

Correspondence Theory of Semantic Information

Marcin Miłkowski

Abstract

A novel account of semantic information is proposed. The gist is that structural correspondence, analysed in terms of similarity, underlies an important kind of semantic information. In contrast to extant accounts of semantic information, it does not rely on correlation, covariation, causation, natural laws, or logical inference. Instead, it relies on structural similarity, defined in terms of correspondence between classifications of tokens into types. This account elucidates many existing uses of the notion of information, for example, in the context of scientific models and structural representations in cognitive science. It is poised to open a new research program concerned with various kinds of semantic information, its functions, and its measurement.

- 1 *Introduction*
- 2 *From Structural Similarity to Information*
- 3 *Defining Informational Structure*
- 4 *Defining Similarity as Accuracy*
 - 4.1 *Similarity as infocorrespondence*
 - 4.2 *From accuracy to truth value*
- 5 *Applications*
 - 5.1 *Scientific models*
 - 5.2 *Cognitive maps*
- 6 *Possible Objections*
 - 6.1 *Similarity versus covariation*
 - 6.2 *Too many correspondences*
 - 6.3 *No propositional content*
- 7 *Conclusions*

1 Introduction

The purpose of the theory of semantic information is to provide an account of how information vehicles, whose information capacity may be accounted for in formal theories of information (Shannon [1948]), can attain a semantic value such as truth. As is well known, Shannon ([1948], p. 349) stressed that ‘semantic aspects of communication are irrelevant to the engineering problem’—which he was busy solving—even if information is valuable to its users only when it has these aspects.¹ These semantic aspects will be understood, for the purposes of this article, in terms of satisfaction conditions. Suppose now that satisfaction is a binary property of pieces of information. Then a declarative piece of semantic information is satisfied if and only if it is true, and an instructional piece of information is satisfied if and only if it is followed or executed.² Thus, the problem to be solved here is to answer the question of what constitutes the satisfaction conditions of a piece of information.

The basic inspiration for the answer defended in this article comes from the correspondence theory of truth, which, in its contemporary formalized version, also uses the notion of satisfaction (Tarski [1933]). In a nutshell, a piece of information is true if and only if it corresponds to what it is about. The rest of this article is concerned with elucidating the notion of correspondence in terms of structural similarity. The issue of how to measure correspondence-based seman-

¹ Appearances to the contrary notwithstanding, Shannon's theory is not independent of these semantic aspects (Isaac [2017]).

² As it will become clear, the notion of satisfaction defended below can be graded; thus, declarative information is satisfied to the degree that it is true (so specifying its truth value in logic requires multivalued logic, such as fuzzy logic (Zadeh [1988])).

tic information is set aside for another occasion. The current account also abstracts away from the use of semantic information, but this is only a strategic idealization. In fact, one must take the users of information into account to apply the correspondence-based theory of information in an appropriately constrained fashion. But, to do this, one must first define the basic semantic notions, which is the task of the current article.

This task stems from a fairly long philosophical tradition and the novelty of the current theory lies in making the notion of semantic information explicitly depend on similarity. The use of similarity in philosophical accounts of representational contents can be traced back at least to Aristotle's metaphor of a signet ring being impressed on a piece of wax. Naturally, Aristotle had nothing to say about the modern notion of information. Neither did Locke, whose claim that signification depends on causality or resemblance set the stage for the ongoing debate on representation. It also influenced Peircean semiotics, in which similarity underlies the notion of icon, which is one of the fundamental types of signs (Short [2007]). Peirce, in turn, had immense impact on biosemantics (Milikan [1984]). In the philosophy of cognitive science, there is now an ongoing debate on whether both causality and resemblance are required to account for mental representation (Ramsey [2007]; Morgan [2013]; Gładziejewski and Miłkowski [2017]; Rupert [2018]; Thomson and Piccinini [2018]; Nirshberg and Shapiro [2020]). In this article, it will be argued that if one understands that there are many kinds of similarity, it becomes clear that at least covariation and correlation can be understood as particular kinds of similarity, which makes similarity more basic. However, these notions can be usefully kept distinct in a pluralist account of meaning (we will return to this point in Section 6.1).

Defenders of structural representations stress that resemblance between the structure of representation vehicles and representation targets is crucial to the representation vehicles' contents (Craik [1943]; Cummins [1996]; Opie and O'Brien [2004]; Gładziejewski and Miłkowski [2017]; Neander [2017]; Shea [2018]). The idea of resemblance as grounding the meaning of representations is also present in the literature on scientific models (Giere [2004]; Weisberg [2013]). While the issue of whether scientific representation is conceptually distinct from mental representation (see Callender and Cohen [2006]) is beyond the scope of this article, there is one thing that similarity accounts of mental and scientific representation share: It is not merely the reliance on structural similarity. Arguably, it is something more: The implicit reliance on semantic information, underwritten by structural similarity.³ This kind of information is called 'correspondence-based semantic information' in this article. Thus, correspondence-based semantic information is a more basic feature that is inherent in what may function as a representation.

Surprisingly, however, while correlation (Dretske [1982]; Shea [2018]) and causation (for example, Dilworth [2008])⁴ were both assumed to be at least partially responsible for the semantic character of information, similarity received no such treatment. Even Neander, in her treatment of informational teleosemantics, does not make the notion of semantic information explicit in the proposed

³ Interestingly, in her account of iconic signs, influential in the Lvov-Warsaw School, Joanna Kotarbińska ([1957]) relies on the notion of information to tackle the problem of too many correspondences (see below, Section 6.2). The current proposal goes the opposite way: it uses the notion of similarity to explicate the notion of information, without relying on the intentional information supplied by sign users.

⁴ Note that while Dretske is commonly interpreted as defending of a notion of semantic information based on causation, he clearly distinguishes between causation and information (Dretske [1982]).

view on structural representation (Neander [2017]). Our task, then, is to provide the account of semantic information that is inherent in both mental and scientific representation, but remains implicit in existing philosophical accounts. This correspondence-based semantic information could be understood as more primitive and basic than mental representation. Arguably, any account of mental representation would be incomplete without a plausible story of how the information is used in an organism (Eliasmith [2005]; Morgan [2013]). This account follows Dretske's recipe for building a theory of mental representation: First understand natural meaning (or natural semantic information) and then account for how it may function by playing a representational role.⁵ Some other accounts of mental representation allow the existence of unused or unexploited representations. For example, Cummins (Cummins *et al.* [2006]) defends the notion of 'unexploited contents', and Shapiro ([1997]) hypothesizes the existence of 'junk representations'. This proposal offers a way to rephrase their claims in terms of semantic information: Namely, there could be unexploited semantic information just as there might be some tree rings that never happen to be exploited as signs. This approach has the virtue of not making the notion of representation dangerously liberal (as critics of unexploited contents suggest; see Ramsey [2007]), while still admitting that there could be information-bearing structures thanks to naturally occurring similarities.

The article proceeds as follows. In Section 2, a simple example is introduced to illustrate the common uses of correspondence in understanding semantics of

⁵ In contrast to Dretske, however, it is not presupposed that semantic information is 'out there' just to be registered by agents; in contrast, correspondence-based information can be teased out from complex interactions among multiple users (just as in locally recurrent natural signs; see Millikan [2004]).

information. This section also lays down the desiderata for the theory of semantic information described in the rest of the article. The major difficulties in making this kind of correspondence explicit are: (1) how to characterize the relata of the correspondence relation in terms consistent with information theory, and (2) how to understand what correspondence is. Consequently, Section 3 defines the relata of the similarity relationship for the purposes of the account of semantic information. In Section 4, it is argued that there are, in principle, an infinite number of similarity types that could give rise to satisfaction conditions. A generic notion of similarity is introduced by stressing its role in so-called surrogative reasoning (Swyer [1991]). It is argued that this role can be captured in terms of conditions of information flow. This generic kind of similarity is subject to further analysis, which implies that there are multiple kinds of correspondence-based semantic information. Section 5 provides two simple applications of the theory. The first concerns scientific models in natural language engineering and the second, structural representations in cognitive science. Section 6 addresses possible objections to the proposed account. Finally, the article is concluded by indicating that the current contribution opens an exciting line of research into various kinds of semantic information, their uses in inferential processes, and the problems in measuring their contents. Hopefully, the present account will contribute to a better understanding of what is implicit in existing theories of representation, both mental and scientific.

