not about the format of their vehicles.

of Shannon information. What characterises the old ('traditional') type of Shannon information in cognition is that it is associated with probability of the *vehicle* occurring conditional on various environmental outcomes. What characterises the new ('probabilistic representation') type of Shannon information in cognition is that it is associated with the probabilities represented inside the representational *content*.

The degree to which these two kinds of Shannon information differ depends on the degree to which the two underlying probability distributions differ. In this section, I argue that the two distributions involve different *outcomes*, different *probability values*, and they must involve different *kinds of probability*.

Different outcomes. For traditional representations, the outcomes are the *possible neural states* and *possible environmental states*. The outcomes are the objective possibilities – neural and environmental – that could occur. What interests Dretske, Millikan, Shea, and others is to know whether a particular neural state from a set of alternatives (N) occurs conditional on a particular environmental state from a set of alternatives (S) occurring.¹³ In contrast, for probabilistic representations the outcomes are possible *represented states of the environment*. These are the ways in which system represents the environment to be. This need not be the same as what is genuinely possible. A cognitive system may be mistaken about what is possible just as it may be mistaken about what is actual. It may represent an environmental outcome that is impossible (e.g., winning a lottery one never entered) and it may fail to represent an environmental state that is possible (e.g., that it is a brain in a vat). Unless the cognitive system represents all and only what is genuinely possible, there is no reason to think that its set of represented outcomes will be the same as the set of objectively possible environmental outcomes. Hence, the outcomes assigned probability need not be the outcomes of S . Moreover, for the probability distributions to coincide, the system would need not only to represent possible environmental states but also to represent outcomes regarding its own neural states. Only in the special case of a cognitive system that (a) represents all and only the possible environmental states and (b) represents all and only its possible neural states, would the two sets of outcomes coincide.

Different probability values. Suppose the cognitive system, maybe through some quirk of design, represents all and only the real environmental and neural possibilities. Even in such a case, the probability *values* associated with outcomes are likely to differ. The probability values that pertain to traditional representations are the objective chances, frequencies, propensities, or some other measure of a neural state occurring conditional on a possible environmental state. What interests Millikan, Shea, and others are these objective probabilistic relations between neural

¹³Or whether an environmental state occurs conditional on some neural state occurring. Each can be exchanged for the other via Bayes' theorem.

states and environment states. In contrast, the probability values that pertain to probabilistic representations are the system's guesses about how probable different represented outcomes are. These are its *estimation* of the probabilities, not the actual objective probabilities. Brains are described as having 'priors' – representations of environmental outcomes as having certain unconditional probabilities of occurring. Brains are also described as having a 'likelihood function' or 'probabilistic generative model' – a representation of the probabilistic relations between outcomes. Researchers are interested in how the brain uses its priors and likelihood function to make an inference about unknown events, or how it updates its priors in light of evidence. All the aforementioned probabilities are subjective probabilities: guesses about the probabilities of outcomes and relations between them. Only a God-like cognitive system, who knows the objective truth, would assign the right probability values to those various represented outcomes and relations. Such a system would have a *veridical* (and a *complete*) probabilistic representation of its environment, its own neural states, and the relations between them. Achieving this may be a goal to which a rational agent aspires, but it is surely a position that few of us achieve.

Different kinds of probability. Even in the case of a God-like cognitive agent, we still have two kinds of probability distribution. In both cases, the $P(\cdot)$ values are probabilities and we are assuming (for the sake of argument) that the $P(\cdot)$ values are numerically identical. But the respective $P(\cdot)$ values measure different *kinds* of probability. For traditional representations, the $P(\cdot)$ values measure objective probabilities. These may be chances, frequencies, propensities, or whatever else that corresponds to the objective probability of the relevant outcome occurring.¹⁴ For probabilistic representations, the $P(\cdot)$ values measure subjective probabilities. These are the cognitive system's representation of how likely the relevant outcomes are to occur. Chances, frequencies, propensities, or similar are not the same as a system's representation of – its guess concerning – an event's probability. For our God-like cognitive agent, the two may agree in terms of their value, but the two are nonetheless distinct. One involves objective probabilities; the other involves subjective probabilities. They just happen to align in terms of values. Subjective probabilities do not become objective probabilities merely because they happen to accurately reflect them. No than a picture of a Komodo dragon becomes a living breathing Komodo dragon if that picture happens to be accurate. One is a representation; the other is a state of the world. One is a distribution of objective probabilities; the other is a system's (perhaps veridical) representation of those probabilities. The two are not the same. Well-known normative principles connect subjective and

