DOI: 10.1111/nous.12404

ARTICLE

Noûs

The Structure of Analog Representation

Andrew Y. Lee[1] | Joshua Myers[2] | Gabriel Oak Rabin[3]

[1] University of Oslo
[2] New York University
[3] New York University, Abu Dhabi

Correspondence
Gabriel Oak Rabin
Email: Gabriel.Oak.Rabin@gmail.com
Andrew Y. Lee
Email: andrew.yuan.lee@gmail.com
Joshua Myers
Email: joshualmyers@gmail.com

Authors listed by alphabetical order of surname.

Abstract

This paper develops a theory of analog representation. We first argue that the mark of the analog is to be found in the nature of a representational system's interpretation function, rather than in its vehicles or contents alone. We then develop the *rulebound structure theory* of analog representation, according to which analog systems use interpretive rules that map syntactic structural features onto semantic structural features. The theory involves three degree-theoretic measures that capture three independent ways in which a system can be more or less analog. We explain how our theory improves upon prior accounts of analog representation, provides plausible diagnoses for novel challenge cases, extends to hybrid systems that are partially analog and partially symbolic, and accounts for some of the advantages and disadvantages of representing analogically versus symbolically.

1 | INTRODUCTION

There is a commonly drawn distinction between analog and symbolic representation. Paradigms of analog representation include mercury thermometers, hand-clocks, photographs, drawings, line graphs, heat maps, and audio recordings. Paradigms of symbolic representation include digital thermometers, digital clocks, binary data, maritime signal flags, traffic lights, and words. The distinction between the analog and the symbolic is appealed to across a diverse range of disciplines, both within philosophy and beyond.[1] But while the difference between the analog and the symbolic appears intuitive, there is a good deal of controversy over how to analyze the distinction. What, exactly, is the mark of the analog?[2]

This paper develops what we call the *rulebound structure theory* of analog representation. The central claim of the theory is that analog systems are those marked by the use of interpretive rules that map syntactic structural features onto semantic structural features. Our goal is to develop this idea into a systematic theory that yields plausible verdicts across a spectrum of cases, including some novel cases that to our knowledge have never been discussed. We develop three degree-theoretic measures—*analog mirroring*, *analog purity*, and *analog mass*—each of which captures a different way a system can be more or less analog. Alongside these first-order claims, we also stake out a methodological position about where the mark of the analog lies: namely, in the nature of a representational system's interpretation function.

The focus on structural features places the rulebound structure theory within the *structural approach* to analog representation. Theories within this approach can be thought of as unified by the core idea that analog representation involves structure corresponding to structure. Other theories within the structural family include density theories (Goodman, 1968; Haugeland, 1981; Schonbein, 2014), magnitude mirroring theories (Beck, 2019; Lewis, 1971; Maley, 2011; Peacocke, 2019), and abstraction theories (Kulvicki, 2015). We will explain how the rulebound structure theory handles a number of challenges facing these other views and develops a more systematic analysis of what it means for structure to correspond to structure.

§2 argues first that the mark of the analog lies in a system's interpretation function (rather than in its vehicles or contents), and then claims that interpretation functions are best thought of as rules mapping inputs to outputs (rather than merely as sets of input-output pairs). **§3** addresses how the idea of isomorphism relates to analog representation and provides a basis for the development of our positive theory. **§4** and **§5** contain the bulk of the rulebound structure theory, where we develop three degree-theoretic measures of analogicity. **§6** extends our theory to hybrid systems that are partly analog and partly symbolic. **§7** explains how our theory relates to and improves upon extant theories of analog representation. **§8** addresses objections. **§9** discusses the advantages and disadvantages of representing analogically versus symbolically.

2 | FOUNDATIONS

We begin with some preliminary remarks on how to think about representational systems and where to look for the mark of the analog.

[1] References to analog representation occur in the philosophy of science (Papayannopoulos, 2020, Sterrett 2017, Trenholme 1994), philosophy of art (Goodman 1968, Lopes 2018, Thomson-Jones 2015), epistemology (Dretske 1981), philosophy of perception (Beck 2019, Block forthcoming, Peacocke 1986), cognitive science (Dehaene 1997, Feigenson et al. 2002, Shagrir 2010), computer science (Siegelmann & Fishman 1998, Vergis et al. 1986, Xue 2016), and electrical engineering (Couch 2013).

[2] The terms 'iconic', 'pictorial', and 'imagistic' are sometimes used instead of 'analog' and the terms 'digital' and 'discursive' are sometimes used instead of 'symbolic'. Some (though certainly not all) uses of these terms aim at the same representational kind we seek to capture. We will remain neutral throughout the paper on any relationships between the analog, the iconic, the pictorial, and the imagistic.

2.1 | The Interpretation Function

Every representational system has vehicles, contents, and an interpretation function. The *vehicles* are the representers, the *contents* are the represented, and the *interpretation function* maps vehicles to contents.[3] Consider a digital clock (a paradigmatic symbolic system): the vehicles are numeral displays, the contents are times of day, and the interpretation function maps '12:01' to one minute after noon. Or consider a mercury thermometer (a paradigmatic analog system): the vehicles are thermometers individuated by mercury heights, the contents are temperatures, and the interpretation function maps 6 cm of mercury to 30°C.

This paper is an investigation into what makes a representation analog. But first we must ask: where should one look for the mark of the analog?[4] The obvious options include looking at a system's vehicles, contents, or interpretation function. While methodological views about where to look are rarely made explicit, vehicular rhetoric dominates the literature. For example, Quilty-Dunn [2017: 61] states that "formats are general types of vehicular structures," and Beck [2019: 323] writes that he will "understand the thesis that perception is analog solely as a claim about perception's vehicles." Others have favored alternative approaches: for example, Haugeland [1991] endorses a content approach and Goodman [1968] defends a hybrid vehicle-content account.

We think that neither vehicle nor content approaches identify the true mark of analog representation. As a preliminary argument, observe that it is possible for (1) the same contents to be represented either analogically or symbolically and (2) the same vehicles to represent either analogically or symbolically. To see the first dissociation, consider a digital clock versus a hand-clock. Both systems represent times of day, so both have the same contents. But the digital clock is symbolic while the hand-clock is analog. To see the second dissociation, consider a spy tapping on a desk to send a message in Morse code versus a geologist tapping on a desk to describe a pattern of seismic wave activity. Even if the two series of taps are indistinguishable, the spy's taps represent symbolically whereas the geologist's taps represent analogically.[5]

Our view is that the mark of the analog is found in a system's interpretation function, rather than in its vehicles or contents alone. Call this the *interpretation function approach*.[6] On this

[3] We use the term 'vehicles' to refer to syntactically individuated vehicle types rather than vehicle tokens. A mercury thermometer whose 6cm of mercury represents 30°C may be token-distinct but still type-identical to a second mercury thermometer whose 6cm also represents 30°C. Note that syntactic types are usually more coarsely individuated than physical types: even if the first thermometer contains one more atom of mercury than the second, that will probably not suffice for a syntactic difference between the two.

[4] Our talk of a search for the mark of the analog designates an investigation into the *grounds*, rather than the *bearers*, of analogicity. We are primarily concerned with what makes something analog, rather than which entities bear the property of being analog. Insofar as we have a view about bearers, we are inclined to think that the primary bearers of analogicity are representational systems. But we also believe that one may permissibly talk of vehicles, contents, or interpretation functions as being analog when they belong to a system that has what we take to be the mark of the analog: namely, the use of rules that map syntactic structural features onto semantic structural features. Thanks to an anonymous referee for helpful discussion of the grounds/bearers distinction.

[5] One might seek to individuate vehicles in a way that permitted physically indistinguishable vehicles, such as the spy's and the geologist's taps, to count as distinct kinds of vehicles. We will simply note that the examples preceding this footnote are illustrative and preliminary. Our argument for the interpretation function, rather than vehicle or content, approach to analog representation is largely supported by the usefulness of that methodology, which will be manifest over the course of the paper.