2 From Structural Similarity to Information

The account proposed here is both novel and already presupposed in everyday usage of the term 'information'. These two features may seem contradictory.




Figure 1.

Nonetheless, talk of information supported by correspondence abounds, not only in the debate over mental or scientific representation, but also in everyday conversations about computer files and other artefacts. The task at hand is to understand why and how this talk is valid in terms of semantic information.

Take the two pictures of an infinity symbol in Figure 1. The smaller infinity symbol (A) is composed of mostly black dots, which remain barely visible on a computer screen because of the pixel size (note that current font-smoothing algorithms introduce coloured pixels to create a visual illusion of smooth shapes). The larger symbol (B) is an enlargement of the smaller one, composed of the same number of squares in the same configuration and larger sizes. It is natural to say that B carries information about A: it allows someone in possession of B to infer the geometric properties of A, for example. But it is also natural to claim that A also carries information about B. If we had only partial access to A and B, we could reconstruct them both completely. Simply put, the point is that A and B correspond to one another, even if they differ in size. Given the truth of many geometrical inferences about both A and B, we tend to ignore the size difference and say that the similarity between A and B is independent of it.

This simple example reflects the everyday usage that is prevalent in our talk of information contained in pictures or photos. A somewhat more complex example of correspondence also underlies models produced by classifications in machine learning: To classify an object in a picture as a face, the artificial neural network relies on extracted features that correspond to (hopefully) all and only faces in pictures (Buckner [2018]). Similarly, upon receiving a digital copy of a article that was flagged as plagiarized from the web, I can reasonably assume that the original file was also plagiarized. Copies retain semantic values of their originals in virtue of similarity, even if they are somewhat distorted by noise along the way (for example, plagiarizers sometimes add typos to avoid detection). When this notion of correspondence is spelled out more precisely, it turns out that it underlies a number of uses of the term ‘information’, such as in scientific models or cognitive processing in spatial navigation.

The information talk concerned with A and B cannot be elucidated easily with the help of the probability-related notions of semantic information. For example, there seems to be no natural law that would explain the co-occurrence of A and B with a probability of one, which is required by Dretske's ([1982]) account of semantic information. A quick objection to this claim might, of course, be that Dretske's requirement is far too strong, hence one could substitute the original formula proposed by Dretske with a relaxed requirement: Shea ([2018], p. 76) defines correlational information in terms of conditional probabilities; for two items, a and b, a's being F carries correlational information about b's being G if and only if $P(Gb|Fa) \neq P(Gb)$. But pictures A and B might not correlate with one another as strongly as they would, for example, with the needs of developers of font-smoothing algorithms. Consequently, even if there

were some correlational information in Gb about Fa, it would be much less informative about Fa than about the needs of developers, even if we are not interested in them right now. They may even not correlate at all: The enlarged, jagged symbol is of little practical use outside the test bed of screen font-smoothing and it could never be produced for this particular symbol. Nevertheless, in this particular case the correlational information does exist (after all, I produced B from A). But this is not true in general. For example, a dictator's doppelganger may be used by a hostile intelligence agency to find out what the dictator looks like although the physical appearance of the doppelganger is statistically independent of the physical appearance of the dictator. We can use such statistically spurious similarities for sound inference. Thus, the doppelganger example does not yield itself easily to an elucidation in terms of probabilities of co-occurrence of two events. It is much more natural to focus on the similarity rather than the objective or subjective probability of the doppelganger's looks and the looks of the dictator occurring together to define this kind of semantic information. This kind of serendipitous similarity might even play a scientific role: The anecdotal inspiration for the discovery of the benzene ring was August Kekulé's dream of the ouroboros. There is obviously no correlational information between the ouroboros and the chemical structure of benzene but they do share an important organizational property. These examples imply that there is a kind of semantic information that is grounded in a two-place correspondence relation. If this is true, we can make sense of several intuitive judgements.

One is related to the issue of coding. Intuitively, there are cases in which A, a series of informational states, is about another series, B: When A encodes B. This occurs only when the encoding is a kind of correspondence relation. If the

encoding is sufficiently reliable, this can produce a chain of informational structures with similar informational contents. Thus, JPEG image files can contain some of the original semantic information inherent in a bitmap image. Moreover, JPEG files, being compressed in a lossy fashion, are only partially similar to original files, which also allows us to quantify the adequacy of the information contained in JPEG files.

Again, the intuition that encoding is deeply related to semantic information is difficult to cash out in terms of most existing accounts of semantic information. Indeed, some claim that the very notion of encoding makes scientific research on cognitive representation deeply flawed (Bickhard and Terveen [1995]; Brette [2019]). But most probability-based accounts of semantic information have virtually nothing to do with encoding. While one could defend a kind of semantic information that is grounded in Shannon information (Isaac [2017]), it would describe contents of such information always in terms of a vector of log probability ratios, which would be unwieldy in cases that are more readily understood in simpler structural (correspondence) terms.

All these intuitions can be elucidated by saying that if similarity obtains between two structures, one structure bears information about another. Before we go on, it is useful to list desiderata that our account of semantic information should fulfil:

- D1. The account should elucidate in virtue of what a piece of semantic information has satisfaction conditions.
- D2. The method of identifying pieces of information should be clear.

- D3. At least some pieces of semantic information should be evaluable as true.
- D4. There should be some quantifiable features of semantic information available to define its measures.
- D5. The account should be as general as possible.

Some comments are in order. While D1 and D2 seem commonly accepted, D3 makes it clear that the present account does not presuppose that false semantic information exists (the possibility of false information is denied by both Dretske ([1982]) and Floridi ([2005], [2007])), or perhaps there is only misinformation, mistaken for real information (which is, by definition, true).⁶ A defender of the concept of false information could stress that information could be more-or-less satisfied, and that there are at least two special truth values that semantic information could have (or more, if the defender is committed to multivalued logics or pluralism about logics). The present account attempts to remain neutral about the issue of falsity. D4 stresses that it should be possible to quantify at least something with regards to semantic information, but this issue is set aside for further inquiry. D5 motivates this account to remain as open as possible. Thus, this article defines generically a class of correspondence-based semantic information, which can be specified further as specific kinds for particular applications. The same desideratum motivates also D3, so as not to exclude various stances on falsity.

⁶ Neander ([2017]) calls this real information 'natural-factive information,' which she distinguishes from 'intentional information'. Following Neander, one could defend the view that falsity is present only in intentional information, which could be (at least partially) grounded in resemblance.

The current account does not require that semantic information be transmitted over any channel, sent by a sender or acquired by a receiver. The issues of control, communication, and function and value of information—in short, of the pragmatics of semantic information—are beyond the scope of the theory presented in this article.⁷ It is one thing to provide an account of semantic information that satisfies reasonable desiderata and quite another to offer a full-blown analysis of all possible uses of this information. Instead of providing a full theoretical analysis of pragmatics, two cases will be analysed in Section 5. However, the uses of information cannot be satisfactorily analysed in complete abstraction from what makes vehicles of information semantically satisfiable. The current contribution focuses only on semantic satisfiability of information.