¹⁴Different theorists in Section 4 take different views on the underlying nature of these objective probabilities. Shea (2007) says the probabilities are objective chances (although he does not say what chances are); Millikan (2000) focuses on the idea that they are frequencies and the resulting reference class problem. No one entertains the hypothesis that they are subjective probabilities.

objective probabilities. However, no matter which normative principles one believes should hold, and no matter whether an ideal agent satisfies them, the two probability distributions are distinct.¹⁵

Two kinds of probability distribution appear in probabilistic models of cognition. Each generates an associated measure of Shannon information. The two corresponding Shannon quantities are distinct: they involve different outcomes, different probability values, and different kinds of probability. This allows us to make sense of two kinds of Shannon information being processed in cognition: two kinds of probability distribution change under probabilistic models of cognition. Cognitive processing involves changes in a system's representational vehicles and changes in a system's probabilistic represented content. Information-processing algorithms that define cognition can be defined over either or both of these Shannon quantities.

6 Relationship between the two kinds of information

My claim in the preceding section is that the two forms of Shannon information are *conceptually* and *logically* distinct. I wish to stress that this does not rule out interesting connections between them. This section briefly highlights some possible connections.

6.1 Bridges from semantic theories

One is likely to be persuaded of some deep conceptual or logical connection between the two kinds of Shannon information if one endorses some form of information-theoretic semantics for probabilistic representations. Probabilistic models of cognition do not to explain how probabilistic representations gets their content. In principle, a probabilistic model is compatible with a range of semantic proposals, including potentially some version of information-theoretic semantics.

Skyrms' or Isaac's theory looks the most promising existing theory to adapt for an information-theoretic account of probabilistic content. Both their approaches already attribute multiple environmental outcomes plus a graded response for each outcome. However, it is not immediately obvious how to proceed. The probability distribution that n represents cannot merely be assumed to be the probability distribution of S . As we saw in Section 5, a probabilistic representation may misrepresent

¹⁵Skyrms agrees: 'objective and subjective information' may be carried by a neural state (2010, pp. 44–45). Skyrms' concern is with the objective probabilities that pertain to neural states and environmental states. However, he agrees that subjective probabilities (and hence, subjective information) may be carried by a neural state *qua* content.

the true possibilities and their probability values. We should allow for the possibility that a system's probabilistic representations may be wrong. A second consideration is that the represented probabilities appear to depend not only on the probabilistic relation of a single representational vehicle to environmental outcomes; they should also depend on what else the system 'believes'. The probability that a system assigns to *there is a face* should not be independent of the probability it assigns to *there is a person*, even if the two outcomes are represented by different neural vehicles. A noteworthy feature of the information-theoretic accounts of Section 3 is that they disregard relationships of probabilistic coherence between representations in assigning representational content. They assign content piecemeal, without considering how the contents may cohere. How to address these two issues and create an information-theoretic semantics for probabilistic representations is presently unclear.

If an information-theoretic semantics for probabilistic representations could be developed, it would provide a conceptual and logical bridge between the two Shannon measures. One kind of information (associated with represented probabilities) could not vary independently of the other (associated with the objective probabilities). The two would correlate in the cases to which the semantic theory applied. Moreover, if the semantic theory held as a matter of conceptual or logical truth, then the correlation between the two Shannon quantities would hold with similar strength. An information-theoretic account of probabilistic representation is likely to provide conceptual or logical connections between the two types of Shannon information in the brain. However, in the absence of such a semantic theory, it is hard to speculate on what those connections are likely to be.

If one is sceptical about the prospects of an information-theoretic account of probabilistic content, one is likely to be less inclined to see conceptual or logical connections between the two kinds of Shannon information. If one endorses Grice's (1957) theory of *non-natural* meaning, then the two measures will look to be conceptually and logically independent. Grice says that in cases of non-natural meaning, representational content depends on human intentions and not, say, the objective probabilities of a physical vehicle occurring in conjunction with environmental outcomes. There is nothing to stop any physical vehicle representing any content, provided it is underwritten by the right intentions. I might say that the proximity of Saturn to the Sun (appropriately normalised) represents the probability that Donald Trump will be impeached. Provided this is underwritten by the right intentions, probabilistic representation occurs. Representation is, in this sense, an arbitrary connection between a vehicle and a content: it can be set up or destroyed at will,

without regard for probabilities of the underlying events.¹⁶ If one endorses Grice's theory of non-natural meaning, there need be no conceptual or logical connections between the probabilities of neural and environmental states and what those states represent. One Shannon measure could vary independently of the other. This is not to say that we will find the Shannon two measures uncorrelated in the brain. But it is to say that if they correlate, that correlation does not flow from the semantic theory.