[6] Examples of interpretation function theorists include Barwise and Hammer [1996], Burge [2018], and Frigario, Giordani, and Mari [2013], though they vary in the representational kinds targeted.

view, analog representation is a matter of *how* vehicles represent contents, rather than what does the representing or what is represented. As a corollary, analogicity is first and foremost a property of representational systems (rather than individual representations). The examples above of spies, geologists, and clocks are evidence in favor of the interpretation function approach. Nevertheless, while our focus on the interpretation function is in some ways a departure from prior vehicle (and content) approaches, this shift constitutes a less radical break from orthodoxy than it may initially appear. Many theorists who employ vehicular rhetoric clearly care about interpretation functions. For example, Beck [2019: 334] argues that analogicity is a matter of vehicles covarying with contents, indicating he is committed to the interpretation function approach despite his aforementioned emphasis on vehicles. Given this, we think that the interpretation function approach better captures what many theorists have had in mind all along.

2.2 | Rules

How are interpretation functions individuated? At minimum, any two interpretation functions with distinct input-output pairs are not identical. The more interesting question is whether there is more to the nature of an interpretation function than its inputs and outputs. In our view, interpretation functions (and thus representational systems) are best individuated by *how* inputs are mapped to outputs. In other words, we must look to the *rules* the system uses to map vehicles to contents.

This idea can be illustrated via a debate from the history of mathematics about the nature of mathematical functions: is a function any set of input-output pairs, or is it a rule for getting from inputs to outputs?[7] To illustrate, consider the following two functions:

$$f(x) = x^2$$

$$g(x) = \sqrt{(x^2+3)(x^2-2)+(6-x^2)}$$

These functions yield the same output for any input, but they use different rules to move from those inputs to those outputs. The prevailing convention in mathematics is *extensionalist*: mathematical functions are individuated purely in terms of input-output pairings. But we think that for the purpose of investigating representational systems, it is better to take a more fine-grained *intensionalist* approach, where interpretation functions are individuated by the rules that are used to map inputs to outputs.[8] On this picture, every representational system has rules by which it maps vehicles to contents, and even representational systems that have the exact same input-output vehicle-content pairs can nevertheless be distinct in virtue of deploying different rules.[9]

[7] For example, Euler [1748] defined a function as a rule—an "analytic expression"—for generating outputs from inputs. Schönfinkel [1924] was more liberal, permitting any arbitrary correspondence to count as a function as long as each input generated only one output. For discussions of this debate, see Kleiner [2009] and Youschkevitch [1976].

[8] The best example we know of a theory of analog representation that revolves around the rules of the interpretation function is Frigario, Giordani, and Mari [2013].

[9] Rules need not be systematic. For example, a brute assignment of each individual vehicle to its content still counts as a rule. It is plausible that many symbolic systems rely on rules of this nature.

To motivate the idea that the rules of the interpretation function matter when evaluating whether a representational system is analog or symbolic, consider the following two systems:

Stipulated Population Map—**Vehicles**: 100 shades of red. **Contents**: 100 population densities. **Interpretation**: Stipulated one-by-one assignment of shades of red to population densities.

Structural Population Map—**Vehicles**: 100 shades of red. **Contents**: 100 population densities. **Interpretation**: If shade_1 is darker than shade_2 then the content of shade_1 (a population density) is denser than the content of shade_2.

These systems map all the same vehicles (inputs) to all the same contents (outputs), but they do so using different rules. *Stipulated Population Map* assigns each red shade one-by-one to a specific population density. *Structural Population Map* assigns contents using one simple rule: "darker shades represent greater population density." Even though the systems have identical input-output structures, *Stipulated Population Map* seems symbolic while *Structural Population Map* seems analog. We believe this felt difference corresponds to a real difference in representational kind.[10] And we think the difference is best explained by the fact that the systems have different interpretive rules. From this point forward, we will take as given that the nature of a representational system depends not only on the input-output structure of a system, but also on the rules that constitute its interpretation function.

We will remain largely neutral on the nature of rules. This is because we want our claims to be compatible with a variety of views about the nature of interpretation functions and about representation more generally. For example, rules need not be understood as linguistically expressible sets of directions. Our use of the term 'rules' is intended merely to designate *how* an interpretation function maps vehicles to contents. Unless one is prepared to deny that there can be distinct representational systems with the same inputs and outputs, and consequently deny that there is any representationally relevant difference between *Stipulated Population Map* and *Structural Population Map*, one is committed to the idea that interpretation functions have rules.

Our claims so far are as follows. When determining whether a representational system is analog, one should look at the system's interpretation function and not only at vehicles or contents. Furthermore, one should look at the *rules* of the interpretation function. These claims are methodological, in the sense that they are claims about *where to look* for the mark of the analog rather than claims about *what to look for*. §4 and §5 develop our positive theory of analog representation, which explains what to look for.

3 | ISOMORPHISM

The unifying theme behind the structural approach to analog representation is that analog representation is a matter of structure corresponding to structure. For many readers, this idea will

[10] For readers who are uncertain about our appeals to intuitions about which systems are more analog, we wish to emphasize that this paper is not a conceptual analysis of the word 'analog' or the concept ANALOG. Instead, we think there is an important difference in representational kind between cases like *Structural Population Map* and *Stipulated Population Map*, and we think the terms 'analog' and 'symbolic' work well for tracking this difference. When we say that one system "seems more analog" than another, we seek to evoke the reader's feeling for that difference. §7.3 discusses related issues.

immediately evoke the notion of isomorphism. An *isomorphism* is, roughly put, a structure-preserving bijective function. Though we do not think that isomorphism provides a satisfactory criterion for analog representation, it will be useful to begin with a discussion of isomorphism and develop our theory from that basis.

Although the term 'isomorphism' is common in the literature on analog representation, there is a surprising lack of uniformity in what is meant by 'isomorphism theory'.[11] In some cases, the intended idea is that a system is analog just in case there is an isomorphism between its vehicles and its contents. We will argue shortly that such a claim is at best incomplete and at worst false. In other cases, invocations of the term 'isomorphism' are simply intended to express the general idea behind all structural approaches: that analog representation involves similar structure between vehicles and contents. We think that such an idea is on the right track. But a satisfying theory demands a more systematic analysis. This section aims to explain more precisely what an isomorphism is, articulate what we take to be the best version of an isomorphism theory of analog representation, and argue that such a theory is nevertheless inadequate.

At first pass, it is natural to characterize the isomorphism theory as endorsing the following criterion: a system is analog just in case there exists an isomorphism between its vehicles and its contents. Consider how in *Mercury Thermometer*, there seems to be an isomorphism between mercury heights and temperatures, whereas in the system for a digital thermometer, there does not seem to be an isomorphism between numerals and temperatures. However, this criterion is overly permissive. Consider the following systems:

Mercury Thermometer—**Vehicles**: 100 mercury heights. **Contents**: 100 temperatures. **Interpretation**: If height_1 is taller than height_2 then the content of height_1 (a temperature) is warmer than the content of height_2.

Scrambled Thermometer—**Vehicles**: 100 mercury heights. **Contents**: 100 temperatures. **Interpretation**: An unsystematic one-by-one bijective assignment of mercury heights to temperatures.

Mercury Thermometer and *Scrambled Thermometer* contain exactly the same vehicles and exactly the same contents. Therefore, if there exists an isomorphism between the vehicles and contents for *Mercury Thermometer*, then there also exists an isomorphism between the vehicles and contents for *Scrambled Thermometer*. But *Mercury Thermometer* is analog while *Scrambled Thermometer* is not. Therefore, whether there exists an isomorphism between vehicles and contents cannot be enough to decide whether a system is analog.

A natural move at this point is to appeal to the fact that the systems map vehicles to contents differently. More specifically, perhaps instead of asking whether there exists some function or other that is an isomorphism between heights and temperatures, we should instead ask about a very specific function: the interpretation function.[12] The interpretation function of *Mercury Thermometer* always maps greater heights of mercury to warmer temperatures; the interpretation function of

[11] See Beck [2015, 2019], Blachowicz [1997], Block [forthcoming], Frigario, Giordani, Mari [2013], Kulvicki [2015], Maley [2011], and Shepard [1978]. Haugeland [1991] calls the isomorphism account the "canonical account."