3 Defining Informational Structure

The first important step in elucidating the notion of correspondence in terms of structural similarity is to define the structure of information vehicles, which is required by D2. Otherwise, the talk of correspondence would remain vague, which is indeed a major objection to the correspondence theory of truth as applied to language, called ‘the problem of matching’ (Rasmussen [2014], pp. 52–6). In the traditional correspondence theory of truth, as applied to natural language, the major difficulty is that propositions, at least *prima facie*, aren’t structurally similar to anything that they might correspond to. An analogous problem

⁷ It is arguable that pragmatics of semantic information is a crucial part of the account of mental and scientific representation. For a recent cogent defence of the idea that semantic information is explanatorily relevant in cognitive science only when exploited, see (Shea [2018]).

exists for the theory of scientific representations as based on similarity or isomorphism (Suárez [2003]). Thus, the correspondence theory of semantic information must provide a reply to the following question: How could semantic information be similar to anything it might correspond to? To answer it, information vehicles will have to be defined by identifying both (1) vehicles of information and (2) their structure. Only then can structural similarity be successfully introduced.

Vehicles of information will be understood, following the influential account of Gabor ([1946]), as quantifiable independent degrees of freedom of physical entities. Their number of degrees of freedom is the simplest measure of information, and information that can be measured this way is called 'structural' (MacKay [1969]). This number is strictly equivalent to the logical dimensionality of a complete description of a vehicle, that is, a description that provides all differences that make a difference for this vehicle. Less formally, it is simply a number of yes/no questions that one has to answer in order to yield a complete description of a vehicle.

The task of defining structure over such vehicles is far from trivial, because most vehicles we use have a complex structure that can be interacted with in multiple ways, and structural information content allows us only to quantify their maximum capacity. For example, one could have a vehicle with eight degrees of freedom. One obvious use of such a vehicle is to encode an 8-bit binary number on it, with every degree of freedom encoding a single bit of the number in the binary code. In such a setting, it is intuitive to think of this vehicle as composed of elementary atomic vehicles that correspond to individual bits. At the same time, this vehicle could be used to encode natural numbers in terms of

the number of 'on' states of the vehicle: If two states of the vehicle are on, the vehicle encodes '2', and so on. In other words, the encoding one uses for this vehicle is, to a large extent, arbitrary, and the number of degrees of freedom determines only the maximum structural information-bearing capacity, not the actual one. The vehicle can also support a more redundant encoding, while some degrees of freedom remain reserved for more reliability (parity checks, for example). Moreover, the machinery that reads the vehicle may not be causally sensitive to all of its physical degrees of freedom, which could further constrain its informational capacity.

Nonetheless, the crucial feature of the notion of structure is that it should be able to reflect the logical order. In the case of writing in the Latin alphabet, two inscriptions in English, 'tab' and 'bat', differ from each other with respect to the order of the individual letters. The same applies to any physical structure of information vehicles: In some settings, physical ordering in time or space should be distinguishable. The problem is that most approaches to defining structure remain open to a multiple of interpretations, which defy attempts to define ordering over vehicles. For example, if one were to define the notion of structure in set-theoretic terms of sets, members, and relations between members of sets, and so on, the problem is that this structure seems to collapse to cardinality, which does not go much further than the notion of maximum capacity. To see this, it is sufficient to note that if one defines the correspondence of two structures in set-theoretic terms, which is most commonly done in terms of a strict isomorphism, these structures are already strictly isomorphic as soon as they have the same cardinality (Newman [1928]).

One solution to this problem is to define the notion of structure in terms of non-extensional relation theory, such as homotopy, which could situate the homotopy theory as foundational (Ladyman and Presnell [2018]). Another is to think of the requirement more abstractly: By 'structure', we mean a certain ordered complex entity, which, in the case of information, should be physically discernible. Two physical tokens have the same structure as long as they can be classified as subsumed under the same type. This observation was used by Barwise and Seligman ([1997], p. 28) in their influential work on information flow. This is how they define their notion of classification:

A classification $A = \langle A, \Sigma_A, \varepsilon_A \rangle$ consists of a set A of objects to be classified, called *tokens* of A , a set Σ_A of objects used to classify the tokens, called the *types* of A , and a binary relation ε_A between A and Σ_A that tells one which tokens are classified as being of which types.

Interestingly, one can classify also a set of tokens that are actually a set of classifications. In other words, some classifications can be stacked on top of one another. In this way, nested complex structures can be defined, as these types may require some order to be retained. But the way in which it is retained remains to be specified in how we defined the binary relation ε_A . These could also be defined in terms of homotopy theory, if one wishes, but if the structure is actually insensitive to ordering, one could specify it set-theoretically (for example, when one uses it in an application that treats it as a grab bag with its elements, or tokens, to be taken in any order).

The kind of similarity we will be interested in when defining correspondence-based information is based on the relation between classifications. But defining relations of the correspondence relation is just the first step in dealing with the problem of matching.

4 Defining Similarity as Accuracy

At least since Goodman's critiques ([1968], [1972]), similarity is approached by philosophers with appropriate caution. In a nutshell, the problem is that as long as one does not specify in which respects two things, a and b , are to be held similar, one can always ascertain that they resemble each other in an indefinite number of ways. This trivializes an unconstrained notion of similarity.

Even worse, the geometrical model of similarity, which was influential in the 20th century (Carnap [1928]; Coombs [1954]; Shepard [1958]), raises a number of problems: In this model, not only is similarity context-dependent, as noted by Goodman, but it also does not seem to correspond to human similarity judgments (Tversky [1977]). But instead of surveying these problems and various alternative formal accounts of similarity (for a recent accessible survey, see Hahn [2014]), the focus here will be on the role of similarity in cases like that presented in Figure 1.

Similarity licenses us to perform valid inferential operations. Take two infinity symbols, A and B . If they are similar in some respect r , then knowing that A has a property that can be subsumed under r allows us to infer that B also has this property (at least to a degree). Swyer ([1991]) has dubbed this kind of reasoning 'surrogate': It allows us to draw inferences about B when we only have access to A . In this case, it is intuitive to say that B carries information

about A . Thus, the kind of similarity we are after is the one involved in surrogative reasoning. To sum up, the generic notion of similarity, as involved in correspondence-based semantic information, can be defined as:

$$r\text{-sim}(A, B) \equiv [r(A) \rightarrow r(B)].$$

By $r\text{-sim}$ we mean a two-place predicate that denotes a relation of similarity under respect r . Note that the right-hand side of the formula contains only a material implication rather than a biconditional, between $r(A)$ and $r(B)$. This is because we do not want to exclude asymmetrical and antisymmetrical accounts of similarity.⁸


This can be further generalized.

4.1 Similarity as infocorrespondence

Barwise and Seligman use the notion of classification to introduce their notion of infomorphism, which is then used to define information flow. Infomorphism is a pair of f functions $\langle f^\wedge, f^\vee \rangle$. For two classifications, $\mathbf{A} = \langle A, \Sigma_A, \vDash_A \rangle$ and $\mathbf{B} = \langle B, \Sigma_B, \vDash_B \rangle$ an infomorphism f obtains if and only if $f^\vee(b) \vDash_A a \equiv b \vDash_B f^\wedge(a)$ for all tokens b of B and all types a of A (see Figure 2).

⁸ Some kinds of similarity, like that involved in lossy compression of data, seem to be best analysed in such asymmetrical terms.

Figure 2.



Barwise and Seligman provide a simple example of an infomorphism (related to an example of sentence segmentation we consider in Section 5.1). It starts with two classifications, *Punct* and *Sent*. The tokens of *Punct* are English punctuation marks (‘,’, ‘.’, ‘!’, and so on). They can be classified as COMMA, PERIOD, and so on. The tokens of *Sent* are English grammatical sentences of types: DECLARIVE, QUESTION, and OTHER. Now, the infomorphism $f : \mathbf{Punct} \rightleftarrows \mathbf{Sent}$ is defined such that f assigns to each sentence token its own terminating punctuation mark. Thus, f assigns DECLARIVE to EXCLAMATION MARK and PERIOD, QUESTION to QUESTION MARK, and OTHER to other types of Punct (Barwise and Seligman [1997], p. 73).

Intuitively, for information to flow in a channel, it should be preserved along the way. This property of preservation can be nicely accounted for in terms of infomorphism. But this property is also related to correspondence-based information: In the case above, A carries information about B and retains all of it. When defending his notion of surrogative reasoning, Swoyer ([1991], p. 472) stresses that structural representation occurs just in case properties and relations are preserved and counter-preserved. This can be also expressed in terms of infomorphisms. In the case of infomorphism occurring between A and B , the information vehicle A tells us ‘the whole truth and nothing but the truth’ about B (Swoyer [1991], p. 472).