6.2 Empirical correlations

Regardless of connections that may flow from one's semantic theory, there are likely to be other, empirical connections between the two kinds of Shannon information in the brain. The nature of these connections will depend on which strategy the brain uses to 'code' its probabilistic representations. This coding scheme describes how the defining features of probabilistic content – probability values, overall shape of the probability distribution, summary statistics like mean or variance – map to physical activity of the brain or physical relations between the brain and environment. The specific scheme that the brain uses for coding its probabilistic content is unknown and the object of much speculation. Proposed options include that the rate of firing of a neuron, the number of neurons firing in a population, the chance of neurons firing in population, or the spatial distribution of neurons firing in a population, is a function of some defining feature of the represented probability distribution (see, for example, Barlow, 1969; Averbeck, Latham and Pouget, 2006; Ma et al., 2006; Deneve, 2008; Fiser et al., 2010; Griffiths, Vul and Sanborn, 2012; Pouget et al., 2013). According to these views, the probability of various neural states occurring varies in some regular way with their represented probability distribution. The consequent relationship between the objective probabilities and represented probabilities may be straightforward and simple or it may be complex, non-linear, and vary in different parts of the brain. The same applies to the relationship between the two Shannon measures. If an experimentalist were to know the coding scheme of the brain, she would be able to reliably infer one measure from the other. But the two kinds of Shannon information would remain conceptually and logically distinct, for the reasons given in Section 5.

Cognitive processing is sometimes defined over the traditional, physical Shannon measure. Saxe, Calderone and Morale (2018) describe how brain entropy during resting state, as measured by fMRI, correlates with general intelligence. Chang et al. (2018) describe how drinking coffee increases brain entropy during resting state.

¹⁶Skyrms argues against this that all meaning is natural meaning. All meaning depends on the physical probabilities that connect vehicles and their content.

Carhart-Harris et al. (2014) describe the relationship between consciousness and brain entropy, and how this changes under the psychedelic drug, psilocybin. Rieke et al. (1999) advocate an entire research programme that examines information-theoretic properties of neural vehicles (spike trains) and their relations to possible environmental outcomes. They argue that Shannon information-theoretic properties of the physical vehicles and environmental outcomes allow us to make inferences about possible and likely computations that the brain uses and the efficiency of its coding scheme. In each of these cases, the Shannon entropy is measured over possible neural vehicles and environmental states, not over their represented content (although the authors suggest that since the two are likely to be correlated via the brain's coding scheme, we can use one to draw conclusions about the other).

In contrast, Feldman (2000) looks at algorithms defined over the information-theoretic properties of the represented content. He argues that the difficulty in learning a new Boolean concept correlates with the information-theoretic complexity of the represented Boolean condition. Kemp (2012) and Piantadosi, Tenenbaum and Goodman (2016) extend the idea to general concept learning. They propose that concept learning is a form of probabilistic inference that seeks to find the represented concept that maximises the probability of the represented data classification. This cognitive process is described as the agent seeking the concept that offers the optimal Shannon compression scheme for the represented data. Gallistel and Wilkes (2016) describe associative learning as a probabilistic inference about the most likely cause of an unconditioned stimulus given the observations. They describe this in terms of Shannon information: the cognitive system starts with priors over hypotheses about causes chosen so as to have 'maximum entropy' (their probability distributions are as 'noisy' as possible consistent with the data); the cognitive system then aims to find the hypotheses that provide optimal compression (that maximise Shannon information) of the represented hypothesis and observed data. In general, theorists move smoothly between probabilistic formulations and information-theoretic formulations of a probabilistic inference when describing a cognitive process. In each of the cases described above, the Shannon information is associated not with the probabilities of certain neural vehicles occurring, but with the represented probability distributions over which inference is performed (although again, one might think that the two are likely to be connected).

6.3 Two versions of the free-energy principle

Friston (2010) invokes two kinds of Shannon information processing in two versions of his 'free-energy principle'.