[12] The only extant account of analog representation we are aware of that requires that the interpretation function itself must be an isomorphism is developed by Frigario, Giordani, and Mari [2013].

Scrambled Thermometer does not. This strategy succeeds in delivering distinct verdicts on the analogicity of *Mercury Thermometer* versus *Scrambled Thermometer*. However, the strategy works only in cases where the two interpretation functions differ extensionally (i.e. with respect to which inputs get mapped to which outputs). The following example illustrates how even systems with extensionally equivalent interpretation functions can differ in analogicity (an idea already alluded to in §2):

Lucky Scrambled Thermometer—**Vehicles**: 100 mercury heights. **Contents**: 100 temperatures. **Interpretation**: A stipulative one-by-one bijective assignment of mercury heights to temperatures, where by luck of the random assignment, each height happens to be assigned to the very same temperature as in Mercury Thermometer.

If we consider only inputs and outputs, then the interpretation function of *Mercury Thermometer* is indistinguishable from that of *Lucky Scrambled Thermometer*. But *Lucky Scrambled Thermometer* is intuitively not analog. This kind of problem arises for any view that considers only the inputs and outputs of interpretation functions. Even when two systems have interpretation functions with identical input-output structures, the way that they map those inputs to outputs can matter for whether or not the systems are analog. In what follows, we explain how this idea—as well as the notion of isomorphism itself—can be made precise.

Let a *space* be an ordered pair of a set of elements and a set of relations on those elements.[13] Mercury heights paired with a *taller-than* relation is a space. Mercury heights paired with a *twice-the-height-of* relation is another space. Let $X = < D_X, R_X >$ and $Y = < D_Y, R_Y >$ be spaces. A function f from D_X to D_Y is an *isomorphism between X and Y* just in case the following three criteria are satisfied:

1. f is a bijection, meaning that f maps every element of D_X to a unique element of D_Y and every element of D_Y is the output $f(x)$ for some x in D_X.
2. R_X and R_Y each contain at least one relation.
3. For every relation r_X in R_X, there is a relation r_Y in R_Y, such that for any elements x_1 and x_2 in D_X, $r_X(x_1, x_2)$ if and only if $r_Y(f(x_1), f(x_2))$, and for every r_Y in R_Y, there is an analogous r_X in R_X.

The third criterion says that every relation in space X has a counterpart relation in space Y, and every relation in Y has a counterpart in X, such that two elements of X bear that relation if and only if their counterparts in Y bear the counterpart relation.[14] If the third criterion is satisfied, then the function *preserves* the relevant relations. Note that while the function f is a function from the set D_X to the set D_Y, f is an isomorphism between the space X and the space Y. In total, there are three ways that there can be a failure of isomorphism: (1) if f is not a bijec-

[13] The phrase "relations on those elements" requires that for each relation, some elements in the domain in fact bear that relation.

[14] These definitions are formulated in terms of binary relations, but they apply straightforwardly to relations of any adicity. We assume the generalized formulation moving forward.

tion, (2) if the spaces contain no relations, or (3) if f does not preserve the relations on the spaces.[15]

Strictly speaking, it does not make sense to talk about isomorphism *simpliciter* between a set of vehicles and a set of contents: isomorphism is a relation between structured *spaces*, rather than a relation between unstructured *sets*. This means that when speaking of isomorphism, it is important to specify the relations (or spaces) with respect to which the isomorphism holds. Even if a function f is an isomorphism between set D_1 paired with relation r_1 to D_2 paired with r_2, f may nevertheless fail to be an isomorphism between D_1 and D_2 paired with other relations. For example, even though $f(x) = e^x$ is an isomorphism from the real numbers under addition to the positive reals under multiplication, it is not an isomorphism from the real numbers under addition to the positive reals under addition. Though this may initially strike some as a technical detail, the point matters for developing a theory of analog representation. Theorists often make claims about isomorphism that leave implicit which relations structure the spaces. This underspecification is problematic: since the physical world is replete with relations, it is easy to find isomorphisms between any two domains. Because of this, the claim that there exists an isomorphism between a system's vehicles and its contents is almost trivial. Moreover, even the more specific claim that the interpretation function itself is an isomorphism between the set of vehicles and the set of contents is incomplete: the relevant relations on the vehicles and the contents must also be specified.

We now turn to a crucial issue. Supposing that isomorphism is important for analog representation, what are the relations with respect to which the isomorphism must hold? Different structural approaches to analog representation will answer this question differently. Some will say magnitude relations, others dense relations, and yet others abstraction relations. By contrast, the rulebound structure theory says that the relevant relations are whichever figure into the rules of the interpretation function. Every system has an interpretation function that maps vehicles to contents via certain rules. Those rules take into consideration certain relations and not others. For example, *Mercury Thermometer* has a rule that says that taller heights represent warmer temperatures. More precisely, let a system's *rulebound relations* be the relations that are codified by the rules of its interpretation function. *Taller-than* is rulebound in *Mercury Thermometer*; *redder-than* is not. On our view, the relations that matter for isomorphism are the system's rulebound relations. This approach suggests the following criterion for analog representation:

> **The Isomorphism Criterion:** The system's interpretation function is an isomorphism between the syntactic space and the semantic space.

The *syntactic space* is the space consisting of all the vehicles of the system paired with all the rulebound relations on those vehicles. The *semantic space* is the space consisting of all the contents of the systems paired with the rulebound relations on those contents. In order to satisfy the isomorphism criterion, a system must have an interpretation function that is a bijection, that codifies at least one relation, and that preserves the relation(s) that it codifies. This makes the above isomorphism criterion much more substantial than the earlier criterion (that there exists some isomorphism between vehicles and contents). *Mercury Thermometer* satisfies the isomorphism

[15] A mathematical example of isomorphism might help precisify these ideas. $f(x) = e^x$ is an isomorphism between the real numbers $\mathbb{R}$ under addition and the positive real numbers $\mathbb{R}^+$ under multiplication. In this case, the domains are the real numbers and the positive real numbers and the relations are addition and multiplication (understood as triadic relations). The first criterion is satisfied because e^x is a bijection from $\mathbb{R}$ to $\mathbb{R}^+$. The second criterion is satisfied because the equality $e^{x+y} = e^x \times e^y$ ensures relation preservation: $x + y = z$ iff $f(x) \times f(y) = f(z)$. Therefore, the spaces are isomorphic.

criterion because its interpretation function is an isomorphism between its syntactic and semantic spaces: *taller-than* and *warmer-than* are rulebound and the interpretation function preserves them. By contrast, neither *Scrambled Thermometer* nor *Lucky Scrambled Thermometer* satisfy the isomorphism criterion because neither system has rulebound relations that can be preserved in the first place. These observations indicate that the isomorphism criterion provides a more accurate, substantive, and precise distinction between the analog and the symbolic than prior accounts that appeal more casually to the notion of isomorphism.

Nevertheless, we still think that the isomorphism criterion is inadequate as a general theory of analog representation. The basic problem is that systems that *nearly* satisfy the isomorphism criterion can still be highly analog. Consider the following examples:

Scrambled Population Map—**Vehicles**: 100 shades of red. **Contents**: 100 population densities. **Interpretation**: A stipulative bijective mapping between shades of red and population density, with no correspondence between darkness of shade and density.

Almost Isomorphic Population Map—**Vehicles**: 100 shades of red. **Contents**: 100 population densities. **Interpretation**: A systematic bijective mapping between shades of red and population density, with darker shades corresponding to greater densities, except that the darkest shade is mapped to the second densest population and the second darkest shade is mapped to the densest.