Figure 3.

But similarity comes in degrees, which means that we may introduce a generalized notion of ‘infocorrespondence’ between A and B in an analogous manner, as a pair of functions. For two fuzzy classifications, $A = \langle A, \Sigma_A, \varepsilon_A \rangle$ and $B = \langle B, \Sigma_B, \varepsilon_B \rangle$, f obtains if and only if $f^\vee(b) \varepsilon_A a \equiv b \varepsilon_B f^\wedge(a)$ for at least some tokens b of B and at least some types a of A (the definition of a fuzzy classification is offered below). Infocorrespondence, in contrast to infomorphism, requires only that the biconditional obtains for at least some tokens b of B and some types a of A and the classifications in question could be fuzzy. By quantifying the number of tokens and types one could measure the degree of similarity that occurs between A and B . For example, in Figure 3, we can see that B is distorted in its lower right corner and does not fully correspond to A .

In Figure 3, the structural information of A is still retained to some degree by B , but it may no longer tell us the whole truth. However, our job is not yet done. Figure 4 features another case of similarity: The bitmap symbol A turns out to be similar to a larger symbol B , which does not appear jagged. In some ways, Figure 4 may seem a better example of similarity than Figure 1 or 3. This is because it seems reasonable to idealize away all of the jagged rectangular




Figure 4.

lines, invisible to the human eye on a smaller scale, as being in reality cases of rounded lines. Technically speaking, then, we could classify these rectangular lines as round by allowing for fuzzy (approximate) classification. In other words, our classifications could be based on fuzzy sets (Zadeh [1965]): Tokens could belong to sets to some degree. And, to a degree, the dots could be described as approximate members of circular lines. Note that without this relaxation of classification, we could not think of vectorized images as similar to bitmap images. This is justified idealization, which consists of removal of unwanted information from source information.

In some ways, Figure 4 may seem a better example of similarity than Figure 1 or 3. This is because it seems reasonable to idealize away all of the jagged rectangular lines, invisible to the human eye on a smaller scale, as being in reality cases of rounded lines. Technically speaking, then, we could classify these rectangular lines as round by allowing for fuzzy (approximate) classification. In other words, our classifications could be based on fuzzy sets (Zadeh [1965]): Tokens could belong to sets to some degree. And, to a degree, the dots could be described as approximate members of circular lines. Note that without this relaxation of classification, we could not think of vectorized images as similar to

bitmap images. This is justified idealization, which consists of removal of unwanted information from source information.

In general, then, while infomorphism seems the upper bound of similarity, the usual case may be infocorrespondence based on a fuzzy classification, in which tokens are classified as belonging to types to a very low degree. The lowest bound is set, of course, when no information is retained between classifications, even when all kinds of idealizations and distortions are included. But more exacting models of similarity can be then added to define various subkinds of fuzzy infocorrespondence by further constraining the definition of a given infocorrespondence (remember we stipulated that we wish to remain as general as possible). In other words, it should be possible to analyse the relationship between classifications in terms of some specific model of similarity.

For example, the feature-matching model introduced by Tversky ([1977]) would not only require us to define features that are shared by A and B , but also to define their differences (which jointly define the contrast between A and B). In the extreme case of infomorphism, there are no distinctive features that differ between the two classifications, which ensures a very high (and symmetric) degree of similarity. Obviously, there is some contrast in Figure 3: B does not retain all information about A .

Similarly, one could apply the constraints above to think of classifications in terms of random variables A and B whose joint distribution $P_{(A, B)}$ and marginal distributions P_A and P_B are known. In this case, a simple natural measure of similarity is mutual information (MI), which measures the mutual dependence between the A and B . MI defines the degree of similarity of the joint dis-

tribution to the product of the marginal distributions. The corollary is that whenever MI is used to measure dependence between certain classifications, there is semantic information inherent in both of these classifications about each other.

The details of how to apply the idea in particular models of similarity need not concern us here because our task is not to define one particular species of correspondence information but to provide a useful elucidation of the general concept. There are multiple other approaches to defining similarity for various purposes, and there is no *a priori* reason to choose any one over another (for a short review of work on similarity measures, see Ashby and Ennis [2007]).⁹

This concludes our solution to the problem of matching, introduced in the previous section: The infocorrespondence relationship holds between classifications. It becomes immediately clear that classifications, in the sense that was introduced by Barwise and Seligman, may match one another, even if one is sceptical of whether propositions may match the structure of physical reality.

Let me stress one feature of the present account: It does not constrain in any way classifications that could stand in an infomorphism or infocorrespondence relation. These could be arbitrarily picked, if so desired. Nonetheless, in complete theories of representation, one usually introduces additional constraints on such classifications, for example to avoid the issue of too many correspondences (see Section 6.2, below). This is because mental and scientific representation involve more than mere semantic information, which implies that the accounts

⁹ For example, self-similarity is not always considered maximally informative, although one could, in principle, think of automorphism of A with A in terms of infomorphism. However, in the definition above, it is left open whether or not $A = B$. In some applications of correspondence information, such as information retrieval, one does not expect that the search key should retrieve documents that are exactly the same as the search key (Amigó *et al.* [2017]).

of mental and scientific representation must provide their own accounts of inaccuracy and mistargeting (Suárez [2004]). Mistargeting is a mistake in which one supposes the target of a representation to be something that it actually does not represent. The arbitrary choice of any two classifications might lead to mistargeting and the present account, for generality purposes D5, does not offer any special solutions, but usually, causal relations or overall reliability of inference over time could be used to justify the claim that a given classification **B** is the target of semantic information **A**. For example, **B** might have been produced by a process that is structured with respect to **A** (causal relation) or over time, serendipitous inferences using **B** turn out to be correct about **A** (as in the doppelganger example). Only when we know that **A**'s target is **B**, we can correctly determine its level of accuracy or detect the error inherent in it.¹⁰ But we can still provide an account of truth of semantic information, which is our next task.

4.2 From accuracy to truth value

Similarity between structures allows one to infer properties of one structure, given another structure. This is what is implied by the fact that **A** carries information about **B** by corresponding to **B**. In the extreme case, the correspondence is simply an infomorphism, but some violations of conditions of infomorphism are also possible. Importantly, as long as infomorphism holds, an inference about **B** as based on **A** will be sound, as long as **A** is a correct classification. In

¹⁰ Error detection may require comparing several classifications $A_1 \dots A_n$ whose target is the same classification **B**; only then can we detect inconsistency among $A_1 \dots A_n$, which is a mark of misrepresentation. See the full account in (Bielecka [2019]).

other words, we may validly infer properties of **B**, given **A**. In such a case, **A** is true about **B**.

Thus, we may define true semantic information in the following way:

SEM-TRUE: If infomorphism obtains between **A** and **B**, then **A** is true about **B** and **B** is true about **A**.

In other words, the satisfaction conditions of **A** about **B** are the same as the conditions under which the infomorphism between both holds. The truth-value—truth—is ascribed to a classification **A** depending on whether the infomorphism between **A** and **B** holds. This shows that the present account satisfies D1.

One could be tempted to define falsity in the same way. Although Dretske ([1982]), Israel and Perry ([1990]), and Floridi ([2005], [2007]) argued that there is no such thing as false information, and false information is simply misinformation, their position is not uncontroversial (Scarantino and Piccinini [2010]). The present account remains neutral with regard to how or whether to define falsity of correspondence information. This is enough to satisfy D3, as introduced above, and neutrality is inspired by D5. This issue is not crucial to our understanding of the basic notion of semantic information as based on correspondence. Nonetheless, before we turn to non-classical truth-values, let me briefly state the issues regarding the falsity of semantic information.