First, Friston says that the free-energy principle is a claim about the probabilistic

inference performed by a cognitive system. He claims that the brain aims to predict upcoming sensory activation and it forms probabilistic hypotheses about the world that are updated in light of errors it makes in making this prediction. Shannon information attaches to the represented probability distributions over which the inference is performed. Friston says that the brain aims to minimise the surprisal of – the Shannon information associated with – new sensory evidence. When the brain is engaged in probabilistic inference, however, Friston says that its neural activity does not represent the full posterior probability distributions. Instead, the brain approximates them with simpler probability distributions, assumed to be Gaussians. Provided the brain minimises the Shannon-information quantity ‘variational free energy’ it will bring these simpler probability distributions into approximate correspondence with the true posterior distributions that a perfect Bayesian reasoner would have (Friston, 2009; Friston, 2010). Variational free energy is an information-theoretic quantity, predicated of the agent’s represented probability distributions, that measures how far those subjective probability distributions depart from the optimal guesses of a perfect Bayesian observer. According to Friston, the brain minimises ‘free energy’ and so approximates an ideal Bayesian reasoner.

Friston make a second, conceptually distinct, claim about cognition (and life generally) aiming to minimise free energy. In this context, the aim is to explain how cognitive (and living) systems maintain their physical integrity and homeostatic balance in the face of a changing physical environment. Cognitive (and living) systems face the problem that their physical entropy tends to increase over time: they generally become more disordered and the chance increases that they will undergo a fatal physical phase transition. Friston says that when living systems resist this tendency, they minimise free energy (Friston and Stephan, 2007; Friston, 2013). However, the free-energy minimised is not that which attaches to the represented, probabilistic guesses of some agent. Instead, it attaches to the objective probabilities of various possible (fatal) physical states of the agent occurring in response to environmental changes. Minimising free energy involves the system trying to arrange its internal physical states so as to avoid being overly changed by probable environmental transitions. The system strives to maintain its physical nature in equipose with likely environmental changes. The information-theoretic free-energy minimised here is defined over the objective distributions of possible physical states that could occur, not over the probability distributions represented by an agent’s hypotheses.

Minimising one free-energy measure may help an agent to minimise the other: a good, Bayesian reasoner is plausibly more likely to survive in a changing physical environment than an irrational agent. But they are not the same quantity. Moreover, any correlation between them could conceivably come unstuck. An irrational agent could depart far from Bayesian ideals but be lucky enough to live in an

hospitable environment that maintains its physical integrity and homeostasis no matter how badly the agent updates its beliefs. Alternatively, an agent may be a perfectly rational Bayesian and update its beliefs perfectly, but its physical environment may change so rapidly and catastrophically that it cannot survive or maintain its homeostasis. Understanding how Friston's two formulations of the free-energy principle interact – that pertaining to represented probabilities and that pertaining to objective probabilities – is ongoing work.¹⁷

7 Conclusion

Traditionally, philosophers have invoked Shannon information as a rung on a ladder to naturalise representation. In this context, Shannon information is associated with the outcomes and probability distributions of physical states and environmental states. This project obscures a novel way in which Shannon information enters into cognition. Probabilistic models treat cognition as an inference over representations of probability distributions. This means that probabilities enter into cognition in two distinct ways: as the objective probabilities of neural vehicles and/or environmental outcomes occurring and as subjective probabilities that describe the agent's expectations. Two types of Shannon information are associated respectively: Shannon information that pertains to the probability of the vehicle occurring and Shannon information that pertains to the represented probabilistic content. The former is conceptually and logically distinct from the latter, just as representational vehicles are conceptually and logically distinct from their content. Various (conceptual, logical, empirical) relations may connect the two kinds of Shannon information in the brain, just as various relations connect traditional vehicles and their contents. Care should be taken, however, not to conflate the two. For, as we know for the distinction between traditional vehicles and content, much trouble lies that way.

Bibliography

Alink, A., C. M. Schwiedrzik, A. Kohler, W. Singer and L. Muckli (2010). "Stimulus predictability reduces responses in primary visual cortex". In: *Journal of Neuroscience* 30, pp. 2960–2966.

¹⁷Colombo and Wright (2018) draw a similar contrast between the two formulations of the free-energy principle. They describe different versions of the free-energy principle as involving 'epistemic' and 'physical' probabilities.