Scrambled Population Map maps shades of red to population densities arbitrarily. *Almost Isomorphic Population Map* maps shades of red to population densities systematically, with exceptions for the two darkest shades and densest populations. Neither of these systems satisfies the isomorphism criterion, since neither function preserves the rulebound relations in the mapping from the syntactic to the semantic space. Yet it is intuitive that *Almost Isomorphic Population Map* is more analog than *Scrambled Population Map*. Our diagnosis is that the difference is due to the fact that *Almost Isomorphic Population Map* nearly satisfies the isomorphism criterion whereas *Scrambled Population Map* does not come close.

The core problem is that isomorphism is all-or-nothing, whereas analogicity comes in degrees. Consequently, the isomorphism criterion—at least by itself—is inadequate for capturing the degreed aspects of analog representation. Still, we think that appeals to isomorphism are on the right track. If analog representation is a matter of structure corresponding to structure, then isomorphism is a promising start towards developing that idea. The preceding discussion sets the stage for our positive theory, where we will define a degree-theoretic variant of the isomorphism criterion that captures a core dimension of analogicity.

Before moving forward, it is worth making a point about the kinds of systems that will serve as our examples.[16] We will focus on systems whose rulebound relations range over

[16] Two simplifying assumptions are worth mentioning. First, we assume *free recombination*, meaning that a vehicle (or content) that takes a value along one parameter (e.g. size) can take any value along a different parameter (e.g. color). Free recombination prohibits a system in which each vehicle has both a height and either a red, blue, or green color, yet a green vehicle of 10cm is not well-formed, even though vehicles with every other height and color combination are. Second, we assume *comparability*, meaning that every pair of vehicles (or contents) is related by each relation. Comparability prohibits a situation in which two vehicles are incommensurable along a parameter (e.g. size). Going forward, we assume that the

whole vehicles and whole contents (e.g. *Mercury Thermometer*'s *taller-than* relation), rather than on systems whose rulebound relations range over parts of vehicles or parts of contents (e.g. a photograph's *darker-than* relation between individual pixels). We focus on the former kinds of systems both because they are simpler and because we think the canonical examples of analog representation mostly fall within that class. However, our measures of analogicity also apply to systems within the latter category. Our measures appeal to rulebound relations, and each measure applies whether those relations relate wholes to wholes or parts to parts.

4 | MIRRORING

Our theory consists of three degree-theoretic measures of analogicity. This section focuses on the first and most central measure: *analog mirroring*. To set the stage, consider again *Scrambled Population Map* (which maps shades of red to population densities arbitrarily) and *Almost Isomorphic Population Map* (which maps shades of red to population densities systematically except for a mismatch between the two darkest shades and two densest populations). Neither of these systems satisfies the isomorphism criterion. Yet *Almost Isomorphic Population Map* seems more analog than *Scrambled Population Map*.

A natural explanation for the difference is that it is easy to whittle down *Almost Isomorphic Population Map* into a system that satisfies isomorphism: all one needs to do is remove a single element from the syntactic space (e.g. the darkest shade of red) and a single element from the semantic space (e.g. the second densest population). By contrast, turning *Scrambled Population Map* into a system that satisfies isomorphism is either difficult or impossible. These observations hint at the basic idea behind analog mirroring: by counting how many elements (vehicles and contents) must be eliminated before isomorphism is achieved, analog mirroring measures how close the system is to satisfying the isomorphism criterion. Because *Almost Isomorphic Population Map* can be made to satisfy the criterion by eliminating a single vehicle and a single content, that system has a very high, but not quite maximal, degree of analog mirroring.

Let us say that a representational system S^- is a *subsystem* of a system S just in case the syntactic and semantic spaces of S^- are subspaces of the syntactic and semantic spaces of S. For example, a subsystem of *Mercury Thermometer* would be a system whose syntactic space includes mercury heights 1–90 (cutting out 91–100), whose semantic space includes temperatures 1–90 (cutting out 91–100), and which retains the *taller-than* and *warmer-than* relations on the remaining heights and temperatures. Intuitively, we can think of a *subsystem* as the original system with some vehicles and contents removed.

Almost every system that has rulebound relations but that does not itself satisfy the isomorphism criterion will nevertheless have multiple subsystems that do.[17] Some of these subsystems will be larger than others. Some (perhaps only one) such system will be maximal: it will be as big as or bigger than any other subsystem that satisfies the isomorphism criterion. *Almost Isomorphic*

systems we target satisfy both free recombination and comparability (on their syntactic and semantic domains). Although we think that some oddball systems do not obey these principles, we do not think that the price of the extra complications in the main presentation of our theory are worth the benefit of accommodating these cases.

[17] Why "almost every system"? Since the definition of a space requires that each relation be instantiated, systems in which the elimination procedure results in a single vehicle or single content (or less) will have an analog mirroring score of 0.

Population Map has two such subsystems: (1) the subsystem that cuts the darkest shade of red and the second densest population, and (2) the subsystem that cuts the second darkest shade and the densest population. Call any maximal subsystem that satisfies the isomorphism criterion a *mirrored subsystem* of the original system.

A system's analog mirroring score is the ratio of the size of its mirrored subsystem(s) to the size of the system itself.[18]

$$\textbf{Analog Mirroring} = \frac{\text{size of mirrored subsystem}}{\text{size of system}}$$

Analog mirroring measures how much of the representational system satisfies isomorphism. This definition has some attractive formal features. Every system will have a mirroring score between 0 and 1. The maximum value—1—occurs when no elements need to be eliminated because the original system itself already satisfies the isomorphism criterion. The minimum value—0—occurs when no subsystem satisfies isomorphism. To get a mirroring score above zero, a system's interpretation function must codify some relations. When a system has no rulebound relations, it automatically gets a mirroring score of 0.

This formalization yields intuitive verdicts across a variety of cases. *Structural Population Map*, which maps darker shades of red to denser populations with no exceptions, has a mirroring score of exactly 1. Since only two elements need to be eliminated from *Almost Isomorphic Population Map*, the mirrored subsystem of *Almost Isomorphic Population Map* is nearly as large as the original system. Consequently, *Almost Isomorphic Population Map* has a mirroring score of nearly 1. By contrast, *Scrambled Population Map* has a mirroring score of 0, since none of its subsystems satisfies isomorphism.

In effect, our analog mirroring measure is a degreed version of isomorphism. We suspect that such a measure can be useful not only for our theory, but also for research outside of analog representation.[19]

5 | PURITY AND MASS

Analog mirroring is a more nuanced measure of analogicity than the isomorphism criterion. Yet it is still not enough to fully capture the respects in which systems can be more or less analog. This section develops our two other measures: *analog purity* and *analog mass*. For clarity of exposition, we assume that the systems discussed in this section have maximum analog mirroring. This assumption is harmless: on our account, a system with non-maximum mirroring inherits its purity and mass values from its mirrored subsystem, which has maximum mirroring by definition. Mirroring may thus be thought of as the gatekeeper to the realm of the analog: without at least some mirroring, a system has no purity or mass, because it has no mirrored subsystem.

[18] In systems whose spaces contain a finite number of elements, it is natural to measure size in terms of cardinality. In systems whose spaces contain an infinite number of elements, we suggest appealing to the Lebesgue measure. See Tao [2011] for more discussion of measures.

[19] As an example, Shea [2014: 140] notes that it would be useful to have a notion of "approximate isomorphism" in developing a theory of mental content. This is exactly what our mirroring measure delivers.

5.1 | Analog Purity

Consider the following system:

City Populations: **Vehicles**: 10 letters of the alphabet, each coming in three different sizes—big, medium, and small. **Contents**: Populations for 10 different cities. **Interpretation**: Each letter represents a different city ('A' represents Athens, 'B' represents Beijing, and so on). The size of the letter represents the approximate population of the city: a big letter represents "more than 10 million people," a medium letter represents "between and 1 and 10 million people," and a small letter represents "less than 1 million people."

City Populations has maximum analog mirroring, since its interpretation function is a bijection that preserves the *larger-than* relation on vehicles via the *has-greater-population-than* relation on contents. However, the system seems less than maximally analog. Here is our diagnosis of this intuition: although *City Populations*' interpretation function preserves all of the system's structure, the system does not utilize much structure to begin with.