One could suggest that if there is no single sound surrogative inference (not even a rough one, based on a fuzzy correspondence) about **B** from **A**, then **A** is false about **B**. This is a very strong requirement given that we do not constrain

a priori the way one performs basic classifications. Yet for arbitrary classifications A and B , one could wonder why, then, A is false about B , and not something else, if there is nothing true of B that could be inferred from A . There are at least three possibilities in such a case. First, one could bite the bullet and claim that A is false about any X if and only if A does not stand in any infocorrespondence with X . Alternatively, one could assert that in such a case, A is simply not about X , without being false at all (because we do not account for mistargeting in the present theory of semantic information). Third, one could assert, like Dretske and Floridi, that, in fact, A misinforms one about X only if additional pragmatic considerations obtain (such that a given piece of information is designed to carry information about X but fails to do so due to some abnormal condition). Thus, misinformation would be possible only in the case of mistargeting, which implies that the full account of misinformation would be possible only in terms of representation rather than mere semantic information.

Suppose now we go beyond classical logic with its two truth-values. Given that similarity is (usually) understood as graded, it is natural to appeal to multi-valued logics in order to model the accuracy of a statement about the correspondence between A and B . One such logic is fuzzy logic (Zadeh [1988]), which is a natural ally of the fuzzy set theory used to define fuzzy classifications. In such a case, the degree of truth would closely track the degree of similarity, thus making semantic information intuitively match quantifiable aspects of similarity D4:

SEM-TRUTH-DEGREE: The truth-value of an information vehicle classified as A that stands in an infocorrespondence C with classification B is determined by the degree of similarity with which C holds.

Naturally, if the logic in question does not have truth-values defined in terms of continuous values, and one's measure of similarity does, one could apply well-known techniques of digitization.

5 Applications

In this section, two simple cases will be introduced. They were chosen mostly because of the inspiration of the current proposal by the accounts of scientific and mental representation. Thus, the first concerns scientific models; the similarity-based conception of scientific representation is one of the dominant positions in the debate (Bartels [2006]; Chakravartty [2010]; Weisberg [2013]). The main opposing view, the inferential account (Suárez [2004]), however, is not as distant from the position defended here as it may initially seem. The inferential conception assumes that representation has two basic features: (1) it is not arbitrary and, by having some properties, it points to its target; (2) it allows competent and informed agents to draw specific inferences about the target. The second case concerns structural representations (hereafter, S-representations), which are considered fruitful in cognitive (neuro)science (Ramsey [2007]; Gładziejewski and Miłkowski [2017]; Thomson and Piccinini [2018]).

5.1 Scientific models

In his account, Weisberg ([2013]) defends the claim that scientific representation is based on similarity, which he analyses in terms of Tversky's ([1977]) feature model. In particular, Weisberg distinguishes two categories that are supposed to be modelled: attributes and mechanisms. These are required to provide his account of model–world relations in terms of weighted feature-matching. Without going into the details of his analyses of particular cases, it suffices to point out that one can simply replace his binary similarity relation with infocorrespondence (as defined above) to arrive at the correspondence account of these cases.

However, not all models concern mechanisms and their attributes. They also need not be explanatory. For example, one problem in computational linguistics is to provide a model of how a written text is segmented into constituent sentences (Mikheev [2005]). The function of this model is to process the text rather than to provide an explanation of how it is done. Yet we can still think that the model can describe human results in the sentence segmentation task. The task is not entirely straightforward owing to the double function of the period, which is added at the end of abbreviations in many languages. The quality of the model is usually measured using two classical metrics: recall and precision. 'Precision' is the ratio of correctly recognized sentence breaks to total sentence breaks found by the model, and 'recall' is the ratio of correctly recognized sentence breaks to total actual sentence breaks in the text. In general, these two are used to assess how the model fits human performance—or whether it does at all. The model's quality can be computed easily in this case because both classifications

are supplied in the digital form (an unsegmented stream of characters versus the text segmented by human annotators). The stream of text (particular tokens) is processed to provide sentence types. Quite straightforwardly, the sentence-breaking task is understood in computational linguistics in terms of classification problems, as the core of the problem is to recognize whether a token (a string of characters) should be subsumed under a single type (full sentence), unlike as in the infomorphism $f: \mathbf{Punct} \rightleftarrows \mathbf{Sent}$ in Section 4.1. Thus, as long as the model provides performance similar to human performance on text A, it contains accurate semantic information about human performance on A. The accuracy is measured using precision and recall (both are usually over 95% for modern off-the-shelf models).

The opponents of the similarity account of modelling stress that the inferential function of models is more crucial than mere structural relations (Suárez [2004]). Nonetheless, one could understand correspondence-based models as specific cases of representation accounted for by the inferential conception. While a large class of models contain semantic information in the correspondence sense, the present account allows for scientific models that cannot be analysed fruitfully in terms of similarity. It does not deny the inferential function; it also stresses the importance of the non-arbitrary character of representation. Indeed, the function of models is to warrant surrogative reasoning, whose quality can be quantified by using a particular similarity measure. The value of computational models of text segmentation lies in the fact that they can be used to segment a string of characters. In other words, they are not mere static descriptions of human performance but can predict this performance for new data. They do this algorithmically (usually by using a set of rules, statistical models, or

deep neural networks); thus, in a sense, these models are inferential machines. Nonetheless, the results of their inferences are not merely valid; they are sound, or approximately true. Hence, the talk of semantic information is fully justified.

The present account of semantic information is not supposed to replace pragmatic considerations that are required to properly understand and use scientific models, particularly those that require a competent user to draw inferences using them. In the case of the sentence segmentation model, the user must be able to prepare the input stream of characters for processing (knowledge of the language and writing system in which the stream is written is indispensable in this process), run the segmentation software (possibly from a command line), and interpret the output stream of characters.

5.2 Cognitive maps

One example of S-representations in contemporary cognitive (neuro)science is cognitive maps in rodents, in particular, as instantiated in place cells in rats. Not only has the representational hypothesis introduced by Tolman ([1948]) been extremely heuristically fruitful (Bechtel [2016]), but it also provides deep explanations of spatial navigation mechanisms. Accurate semantic information about spatial location determines the level of success in navigation tasks (for an accessible introduction to this kind of representation, see Shea [2018], pp. 113–6).

In particular, rodents plan future paths, which is reflected in the future-oriented navigational activity of place cells in the hippocampus in the brains of rats. This activity was directly observed in an elegant experiment (Pfeiffer and

Foster [2013]). As it turns out, rats pause before taking a journey. During that pause, place cells emit sharp-wave-ripple events: irregular bursts of brief (100–200 ms) large-amplitude and high-frequency (140–200 Hz) activity. These are distinct from regular spikes in place cells. Using an algorithm proposed earlier for decoding similar neural events (Davidson *et al.* [2009]), Pfeiffer and Foster were able to show that place cells are used to represent the future journey of the rat to the location of a previously observed reward. Another experiment showed that these cells can fire in response to food in an observed location (Ólafsdóttir *et al.* [2015]). Indeed, these events reflect future behavioural paths.

In this case, the bearers of S-representations are the bursts of neural activity in place cells, which correspond to possible future locations of the animal. A number of researchers could not find evidence for the topographical organization of place cells, which would license the talk of isomorphism between the firing of place cells and the physical space (for a critical review, see França and Monserrat [2019]). In the topographical organization, neighbouring neurons tend to represent ‘neighbouring stimuli’, for example, adjacent tones in the auditory cortex or adjacent locations in the field of view in the visual cortex. However, a recent hypothesis, defended by França and Monserrat, is that these cells are organized topographically, but the map is not about the physical space; instead, it is about an abstract multidimensional space that contains receptive fields of place cells. These cells seem to encode multisensory perceptual cues for individual places and integrate the animal’s goals and attentional demands. While this hypothesis is not yet the mainstream account of the hippocampus function, it incorporates the existing evidence about tuning that is observed when recording place cell activity. If it is true, then place cells are literally a

map, but not a map of a physical space. But if it is false, then place cells are locally topographically organized to a much smaller degree, which depends on the animal's context. The context of future-oriented navigation, however, is somewhat special: Shea ([2018], p. 115), while assuming that place cells (when activated online) do not form a topographical map, observes that the co-activation present in the future-oriented navigational activity involves encoding the spatial structure, which is a kind of S-representation.