- Averbeck, B. B., P. E. Latham and A. Pouget (2006). “Neural correlations, population coding and computation”. In: *Nature Reviews Neuroscience* 7, pp. 358–366.
- Bar-Hillel, Y. and R. Carnap (1964). “An outline of a theory of semantic information”. In: *Language and Information*. Reading, MA: Addison-Wesley, pp. 221–274.
- Barlow, H. B. (1969). “Pattern recognition and the responses of sensory neurons”. In: *Annals of the New York Academy of Sciences* 156, pp. 872–881.
- Carhart-Harris, Robin, Robert Leech, Peter Hellyer, Murray Shanahan, Amanda Feilding, Enzo Tagliazucchi, Dante Chialvo and David Nutt (2014). “The entropic brain: A theory of conscious states informed by neuroimaging research with psychedelic drugs”. In: *Frontiers in Human Neuroscience* 8, pp. 1–22.
- Chang, D., D. Song, J. Zhang, Y Shang, Q. Ge and Z. Wang (2018). “Caffeine caused a widespread increase in brain entropy”. In: *Scientific Reports* 8, p. 2700.
- Clark, A. (2013). “Whatever next? Predictive brains, situated agents, and the future of cognitive science”. In: *Behavioral and Brain Sciences* 36, pp. 181–253.
- Colombo, M. and P. Seriès (2012). “Bayes on the brain—On Bayesian modelling in neuroscience”. In: *The British Journal for the Philosophy of Science* 63, pp. 697–723.
- Colombo, M. and C. Wright (2018). “First principles in the life sciences: The free-energy principle, organicism, and mechanism”. URL: <http://philsci-archive.pitt.edu/14844/>.
- De Finetti, B. (1990). *Theory of Probability*. Vol. 1. New York, NY: Wiley & Sons.
- Deneve, S. (2008). “Bayesian spiking neurons I: Inference”. In: *Neural Computation* 20, pp. 91–117.
- Dretske, F. (1981). *Knowledge and the Flow of Information*. Cambridge, MA: MIT Press.
- (1988). *Explaining Behavior*. Cambridge, MA: MIT Press.
- (1995). *Naturalizing the Mind*. Cambridge, MA: MIT Press.
- (2009). “Information-theoretic semantics”. In: *The Oxford Handbook of Philosophy of Mind*. Ed. by A. Beckermann, B. P. McLaughlin and S. Walter. Oxford University Press, pp. 381–393.
- Egan, F. (2010). “Computational models: a modest role for content”. In: *Studies in History and Philosophy of Science* 41, pp. 253–259.

- Egner, T., J. M. Monti and C. Summerfield (2010). “Expectation and surprise determine neural population responses in the ventral visual system”. In: *Journal of Neuroscience* 30, pp. 16601–16608.
- Eliasmith, C. (2005). “Neurosemantics and categories”. In: *Handbook of Categorization in Cognitive Science*. Ed. by H. Cohen and C. Lefebvre. Amsterdam: Elsevier, pp. 1035–1055.
- Feldman, J. (2000). “Minimization of Boolean complexity in human concept learning”. In: *Nature* 407, pp. 630–633.
- (2012). “Symbolic representation of probabilistic worlds”. In: *Cognition* 123, pp. 61–83.
- Fiser, J., P. Berkes, G. Orbán and M. Lengyel (2010). “Statistically optimal perception and learning: From behavior to neural representations”. In: *Trends in Cognitive Sciences* 14, pp. 119–130.
- Floridi, L. (2011). *The Philosophy of Information*. Oxford: Oxford University Press.
- Friston, K. (2009). “The free-energy principle: a rough guide to the brain?” In: *Trends in Cognitive Sciences* 13, pp. 293–301.
- (2010). “The free-energy principle: A unified brain theory?” In: *Nature Reviews Neuroscience* 11, pp. 127–138.
- (2013). “Life as we know it”. In: *Journal of the Royal Society Interface* 10, p. 20130475.
- Friston, K. and K. E. Stephan (2007). “Free-energy and the brain”. In: *Synthese* 159, pp. 417–458.
- Gallistel, C. R. and J. T. Wilkes (2016). “Minimum description length model selection in associative learning”. In: *Current Opinion in Behavioral Sciences* 11, pp. 8–13.
- Grice, P. (1957). “Meaning”. In: *Philosophical Review* 66, pp. 377–388.
- Griffiths, T. L., N. Chater, C. Kemp, A. Perfors and J. B. Tenenbaum (2010). “Probabilistic models of cognition: exploring representations and inductive biases”. In: *Trends in Cognitive Sciences* 14, pp. 357–364.
- Griffiths, T. L., E. Vul and A. N. Sanborn (2012). “Bridging levels of analysis for probabilistic models of cognition”. In: *Current Directions in Psychological Science* 21, pp. 263–268.
- Gross, C. G. (2007). “Single neuron studies of inferior temporal cortex”. In: *Neuropsychologia* 46, pp. 841–852.