Analog purity aims to capture this further dimension of analogicity by measuring the degree to which the system's elements are structurally distinct from each other. More precisely, two elements are *structurally distinct* whenever they have different structural profiles, where two elements have the same *structural profile* just in case they bear exactly the same rulebound relations to every other element of the system. In *City Populations*, big-'A' and big-'B' have the same structural profile (and hence are not structurally distinct). This is because big-'A' and big-'B' bear the *larger-than* relation to exactly the same elements. By contrast, big-'A' and small-'A' have distinct structural profiles (and hence are structurally distinct): small-'A' bears the *larger-than* relation to no other letters whereas big-'A' bears the *larger-than* relation to every small and medium letter.

Let us say that two elements have distinct *nominal profiles* just in case they are representationally distinct despite having the same structural profile. In *City Populations*, nominal profiles correspond to letter types. Although big-'A' and big-'B' have the same structural profile, they are still representationally distinct: big-'A' represents Athens while big-'B' represents Beijing. To illustrate further, big-'A' and small-'B' have distinct nominal and distinct structural profiles (different letter, different size), whereas small-'A' and 'big-'A' have identical nominal but distinct structural profiles (same letter, different size). The nominal profiles can be counted by taking the total number of elements in the space and dividing by the number of structural profiles.

The purity of a system is determined by comparing the number of structural and nominal profiles:

$$\textbf{Analog Purity} = \frac{\text{number of structural profiles}}{\text{number of structural profiles} + \text{number of nominal profiles}}$$

Analog purity is designed to capture the degree to which a system's elements are structurally distinct from each other. A system has maximum purity exactly when, in each of the syntactic and semantic domains, every element is structurally distinct from every other element. In such a system, each element has a unique structural profile. The more a system diverges from this ideal, the lower its purity. In systems that fall short of maximum purity, some elements have the same structural profile but different nominal profiles. These are exactly the elements that diminish a system's purity score.[20]

Purity is a degree-theoretic measure of analogicity. Like mirroring, it has a maximum value: 1. This maximum occurs for all and only systems in which each element has a unique structural profile. Purity also has a minimum value: 0. This minimum occurs in systems that have no rule-bound relations and thus no structural profiles. This minimum value also occurs for any systems that have zero mirroring, since purity is defined over a representational system's mirrored subsystem.

A technical point is worth mentioning here. Because the subsystems under consideration have maximum analog mirroring, the syntactic and semantic space are guaranteed to mirror each other perfectly. As a result, there will always be the same number of structural profiles for both the syntactic space and the semantic space. In *City Populations*, the semantic structural profiles correspond to (represented) populations, and the syntactic structural profiles correspond to letter sizes. Whether we count syntactically or semantically, we get the result that *City Populations* has a purity score of 3/13, since on the syntactic side there are 3 letter sizes and 10 letters while on the semantic side there are 3 populations and 10 cities.[21]

The unifying theme of the structural approach is that analog representation is a matter of structure corresponding to structure. In *City Populations*, all of the system's structure corresponds to structure, and as a result the system has maximum mirroring. However, the system does not have much structure in the first place: most of the differences between *City Population*'s elements are nominal, rather than structural. In other words, little of the representational work done by *City Populations* is the result of structure corresponding to structure. By contrast, in *Mercury Thermometer* each element has a unique structural profile. The syntactic profiles are generated by height of mercury and the semantic profiles by temperature. As a result, *Mercury Thermometer* has no nominal profiles, so *Mercury Thermometer* has maximum purity. This illustrates how the purity measure captures the fact that *City Populations* is less analog, in at least one respect, than *Mercury Thermometer*, despite the fact that both have maximum mirroring.

We can also make useful comparisons between systems that have less than maximum purity. For example, consider a variant on *City Populations* that uses 10 letters coming in 10 sizes (instead of 3) to represent 10 cities of 10 different populations (instead of 3). This system has a greater number of structural profiles than *City Populations* but the same number of nominal profiles. As a result, *City Populations* has less purity (3/13) than the expanded system (10/20). This is the intuitively correct result.

[20] In the special case where every element of a system has a unique structural profile, we take the system to have zero nominal profiles. There is perhaps a sense in which a maximally pure system, where every element has a unique structural profile, contains a single "nominal profile." However, counting nominal profiles in this way would fail to serve its main role: namely, gauging "the extent to which the system contains distinct elements that have the same structural profile." In fact, no system would ever have maximum purity under this approach. Since we use the term 'nominal profile' to denote those non-structural features that make for a representational difference between elements, we do not count trivial features that all syntactic or semantic elements share, such as "is a letter", as nominal.

[21] If there could be systems with infinitely many structural profiles but finitely many nominal profiles, then such systems would count as maximally pure given our definition (for such systems, the proportion of structural profiles to total profiles is 1). Nevertheless, some readers may have the intuition that such systems ought not count as maximally pure. In our view, such systems reveal a divergence between two equally good ways of measuring purity: namely, (1) by counting structural and nominal profiles (which is the measure we have defined), and (2) by counting structural and nominal features (such as letter or size). We focus only on the first formulation because doing so greatly simplifies the technical exposition with little loss. However, for those concerned about the kinds of infinitary systems described above, the second formulation provides a natural extension of our theory that accommodates the desired intuition.

5.2 | Analog Mass

The following pair of examples motivates the need for a third measure—analog mass—alongside analog mirroring and analog purity.

Sparse Heat Map—**Vehicles**: 10 colors ranging from pure red to pure blue. **Contents**: 10 temperatures ranging from 0°C to 100°C. **Interpretation**: Redder colors map to warmer temperatures, bluer colors map to colder temperatures.

Rich Heat Map—**Vehicles**: 1000 colors ranging from pure red to pure blue. **Contents**: 1000 temperatures ranging from 0°C to 100°C. **Interpretation**: Redder colors map to warmer temperatures, bluer colors map to colder temperatures.

These systems map redder colors to warmer temperatures and bluer colors to colder temperatures. *Sparse Heat Map* uses 10 colors to represent 10 temperatures; *Rich Heat Map* uses 1000 colors to represent 1000 temperatures. Since the *redder-than* relation on colors corresponds to the *warmer-than* relation on temperatures, both systems have maximum mirroring. Since each element of each system has a unique structural profile, both systems have maximum purity. Yet even though the systems have identical mirroring and purity scores, there remains a way in which *Rich Heat Map* is more analog than *Sparse Heat Map*.

Here is our diagnosis. Mirroring and purity both provide proportional measures of how much of the system's representational work is done by structure representing structure. But degree of analogicity is also a function of the absolute amount of representational work done by structure representing structure. Analog mass captures this idea:

$$\textbf{Analog Mass} = \text{number of structural profiles}$$

Mass is simply the numerator of purity. It can be calculated as outlined in the previous section. To return to our example: *Sparse Heat Map* has 10 structural profiles; *Rich Heat Map* has 1000.[22] Because of this, although the two systems are equivalent with respect to their mirroring and purity scores, *Rich Heat Map* has greater analog mass than *Sparse Heat Map*. The mass measure allows us to capture the respect in which the 1000-color map is more analog than the 10-color map. Note that unlike mirroring and purity, mass has no maximum: there is no limit to the number of structural profiles that a system can have.

5.3 | The Rulebound Structure Theory

The unifying theme of the structural approach is that analog representation is a matter of structure corresponding to structure. We have developed and precisified this intuitive idea into three mea-

[22] There is a way of counting profiles where *Sparse Heat Map* has 20 structural profiles: 10 syntactic structural profiles (colors) and 10 semantic structural profiles (temperatures). However, because analog mass (as well as analog purity) applies only to mirrored subsystems, the numbers of structural and nominal profiles will always be identical between the syntactic and semantic spaces. We count in the way that yields 10 structural profiles because we find it simpler. Since each method of counting generates the same comparative rankings between systems, nothing turns on this choice.

sures: analog mirroring, analog purity, and analog mass. Each measure captures a distinct way in which a system can live up to that core slogan. In particular, analog mirroring measures the extent to which structure corresponds to structure, analog purity measures the proportion of the representational work done by structure corresponding to structure, and analog mass measures the absolute amount of representational work done by structure corresponding to structure.