Thus, in both cases, with the topographical organization in their online activity or not, place cells are bearers of semantic information in the sense defended in the current article: The activity of the hippocampus encodes future locations of the animal. By using dedicated software, researchers can distinguish individual units (parts of the classification) whose activity can then be analysed to recover spatial information. This is done by dividing possible positions into 2 cm bins and calculating the position of tuning curves in place cells—which underlie the basic operation of place cells in navigation. Then, spike trains (during short-wave-ripple events) and tuning curves are processed by a decoding algorithm to estimate rate position.¹¹

In general, decoding algorithms are used in neuroscience to recover the relationship of neural activity to whatever the neural activity is hypothesized to

¹¹ The probability of the animal's position (pos) across M total position bins given a time window (τ) containing neural spiking (spikes) is $(\Pr(\text{pos}|\text{spikes}) = U / \sum_{j=1}^M U)$, where

$$(U = (\prod_{i=1}^N f_i(\text{pos})^{n_i}) e^{-\tau \sum_{i=1}^N f_i(\text{pos})})$$

and $f_i(\text{pos})$ is the position tuning curve of the i -th unit, assuming independent rates and Poisson firing statistics for all N units and a uniform prior over position (Pfeiffer and Foster [2013], p. 80).

reflect. While they are based on various mathematical methods, the purpose of the algorithms is not to explain how that semantic information is processed by the brain, but to discover the relation underlying a particular infomorphism (or infocorrespondence).

6 Possible Objections

In this section, three possible objections to the proposed account will be analysed in turn. One could also raise other objections by noting that the account resembles the correspondence theory of truth, which some consider to be seriously flawed. Here, it will be assumed that this theory can be salvaged, which implies that such objections are not of the highest importance (for a recent defence of the correspondence theory of truth, see Rasmussen [2014]).

6.1 Similarity versus covariation

It could be claimed that structural similarity boils down to covariation. But this is not the case. For example, a map of Rome may carry correspondence-based information about the geometrical layout of ancient buildings, but it does not covary with how the city changes over time. Crucially, correspondence-based information can remain static, which means that the information vehicle remains unchanged.

Thus, while covariance between two items does imply that there is at least some correspondence between them, which implies some kind of similarity, not all kinds of similarity imply covariance. Likewise, not all kinds of similarity are based on nomological necessities like those required by Dretske ([1982]) in his

influential account of semantic information or on correlational information in Shea's ([2018]) sense.

A related point was discussed at length in the debate over cognitive representation. Indicators (sometimes also called receptors), understood as representations relying on covariance or correlation, are contrasted with S-representations whose contents depend on similarity. Against Ramsey, who claimed that unlike S-representations, indicators are not genuinely representational, Morgan argued that indicators are S-representations (Morgan [2013]). In response to Morgan, Gładziejewski and Miłkowski ([2017]) defended the view that S-representations are distinct from indicators. Their position was subsequently criticized by defenders of indicators as contentful representations (Rupert [2018]; Nirshberg and Shapiro [2020]). It seems to be the case, however, that while one could understand at least some S-representations as carrying correlational information, not all indicators—taken as ensembles or arrays—are structure-preserving representations. Thus, in an important respect, Morgan is mistaken in thinking that there's '*simply no distinction* between receptors and structural representations' (Morgan [2013], p. 222). For example, Shea provides an example of various vervet monkey calls that are not structured in such a way that there is a clear ordering or relationship among the calls that would play a representational role (Shea [2018], p. 119). Moreover, as we observed in Section 1, some similarities may be serendipitous and therefore cannot be understood in terms of correlational information (witness the doppelganger and ouroboros examples). At the same time, an account of mental representation in structural terms, which also relies on causal relation with representation targets, is both compatible with the present account and much current psychological theorizing (for a defence of

a causal-structural account of mental representation, see for example (Isaac [2013])).

6.2 Too many correspondences

It is a well-known fact that all attempts to define similarity in a structural fashion open a Pandora's box: Basically anything can be shown to be similar to anything else (Goodman [1951], [1972]). A similar argument was proposed by Watanabe ([1972]) in his Ugly Duckling Theorem, which states that the ugly duckling can be shown to be as similar to a swan as two ducklings are to each other. More generally, it shows that without restrictions, any two objects can share an arbitrarily large number of properties. Similarly, structural correspondence can be shown to hold trivially between any two structures of the same cardinality (Newman [1928]).

In other words, if similarity is so cheap, why bother with semantic information grounded in similarity? This problem was also noted for S-representations in cognitive science (Ramsey [2007], pp. 93–6). The answer lies in how the problem is posed: When we are interested in surrogative reasoning, we are not interested in just any kind of resemblance between two classifications. We are interested in those that would allow us to infer interesting facts about one thing based on the information contents of another. In such a case, the classification must be grounded before one starts to evaluate whether semantic information is accurate by assessing the degree of similarity. This is usually done for any domain in which similarity is used for such inferences. For example, if a deep neural network is supposed to classify a set of images and determine

whether they contain a picture of a known criminal, the picture must first be specified.

Note, however, that this is a pragmatic consideration of how one uses this kind of semantic information (and this is what Ramsey suggests in his solution). In principle, the present account does not preclude use of the same information-bearing structure to carry information about things that are totally disparate. But this is not a bug: This is a feature.

Consider a simple example. A known trick in cryptography is to use a salient feature to embed a message in another message without any encryption, for example, by embedding the face of a spy in a pornographic photo. This is known as steganography. In such a situation, the same information vehicle carries multiple messages in its physical structure: one overtly visible (and used for distraction), and another missed by most observers. In such a case, one need not decide beforehand that the huge number of degrees of freedom of picture must bear similarity to a single target by being part of the same classification.

More generally, the same vehicle may be used to convey several messages at the same time, and multiple receivers could use its properties in multiple ways by ignoring some properties of the vehicle as irrelevant (virtually cross-classifying its properties). The current account has the virtue of allowing this to happen. The same pertains to surrogative inferences on correspondence-based information: These need not be practically useful or even available for any user of information. It is only required that they be logically possible.

6.3 No propositional content

Finally, it could be pointed out that the current theory does not provide the propositional content of the semantic information at hand. The correspondences hold between classifications, after all, and not between propositions and states of affairs.

While it is certainly true that structural correspondence is not considered in particular to hold between linguaform entities and whatever these entities may describe, but between classifications, every part of the correspondence involved is truth-evaluable, which is considered the essential property of propositions (on most accounts of what propositions are). For example, if any part of the icon in Figure 1 is garbled, as in Figure 3, the correspondence holds only partially. Thus, there are parts of the icon that are no longer true of another icon. This means that one can ascribe a truth value to the classification in question. Still, the correspondence relations in Figures 1 and 3 are relatively simple. But it is, in principle, possible to analyse this correspondence in terms of rules (in this case, a geometric rule of scaling is enough). These rules could have more complex roles for constituents of classifications, in particular if they are also complex structures. No principle precludes the use of this framework to formulate something akin to Tractarian pictorial semantics for a formal language. But it is not the aim of the current account to defend such Tractarian semantics; this proposal is mostly concerned with S-representations in general.

One caveat is, however, in order: Not all kinds of correspondence-based semantic information may support classical inference operations. For example, iconic representations, such as pictures, cannot express the simple notion of ne-

gation (Johnson-Laird and Byrne [1991]). In general, connectives are not univocally expressible in a pictorial manner. Thus, operations such as joining or dividing pictures cannot be understood in terms of logical connectives such as conjunction introduction or simplification. This is why complex mathematical diagrams support inferential operations in virtue of a combination of pictorial and symbolic representations, whose cognitive functioning is still poorly understood, even for such foundational systems as Euclidean geometry (Hohol [2019]). As a result, applying correspondence-based semantic information to natural or artificial languages requires special attention to inferential operations, if classical logical rules are to be supported in surrogate reasoning based on this information. In particular, the introduction of arbitrary notation for propositional structure, including logical connectives, seems inevitable.