- Isaac, A. M. C. (forthcoming). “The semantics latent in shannon information”. In: *The British Journal for the Philosophy of Science*. DOI: [10.1093/bjps/axx029](https://doi.org/10.1093/bjps/axx029).
- Kanwisher, N., J. McDermott and M. M. Chun (1997). “The fusiform face area: A module in human extrastriate cortex specialized for face perception”. In: *Journal of Neuroscience* 17, pp. 4302–4311.
- Kemp, C. (2012). “Exploring the conceptual universe”. In: *Psychological Review* 119, pp. 685–722.
- Knill, D. C. and A. Pouget (2004). “The Bayesian brain: the role of uncertainty in neural coding and computation”. In: *Trends in Neurosciences* 27, pp. 712–719.
- Logothetis, N. K. and D. L. Sheinberg (1996). “Visual object recognition”. In: *Annual Review of Neuroscience* 19, pp. 577–621.
- Ma, W. J., J. M. Beck, P. E. Latham and A. Pouget (2006). “Bayesian inference with probabilistic population codes”. In: *Nature Neuroscience* 9, pp. 1432–1438.
- MacKay, D. J. C. (2003). *Information Theory, Inference, and Learning Algorithms*. Cambridge: Cambridge University Press.
- Marr, D. (1982). *Vision*. San Francisco, CA: W. H. Freeman.
- Millikan, R. G. (1984). *Language, Thought and Other Biological Categories*. Cambridge, MA: MIT Press.
- (2000). *On Clear and Confused Ideas*. Cambridge: Cambridge University Press.
- (2001). “What has natural information to do with intentional representation?” In: *Naturalism, Evolution and Mind*. Ed. by D. Walsh. Cambridge: Cambridge University Press, pp. 105–125.
- (2004). *The Varieties of Meaning*. Cambridge, MA: MIT Press.
- Papineau, D. (1987). *Reality and Representation*. Oxford: Blackwell.
- Piantadosi, S. T., J. B. Tenenbaum and N. D. Goodman (2016). “The Logical Primitives of Thought: Empirical Foundations for Compositional Cognitive Models”. In: *Psychological Review* 123, pp. 392–424.
- Pouget, A., J. M. Beck, W. J. Ma and P. E. Latham (2013). “Probabilistic brains: Knows and unknowns”. In: *Nature Neuroscience* 16, pp. 1170–1178.
- Ramsey, F. P. (1990). *Philosophical Papers*. Ed. by D. H. Mellor. Cambridge: Cambridge University Press.
- Ramsey, W. M. (2016). “Untangling two questions about mental representation”. In: *New Ideas in Psychology* 40, pp. 3–12.

- Rieke, F., D. Warland, R. R. van Steveninck and W. Bialek (1999). *Spikes*. Cambridge, MA: MIT Press.
- Saxe, G. N., D. Calderone and L. J. Morale (2018). “Brain entropy and human intelligence: A resting-state fMRI study”. In: *PLoS ONE* 13, e0191582.
- Scarantino, A. and G. Piccinini (2010). “Information without truth”. In: *Metaphilosophy* 41, pp. 313–330.
- Shea, N. (2007). “Consumers need information: Supplementing teleosemantics with an input condition”. In: *Philosophy and Phenomenological Research* 75, pp. 404–435.
- (2013). “Naturalising representational content”. In: *Philosophy Compass* 8, pp. 496–509.
- (2014). “Exploitable isomorphism and structural representation”. In: *Proceedings of the Aristotelian Society* 114, pp. 123–144.
- Skyrms, B. (2010). *Signals*. Oxford: Oxford University Press.
- Sprevak, M. (2013). “Fictionalism about neural representations”. In: *The Monist* 96, pp. 539–560.
- Stegmann, U. E. (2015). “Prospects for probabilistic theories of natural information”. In: *Erkenntnis* 80, pp. 869–893.
- Tenenbaum, J. B., C. Kemp, T. L. Griffiths and N. D. Goodman (2011). “How to Grow a Mind: Statistics, Structure, and Abstraction”. In: *Science* 331, pp. 1279–1285.
- Usher, M. (2001). “A statistical referential theory of content: Using information theory to account for misrepresentation”. In: *Mind and Language* 16, pp. 311–334.
- Wiener, N. (1961). *Cybernetics*. 2nd ed. New York, NY: Wiley & Sons.