Some readers might have intuitions that our measures are unable to capture. In our view, there are many measures that can track intuitions about relative analogicity. Rather than trying to capture all intuitions expressible using the phrase 'more analog', we have focused on what we take to be three fundamental, general, and joint-carving measures. Other readers may wonder how our measures combine to yield a total analogicity score. But we wish to resist the idea that there is any objective fact of the matter about how to weigh these measures against one another. Our view is that these measures identify distinct dimensions along which representational systems can be more or less analog, each of which can vary independently of the others. Each of these principles tracks different families of intuitions about degree of analogicity. Relatedly, our principal focus is on generating the right comparisons as to which systems are more or less analog along the relevant dimensions, rather than on the absolute numbers outputted by the measures.

While some discussions of analog representation take analogicity to come in degrees, others characterize it as a binary property. Our focus has been on the degreed aspects of analogicity, but we can extract several binary conceptions of analogicity from our degree-theoretic measures. First, both analog mirroring and analog purity have a maximum value. We can say that a system is *perfectly analog* just in case it has both maximum mirroring and maximum purity. Nearly all paradigms of analog representation, including mercury thermometers, hand clocks, and photographs, are perfectly analog. Second, mirroring, purity, and mass all have a minimum value. We can say that a system is *perfectly symbolic* when it has minimum (i.e. zero) mirroring, and thus zero purity and mass. Conversely, when a system has non-zero mirroring, purity, and mass, it is not perfectly symbolic, and is thus at least somewhat analog.[23]

6 | HYBRID SYSTEMS

This section explains how our framework can be used to analyze hybrid systems, an often overlooked class of representations. A *hybrid system* is a representational system that has some features that represent analogically and some features that represent symbolically. As an example, consider *Colored Thermometer*, in which a colored column represents both temperature and cloud cover:

Colored Thermometer—**Vehicles**: Colored columns coming in 3 colors and 100 heights. **Contents**: Combinations of temperature and cloud coverage. **Interpretation**: Taller columns represent warmer temperatures. Blue represents clear skies, yellow represents partly cloudy, and red represents overcast.

Theories of analog representation tend to either deliver the wrong results for hybrid systems (by classifying them as analog or symbolic *simpliciter*) or do not suggest any obvious way of handling

[23] These ways of extracting binary measures avoid imposing arbitrary thresholds on our degreed measures. We believe these binary measures carve the relevant representational kinds in a principled and intuitively satisfying way. Nevertheless, those who disagree could still use our measures but adopt a different threshold marking what is required for a system to count as analog.

them (because the theory is not explicit enough to deliver a verdict). In contrast, our mirroring, purity, and mass measures yield plausible verdicts when applied to hybrid systems. First, *Colored Thermometer* has maximum mirroring, since the interpretation function is an isomorphism under the relations codified by the rules of the system (the *taller-than* and *warmer-than* relations). Second, *Colored Thermometer* has the same degree of mass as *Mercury Thermometer*, since both systems have the same number of structural profiles. But third, *Colored Thermometer* suffers with respect to purity because it has nominal profiles (colors and cloud coverages) while *Mercury Thermometer* does not. Consequently, our theory predicts that *Colored Thermometer* is less analog than *Mercury Thermometer*.

We believe these verdicts are correct, but there remains a further explanandum. *Colored Thermometer* seems to represent temperature analogically (via height) and cloud coverage symbolically (via color). Since our measures track degrees of analogicity for whole systems, our theory is not yet sensitive to the fact that some features of a system can represent analogically while other features represent symbolically. However, there is a natural way of extending our theory to accommodate this datum. To forecast: our theory will claim that the degree to which a feature represents analogically is determined by the degree of analogicity of the partial system generated by partitioning according to that feature.

We start by introducing the notion of a partial system. Informally, partial systems result from abstracting away from some of the original system's representational distinctions. For example, one partial system of *Colored Thermometer* groups columns by height and ignores color. A *partial system* of a representational system is generated by partitioning each of the syntactic and semantic spaces in a manner that retains the interpretation function. A *partition* of any set is a division of that set's elements into subsets—called *cells*—such that every element occurs in exactly one subset. A partition *retains* the interpretation function if and only if any two vehicles of the original syntactic space occupy the same cell in the partitioned syntactic space just in case their contents (in the original system) also occupy the same cell (in the partitioned semantic space). Partitioning and retainment ensure that the resulting partial system is a "partial version" of the original whole.

The elements of the partial system's syntactic and semantic spaces are *sets* of syntactic or semantic elements of the original system. For example, the set containing a blue, yellow, and red column of 10 cm height is one syntactic element of the "color ignoring" partial system of *Colored Thermometer* mentioned above. In order to apply our measures to the partial system, we need to define the relations that structure it. Let us say that two cells A and B of the partial system bear relation *r* to each other just in case every element in A bears *r* to every element in B. This definition entails that the set containing blue, yellow, and red columns of 10 cm height bears the *taller-than* relation to the set containing blue, yellow, and red columns of 8 cm height. In other words, the relations that structure the original system are "exported" in the straightforward way to the partial system. Importantly, we can apply our measures of analog mirroring, purity, and mass to partial systems so defined.

Every representational system has many partial systems. But the interesting partial systems are those that correspond to what we would intuitively call the *features* of the system. *Colored Thermometer* has two features: color and height. First, when the system is partitioned according to color (ignoring height), the result is a partial system with only three syntactic elements—blue, yellow, and red—and three semantic elements—clear skies, partly cloudy, and overcast. Because there are no rulebound relations between these elements, our theory predicts that this partial system has zero analogicity. Second, when the system is partitioned according to height (ignoring color), the result is a partial system that has as many syntactic elements as there are syntactically distinct heights. In this partial system, each syntactic cell contains three elements: a blue, a yellow, and a red column of identical height. When we apply our measures, the partial system receives maximum mirroring and maximum purity. In the terminology introduced earlier, the height partial system is perfectly analog and the color partial system perfectly symbolic. This entails that in *Colored Thermometer*, height represents analogically and color symbolically, which is exactly the verdict a theory should deliver.

This approach to partial systems generalizes in straightforward ways. Recall *City Populations*, the system we used to motivate the analog purity measure. This system has two features: letter character (representing city) and letter size (representing population). The size partial system has maximum purity (and is perfectly analog), since each of its cells corresponds to a distinct structural profile (though the system still has relatively low mass). The letter partial system is perfectly symbolic, since its cells do not stand in any rulebound relations to each other.

This way of thinking about partial systems allows the rulebound structure theory to answer not only the question "Which representational systems are analog and to what degree?" but also "Which features within a given system represent analogically and to what degree?" Both questions are important. Our theory's ability to answer both is evidence of its power and flexibility.

7 | OTHER THEORIES OF ANALOG REPRESENTATION

This section reviews how the rulebound structure theory relates to several prominent theories of analog representation.

7.1 | Density

According to *density theories*, a representational system is analog exactly when it is *syntactically* and *semantically dense*.[24] In brief, a syntactically dense system is one in which between any two syntactically well-formed vehicles there is a third, while a semantically dense system is one in which between any two contents there is a third. In the ensuing discussion, we will simply grant that the density theorist has a principled way of characterizing what it is for one element to lie "between" two other elements.

While we think there is an important grain of truth in density theories, we believe that they are extensionally inadequate. First, density theories are too exclusive. Many non-dense systems are analog. Consider again *Structural Population Map*, where 100 red shades represent 100 population densities and darker shades indicate higher population density. This system has only 100 red shades, meaning that there are no syntactically distinct shades in between (say) red_{34} and red_{35}; intermediate densities cannot be represented. Although this system is neither syntactically nor semantically dense, it is nevertheless intuitively analog. The rulebound structure theory counts *Structural Population Map* as having maximum mirroring and purity: the interpretation function is an isomorphism and every representational difference arises from a structural difference.