7 Conclusions

In this article, a novel account of semantic information, analysed in terms of correspondence, was proposed. The account elucidates the satisfaction conditions of information in terms of conditions under which a similarity between classification holds. When complete similarity obtains, the semantic information is evaluated as true. It was stressed that there are multiple ways one can analyse similarity between classifications, but that the upper bound of similarity can be understood in terms of infomorphisms. In other cases, classifications may be based merely on the membership of tokens in fuzzy sets, and the informational relation between sets can be partial. This makes this account general

but not overly liberal because it is required that semantic information be useful for surrogative reasoning.

Further work is required to analyse in detail how various models of similarity give rise to various kinds of semantic information, in particular, with reference to the way information is measured. It was indicated, for example, that depending on the domain of use, one can think of self-similarity as informative or not; in information retrieval, one does not think of search key terms as informative about themselves. Therefore, specific domains may require specific approaches to measuring semantic information as defined in terms of correspondence.

At the same time, most extant approaches to similarity, from computer science through information science to cognitive psychology, offer quantitative methods to measure similarity, which allows one to quantify the accuracy of semantic information in question. However, it is a matter of further research to provide measures of (particular kinds of) semantic information.

When motivating the enterprise of defining semantic information in terms of correspondence, we stressed that while it is related to previous work on mental and scientific representation, it focuses on its basic ingredient, which has hitherto remained only implicit. Let us now elucidate why it is helpful to start with semantic information first, and then build a theory of representation on its grounds. One of the major objections against the accounts of scientific representation as based on similarity (and isomorphism) is that similarity is symmetric and reflexive (Suárez [2003]). According to the definition of infocorrespondence, it is non-symmetric and non-reflexive, thus it allows symmetric and asymmetric, as well as reflexive and irreflexive kinds of similarity. Notice, however,

that one could also define semantic information in terms of a symmetric relation of covariance, which would lead to the same problem. So, would a proponent of scientific representation be in trouble if her account of representation relied on a kind of information which is based on structural similarity, and whose relational properties differ from those of representation? Not really. Semantic information may be symmetric, but the mental or scientific representation that relies on this information need not be, because representation can be considered not a binary but at least a ternary relation, which includes the user of the representation, such as *x-uses-semantic-information-y-to-represent-entity-z*. In more detailed analyses of representational function, such as in teleosemantics (Millikan [1984]), multiple users (senders and receivers, or producers and consumers) are usually introduced, along with other aspects of the information flow, which technically make representation *n*-ary, where *n* is typically larger than five. The logical properties of semantic information, which grounds the use of information vehicles to target and describe entities, need not carry over to logical properties of this *n*-ary relation of representation.

The correspondence-based semantic information justifies the practice of relying on similarity between pieces of information for inferential purposes, even if these pieces are statistically independent. It may also be used to defend the practice of experimental studies of neural encoding as studies of the contents of neural representations (for a recent review, see Kriegeskorte and Kievit [2013]). Thus, it is descriptively valid in two cases that are sometimes difficult to cash out in terms of covariance, correlation, or nomological necessity.

The proposed account may have an air of paradox to it: On one hand, it is widely assumed in our everyday information talk of information as true and

inherent in computer files, not to mention the rich discussion of scientific and mental representation in philosophy and cognitive (neuro)science, but on the other, it has never been defended before as an approach to semantic information. The purpose of this article is to remove the air of paradox and provide the first steps toward further research on the semantics of correspondence of information. Its novelty lies in allowing various kinds of similarity to ground conditions of satisfaction and in providing a general solution to the problem of matching. It also provides an analysis of the basic ingredient of structural representation, both mental and scientific.

The present account suggests a division of labour between theories of semantic information and theories of representation: While an account of semantic information should provide answers to foundational questions, such as the matching problem, which is required by D2 and D3, the account of representation may rely on the foundational work, provide further constraints, and clarify specific uses of particular kinds of similarity in a given domain. The previous work did not introduce this division of work for structural representations, and foundational work—regarding the kinds of possible similarity relationships or the matching problem—was usually kept minimal, so as to enable further discussion of issues related to scientific or mental representation. The current proposal may, therefore, drive further exploration of various kinds of representation as relying on various types of correspondence-based information.

Acknowledgements

This article is dedicated to the memory of John Collier. The author wishes to thank Witold Hensel, Tomasz Korbak, Paweł Gładziejewski, Tadeusz Ciecierski, Andrzej Gecow, Julian Zubek, Joanna Rączaszek-Leonardi, Joanna Guzowska-Loeb, Jacek Malinowski, two anonymous referees, and the editor, as well as audience members of the seminar of the Human Interactivity and Language Lab (University of Warsaw) and conference Философия, логика, култура (Kiten, Bulgarian Academy of Sciences) for invaluable feedback. The work in this article was funded by a National Science Centre (Poland) research grant under the decision 2014/14/E/HS1/00803.

*Institute of Philosophy and Sociology
Polish Academy of Sciences
Warsaw, Poland
mmilkows@ifispan.edu.pl*

References

- Amigó, E., Giner, F., Gonzalo, J. and Verdejo, F. [2017]: 'An Axiomatic Account of Similarity', *SIGIR '17: Proceedings of the 40th International ACM SIGIR Conference on Research and Development in Information Retrieval*, available at <doi.org/10.1145/3077136.3084369>.
- Ashby, F. G. and Ennis, D. M. [2007]: 'Similarity Measures', *Scholarpedia*, **2**, p. 4116.
- Bartels, A. [2006]: 'Defending the Structural Concept of Representation', *Theoria*, **21**, pp. 7–19.
- Barwise, J. and Seligman, J. [1997]: *Information Flow: The Logic of Distributed Systems*, Cambridge: Cambridge University Press.

- Bechtel, W. [2016]: 'Investigating Neural Representations: The Tale of Place Cells', *Synthese*, **193**, pp. 1287–321.
- Bickhard, M. H. and Terveen, L. [1995]: *Foundational Issues in Artificial Intelligence and Cognitive Science: Impasse and Solution*, Amsterdam: North-Holland.
- Bielecka, K. [2019]: *Błądzą, więc myślę: co to jest błędna reprezentacja?*, Warszawa: Wydawnictwa Uniwersytetu Warszawskiego.
- Brette, R. [2019]: 'Is Coding a Relevant Metaphor for the Brain?', *Behavioral and Brain Sciences*, **42**, p. E215.
- Buckner, C. [2018]: 'Empiricism without Magic: Transformational Abstraction in Deep Convolutional Neural Networks', *Synthese*, **195**, pp. 5339–72.
- Callender, C. and Cohen, J. [2006]: 'There Is No Special Problem about Scientific Representation', *Theoria*, **21**, pp. 67–85.
- Carnap, R. [1928]: 'The Logical Structure of the World; Pseudoproblems in Philosophy', Berkeley, CA: University of California Press.
- Chakravartty, A. [2010]: 'Informational versus Functional Theories of Scientific Representation', *Synthese*, **172**, pp. 197–213.
- Coombs, C. H. [1954]: 'A Method for the Study of Interstimulus Similarity', *Psychometrika*, **19**, pp. 183–94.
- Craik, K. [1943]: *The Nature of Explanation*, Cambridge: Cambridge University Press.
- Cummins, R. [1996]: *Representations, Targets, and Attitudes*, Cambridge, MA: MIT Press.
- Cummins, R., Blackmon, D., Byrd, D., Lee, A., May, C. and Roth, M. [2006]: 'Representation and Unexploited Content', in G. McDonald and D. Papineau (eds), *Teleosemantics*, New York: Oxford University Press.
- Davidson, T. J., Kloosterman, F. and Wilson, M. A. [2009]: 'Hippocampal Replay of Extended Experience', *Neuron*, **63**, pp. 497–507.
- Dilworth, J. [2008]: 'Semantic Naturalization via Interactive Perceptual Causality', *Minds and Machines*, **18**, pp. 527–46.
- Dretske, F. I. [1982]: *Knowledge and the Flow of Information*, Cambridge, MA: MIT Press.