[24] A density theory could also rely on either syntactic density or semantic density alone. The canonical density theory, Goodman [1968], requires both, so we operate with that conception. The arguments we offer against such a theory generalize easily to other density accounts. See Schonbein [2014] for a syntactic density account.

Second, density accounts are too inclusive. They incorrectly count certain systems as analog. Consider *Scrambled Population Map*:

Scrambled Population Map—**Vehicles**: Infinitely many shades of red. **Contents**: Infinitely many population densities. **Interpretation**: Each shade of red is arbitrarily mapped to a different population density so that there is a non-structure-preserving bijection between shades of red and population densities.

Suppose that the set of red shades is dense. Because each shade represents a different population density, *Scrambled Population Map* is syntactically and semantically dense. But the mapping between shades and population densities is completely arbitrary; no relation between population density can be read off the relations between red shades. Therefore, density accounts incorrectly count *Scrambled Population Map* as analog. By contrast, the rulebound structure theory provides a diagnosis for why the system is not analog: in particular, the system fails to achieve even a low degree of analog mirroring.

Though we have argued that density accounts are extensionally inadequate, we think that density accounts still track an important truth. In order for a dense system to be interpretable, it must be analog in our sense. That is, dense systems must have interpretation functions that preserve structure. Imagine trying to interpret *Scrambled Population Map*. Perhaps you think you see red_{42}, but might that be $\text{red}_{42.1}$ or even $\text{red}_{42.01}$? Since the system's interpretation function does not preserve structure, there are no constraints on how the structure of the vehicles maps to the structure of the contents. It may be that red_{42} represents 670 people/km^2, $\text{red}_{42.1}$ represents 11 people/km^2, and $\text{red}_{42.01}$ represents 927 people/km^2. Since you cannot be sure whether you are seeing red_{42} or instead one of the infinitely many nearby shades, you could easily be radically wrong about the content. By contrast, if the interpretation function did preserve structure, then nearby shades of red would represent nearby population densities. In such a scenario, even if you were unsure of the exact identity of a vehicle, you could still be relatively confident of the content. These observations indicate that mirroring secures the interpretability of dense systems. As a result, dense systems will tend to be analog.[25]

7.2 | Magnitude Mirroring

According to *magnitude mirroring theories*, analog representation is a matter of magnitudes of the vehicle mapping to magnitudes of the content.[26] For example, Peacocke [2019: 52] writes that "analogue representation is representation of magnitudes, by magnitudes." Examples of magnitudes include length, mass, temperature, and duration. The defining feature of magnitudes is that

---

[25] This is similar to a feature that Kulvicki [2015] highlights: namely, that abstractions over syntax correspond to abstractions over semantic values. We agree with Kulvicki that density is a symptom (rather than the source) of analogicity. But we also think that Kulvicki's explanation (in terms of abstraction) can likewise be explained by a more fundamental theory of analog representation. See Kulvicki [2015: §2–3] for an insightful discussion of the connection between density, abstraction, and interpretation.

[26] Magnitude mirroring theories include Beck [2015, 2019], Block [forthcoming], Lewis [1971], Maley [2011], and Peacocke [2019].

they come in greater and lesser extents. For example, a temperature of 6°C is greater than 2°C but less than 7°C.

In our view, magnitude mirroring theories are on the right track but cast too narrow a net. To begin, observe that analog representation need not involve magnitudes. Consider hue: it does not make sense to say that x is greater than y with respect to hue, so hue is not a magnitude. But it is obvious that systems in which hue values of the vehicle map to hue values of the content (e.g. photographs) ought to count as analog. The magnitude theorist could respond that magnitudes need not be linearly ordered. To some extent, this is a verbal question about the term 'magnitude'. However, magnitudes are standardly defined in metaphysics and the philosophy of science as coming in lesser or greater extents. As examples, Helmholtz [1887: 17] says that "a magnitude is a quality of objects that is amenable to ordering from smaller to greater," Kyburg [1997: 381] writes that magnitudes are the values of "quantity function[s]" and "reflect the structure of the real [numbers]," and Lewis [2009: 206] defines magnitudes as "properties that admit of degree."

More importantly, even if one permits magnitudes to lack linear ordering, existing magnitude accounts do not explain how magnitude mirroring works in these cases. For example, Beck [2019: 334] says that analog representation involves "the representation of one magnitude by a second magnitude such that the second magnitude has the function of increasing or decreasing with the first." These accounts do not generalize in any straightforward way to non-linearly ordered structures, since those kinds of systems do not yield a meaningful notion of increasing or decreasing. By contrast, the rulebound structure theory of analog representation applies just as straightforwardly to systems without magnitudes as it does to systems with magnitudes.[27]

Finally, magnitude mirroring theories are usually all-or-nothing: either magnitudes represent magnitudes or not. However, we have presented a number of cases illustrating how analog representation comes in degrees. In *City Populations* the size magnitude represents the population magnitude. The magnitude mirroring theorist must say that this system is analog simpliciter. But *City Populations* (in which letters represent cities) does not seem as analog as, say, *Mercury Thermometer*. Our theory captures this difference by generating the result that while both systems have maximum mirroring, *City Populations* has relatively low purity and mass.

7.3 | Abstraction

According to *abstraction theories*, analog representation is a matter of abstractions over vehicles corresponding to abstractions over contents (Kulvicki 2015). For example, these theories contend that *Mercury Thermometer* is analog because abstractions over vehicles (e.g. mercury heights between 5cm and 6cm) correspond to abstractions over contents (e.g. temperatures between 25°C and 30°C).

Though we believe abstraction theories are on the right track, we also think they identify a symptom, rather than the source, of analog representation. A theory of analog representation ought to explain why analog systems enable abstractions; the fact that a system yields these abstractions is unlikely to be a brute feature of that system. Our view is that these abstractions are possible because of the nature of the interpretation functions of analog systems. We have argued

[27] Beck [2019] argues that Weber's Law is best explained by taking perception to involve magnitude mirroring, and that this result is evidence that perception is analog. We are sympathetic to Beck's arguments. Since the rulebound structure theory takes magnitude mirroring to be one way for a system to be analog, we agree that evidence of magnitude mirroring is evidence of analogicity.

that analog systems have interpretation functions that preserve the structure induced by their rulebound relations. For example, *Mercury Thermometer*'s interpretation function maps *taller-than* to *warmer-than*. Because of this, abstractions over height correspond to abstractions over temperature. In other words, our account identifies the underlying feature of analog systems that grounds the core component of abstraction theories.

There are also extensional differences between our theory and the abstraction theory. On the abstraction theory, a system is analog to the extent that it enables the relevant abstractions. On the rulebound structure theory, a system is analog to the extent that its interpretive rules map syntactic structural features onto semantic structural features. Rulebound structure preservation entails abstractability, but abstractability does not entail rulebound structure preservation. Consider again *Stipulated Population Map*, which has 100 red shades mapping to 100 population densities such that darker shades represent greater densities. For this system, abstractions over color correspond to abstractions over population. However, *Stipulated Population Map* is not analog on our account because its interpretation function does not operate on any rulebound relations. Instead, the interpretation function stipulatively maps each red shade, one by one, to a specific population density. Because of the lack of rulebound relations, the system has zero mirroring, and by consequence fails to satisfy our criteria for analog representation. This difference between the abstraction theory and the rulebound structure theory is due to how each theory deals with the interpretation function. For the abstraction theory, all that matters are the inputs and outputs of the interpretation function. Because of this, the abstraction theory is not sensitive to any intensional differences between interpretation functions. By contrast, the rulebound structure theory looks "inside" the interpretation function. As a result, it can distinguish between (for example) *Stipulated Population Map* and *Structural Population Map*, which map the same inputs to the same outputs but which use different rules.