- Eliasmith, C. [2005]: 'A New Perspective on Representational Problems', *Journal of Cognitive Science*, **6**, pp. 97–123.
- Floridi, L. [2005]: 'Is Semantic Information Meaningful Data?', *Philosophy and Phenomenological Research*, **70**, pp. 351–70.
- Floridi, L. [2007]: 'In Defence of the Veridical Nature of Semantic Information', *European Journal of Analytic Philosophy*, **3**, pp. 31–41.
- França, T. F. A. and Monserrat, J. M. [2019]: 'Hippocampal Place Cells Are Topographically Organized, but Physical Space Has Nothing to Do with It', *Brain Structure and Function*, **224**, pp. 3019–29.
- Gabor, D. [1946]: 'Theory of Communication, Part 1: The Analysis of Information', *Journal of the Institution of Electrical Engineers Part 3*, **93**, pp. 429–41.
- Giere, R. N. [2004]: 'How Models Are Used to Represent Reality', *Philosophy of Science*, **71**, pp. 742–52.
- Gładziejewski, P. and Miłkowski, M. [2017]: 'Structural Representations: Causally Relevant and Different from Detectors', *Biology and Philosophy*, **32**, pp. 337–55.
- Goodman, N. [1951]: *The Structure of Appearance*, Cambridge, MA: Harvard University Press.
- Goodman, N. [1968]: *Languages of Art: An Approach to a Theory of Symbols*, Indianapolis, IN: Bobbs-Merrill.
- Goodman, N. [1972]: *Problems and Projects*, Indianapolis, IN: Bobbs-Merrill.
- Hahn, U. [2014]: 'Similarity', *Wiley Interdisciplinary Reviews: Cognitive Science*, **5**, pp. 271–80.
- Hohol, M. [2019]: *Foundations of Geometric Cognition*, New York: Routledge.
- Isaac, A. M. C. [2013]: 'Objective Similarity and Mental Representation', *Australasian Journal of Philosophy*, **91**, pp. 683–704.
- Isaac, A. M. C. [2017]: 'The Semantics Latent in Shannon Information', *British Journal for the Philosophy of Science*, available at doi.org/10.1093/bjps/axx029.
- Israel, D. and Perry, J. [1990]: 'What Is Information?', in P. Hanson (ed.), *Information, Language, and Cognition*, Vancouver: University of British Columbia Press, pp. 1–19.

- Johnson-Laird, P. N. and Byrne, R. M. J. [1991]: *Deduction*, Hove: L. Erlbaum Associates.
- Kotarbińska, J. [1957]: 'Pojęcie znaku', *Studia Logica*, **6**, pp. 57–143.
- Kriegeskorte, N. and Kievit, R. A. [2013]: 'Representational Geometry: Integrating Cognition, Computation, and the Brain', *Trends in Cognitive Sciences*, **17**, pp. 401–12.
- Ladyman, J. and Presnell, S. [2018]: 'Does Homotopy Type Theory Provide a Foundation for Mathematics?', *British Journal for the Philosophy of Science*, **69**, pp. 377–420.
- MacKay, D. M. [1969]: *Information, Mechanism, and Meaning*, Cambridge, MA: MIT Press.
- Mikheev, A. [2005]: 'Text Segmentation', in R. Mitkov (ed.), *The Oxford Handbook of Computational Linguistics*, Oxford: Oxford University Press, pp. 201–18.
- Millikan, R. G. [1984]: *Language, Thought, and Other Biological Categories: New Foundations for Realism*, Cambridge, MA: MIT Press.
- Millikan, R. G. [2004]: *Varieties of Meaning: The 2002 Jean Nicod Lectures*, Cambridge, MA: MIT Press.
- Morgan, A. [2013]: 'Representations Gone Mental', *Synthese*, **191**, pp. 213–44.
- Neander, K. [2017]: 'A Mark of the Mental: In Defense of Informational Teleosemantics', Cambridge, MA: MIT Press.
- Newman, M. H. A. [1928]: 'Mr. Russell's "Causal Theory of Perception"', *Mind*, **37**, pp. 137–48.
- Nirshberg, G. and Shapiro, L. [2020]: 'Structural and Indicator Representations: A Difference in Degree, Not Kind', *Synthese*, available at doi.org/10.1007/s11229-020-02537-y.
- Ólafsdóttir, H. F., Barry, C., Saleem, A. B., Hassabis, D. and Spiers, H. J. [2015]: 'Hippocampal Place Cells Construct Reward Related Sequences through Unexplored Space', *ELife*, **4**, p. e06063.
- O'Brien, G. and Opie, J. [2004]: 'Notes towards a Structuralist Theory of Mental Representation', in H. Clapin, P. Staines and P. Slezak (eds), *Representation in Mind: New Approaches to Mental Representation*, Amsterdam: Elsevier, pp. 1–20.

- Pfeiffer, B. E. and Foster, D. J. [2013]: 'Hippocampal Place-Cell Sequences Depict Future Paths to Remembered Goals', *Nature*, **497**, pp. 74–9.
- Ramsey, W. M. [2007]: *Representation Reconsidered*, Cambridge: Cambridge University Press.
- Rasmussen, J. [2014]: *Defending the Correspondence Theory of Truth*, Cambridge: Cambridge University Press.
- Rupert, R. D. [2018]: 'Representation and Mental Representation', *Philosophical Explorations*, **21**, pp. 204–25.
- Scarantino, A. and Piccinini, G. [2010]: 'Information without Truth', *Metaphilosophy*, **41**, pp. 313–30.
- Shannon, C. [1948]: 'A Mathematical Theory of Communication', *The Bell System Technical Journal*, **27**, pp. 379–423, 623–56.
- Shapiro, L. [1997]: 'Junk Representations', *British Journal for the Philosophy of Science*, **48**, pp. 345–62.
- Shea, N. [2018]: *Representation in Cognitive Science*, New York: Oxford University Press.
- Shepard, R. N. [1958]: 'Stimulus and Response Generalization: Tests of a Model Relating Generalization to Distance in Psychological Space', *Journal of Experimental Psychology*, **55**, pp. 509–23.
- Short, T. L. [2007]: *Peirce's Theory of Signs*, Cambridge: Cambridge University Press.
- Suárez, M. [2003]: 'Scientific Representation: Against Similarity and Isomorphism', *International Studies in the Philosophy of Science*, **17**, pp. 225–44.
- Suárez, M. [2004]: 'An Inferential Conception of Scientific Representation', *Philosophy of Science*, **71**, pp. 767–79.
- Swoyer, C. [1991]: 'Structural Representation and Surrogate Reasoning', *Synthese*, **87**, pp. 449–508.
- Tarski, A. [1933]: *Pojęcie prawdy w językach nauk dedukcyjnych*, Warszawa: Towarzystwo Naukowe Warszawskie.
- Thomson, E. and Piccinini, G. [2018]: 'Neural Representations Observed', *Minds and Machines*, pp. 1–45.
- Tolman, E. C. [1948]: 'Cognitive Maps in Rats and Men', *Psychological Review*, **55**, pp. 189–208.

Tversky, A. [1977]: 'Features of Similarity', *Psychological Review*, **84**, pp. 327–52.

Watanabe, S. [1972]: *Knowing and Guessing: A Quantitative Study of Inference and Information*, New York: Wiley.

Weisberg, M. [2013]: *Simulation and Similarity: Using Models to Understand the World*, New York: Oxford University Press.

Zadeh, L. A. [1965]: 'Fuzzy Sets', *Information and Control*, **8**, pp. 338–53.

Zadeh, L. A. [1988]: 'Fuzzy Logic', *Computer*, **21**, pp. 83–93.