For those with residual feelings that *Stipulated Population Map* is analog, there is a straightforward way of explaining away that intuition: *Stipulated Population Map* has a feature—namely, abstraction—that is often a symptom of genuine analog representation. Moreover, if one is presented with only the input-output structure of *Stipulated Population Map*, it is natural to assume that the system has the rules of *Structural Population Map*. However, we think that once one recognizes the difference between *Stipulated Population Map* and *Structural Population Map*, there is clearly a respect in which the former is symbolic and the latter analog.

7.4 | Surveying the Structural Approach

All of the theories we have discussed—density theories, magnitude mirroring theories, abstraction theories, as well as our own rulebound structure theory—fall under what we have called the "structural approach" to analog representation. The unifying theme of the approach is that analog representation is a matter of structure corresponding to structure.[28] Now that we have discussed

[28] Although theories of *structural representation* do not target analog representation, they also frequently appeal to the idea of structure corresponding to structure. For example, Shea [2014: 123) writes that structural representation occurs when "a relation between representational vehicles represents a relation between the entities they represent." (See also Ramsey [2007: 77-92], Shagrir [2012], and Swoyer [1991]). However, even setting aside the fact that they aim at different targets, structural representation theories remain importantly different from structural theories of analog representation. First, structural representation requires that the relevant relations serve as inputs to further processing. By contrast, the rulebound structure theory requires only that the relations occur in the rules of the interpretation function. Other structural

each of these theories individually, it is worth making some general remarks about the structural approach. Any structural theory should answer the following two questions:

Structure Question: What counts as *structure*?
Correspondence Question: What counts as structure *corresponding to* structure?

To answer the structure question, most theories within the structural approach appeal to a certain *kind* of structure: density theories appeal to dense structure, magnitude mirroring theories appeal to magnitudinal structure, and abstraction theories appeal to abstractability structure. We have argued that the mark of the analog does not lie with any particular kind of structure. Instead, we have adopted a liberal approach, where *any kind* of rulebound structure can be used to represent analogically. We provided a story of where this structure comes from: namely, from the rules of the interpretation function. On our picture, the structure of a representational system is determined by the relations codified by its interpretive rules.

To answer the correspondence question, most theories appeal to the idea of structure "mirroring" structure, where the notion of mirroring is either left intuitive or cashed out in terms of isomorphism. We have argued that simple isomorphism accounts fail because isomorphism is all-or-nothing while analogicity comes in degrees. In response, we have developed a precise degree-theoretic mirroring measure, which we believe captures an intuitive sense of mirroring.

Furthermore, we have argued for the surprising conclusion that even systems with perfect mirroring can be more or less analog. For example, the syntactic structure of *City Populations* maximally mirrors its semantic structure. Yet the system remains less than perfectly analog because many of its features represent symbolically. This claim holds not only for the technical analog mirroring measure we have developed, but also for pre-theoretical conceptions of mirroring. In light of this, we think that an adequate answer to the correspondence question should not only explain what is meant by "structure corresponding to structure," but also provide a way of measuring the proportion of and absolute amount of representational work done by structure corresponding to structure. This is exactly what our analog purity and analog mass measures do. To our knowledge, no other theory of analog representation has clearly identified these other aspects of analogicity.

It is worth noting that our answers to the structure and correspondence questions are somewhat separable. While we favor the view that any rulebound structure can contribute to analogicity, it is possible to reject our intensionalist framework but endorse our degree-theoretic measures. For example, one could accept that mirroring, purity, and mass provide good answers to the correspondence question yet still think that only magnitudinal structure matters for analog representation. Conversely, while we favor mirroring, purity, and mass as measures of what it is for structure to correspond to structure, it is possible to reject our measures yet endorse our intensionalist framework and our rulebound account of structure. For example, one could accept that any rulebound structure can contribute to analogicity yet still favor an isomorphism account of correspondence. These considerations show how the components of the rulebound structure theory are relatively modular. We believe the complete package yields an intuitive and satisfying theory of analog rep-

theories of analog representation do not require even that much. Second, analog representation and structural representation are extensionally different: mercury thermometers are paradigms of analog representation, but do not count as structural representations at all (since users do not normally interpret mercury thermometers by processing relations between pairs of mercury heights). Third, the vehicles of structural representation must be complex: "some of the proper parts of the structure that does the representing must themselves be representations" (Shea 2014: 135). By contrast, analog representations have no such complexity requirement.

resentation. But given the modularity of our view, those who diverge from us on some issues can still adopt other components of our theory.

8 | OBJECTIONS

We now consider objections.

8.1 | The Gruesome Structure Objection

Consider *Gruesome Thermometer*, which has the same vehicles and contents as *Mercury Thermometer*, but which maps vehicles to contents in a seemingly random way, with no correspondence between *taller-than* and *warmer-than*. Intuitively, *Gruesome Thermometer* is not analog. Yet *Gruesome Thermometer* can be understood in a way that forces the rulebound structure theory to give the system high mirroring, purity, and mass scores.

Suppose the interpretation function for *Gruesome Thermometer* operates on the *schwarmer-than* relation, where 92°C is *schwarmer-than* 70°C, which is *schwarmer-than* 80°C, and so on. The *schwarmer-than* relation is gerrymandered so as to generate the result that *Gruesome Thermometer* has maximum mirroring, maximum purity, and high mass. This makes *Gruesome Thermometer* come out just as analog, on all the measures, as the more familiar *Mercury Thermometer*.[29] More generally, the rulebound structure theory may seem to classify as analog any representational system whose interpretation function codifies and preserves gruesome structural relations, no matter how symbolic such a system seems. The objection is that this classification is incorrect.

We will provide two potential responses to the gruesome structure objection. Which response one favors will depend on prior metaphysical and metasemantic commitments. Since addressing those commitments in detail would take us too far astray, we will simply articulate the responses without taking a stance on which is better.

The first option adopts a *naturalness constraint*, which says that only natural rulebound relations count toward analogicity. The idea behind naturalness is that some relations (e.g. *warmer-than*) carve nature at its joints whereas others (e.g. *schwarmer-than*) do not. Though there is disagreement on the nature of naturalness, there is usually agreement on which properties are (and are not) natural. Under this first option, even if gruesome relations like *schwarmer-than* can be codified by interpretation functions, they do not make a system more analog. On such an approach, analog representation is not merely a matter of structure corresponding to structure, but rather a matter of natural structure corresponding to natural structure. This option yields the verdict that *Gruesome Thermometer* is not analog at all.[30]

[29] Readers might notice that *Gruesome Thermometer* is similar to our earlier example *Scrambled Thermometer*. These systems have the same input-output structure, but they differ in their rules: *Gruesome Thermometer* has gruesome rulebound relations while *Scrambled Thermometer* has no rulebound relations at all.

[30] See Lewis[1983] and Dorr and Hawthorne [2013] on what naturalness is. See Lewis [1984] and Williams [2015] for the use of naturalness as a metasemantic constraint. The idea that something like naturalness divides the analog from the symbolic (or digital) is endorsed by Lewis [1971], and is also a theme in work on the distinction that appeals to the notions of convention and natural correspondence (see Giardino and Greenberg [2015, §1.1]). On these approaches, analog (or iconic) representation involves natural correspondence, while symbolic representation is wholly conventional.

The second option accepts that *Gruesome Thermometer* is analog and appeals to a metasemantic error theory to explain why such a system strikes us as less analog than natural systems. The basic idea motivating this response is that *Gruesome Thermometer* is just as analog as *Mercury Thermometer*, but merely seems less analog to us because of the kinds of relations we humans happen to use to carve up reality. Even if any relations whatsoever can in principle be codified by interpretation functions to yield analog representational systems, it is plausible that only analog systems that map natural relations to natural relations are amenable to creation and interpretation by humans. If humans were to create a system whose inputs and outputs were identical to *Gruesome Thermometer*'s, such a system would use purely stipulative rules rather than systematic rules operating over gruesome relations. Because of this, there is a danger that one's intuitions about *Gruesome Thermometer* are really tracking features of *Stipulated Gruesome Thermometer*, a system that our theory would classify as symbolic. On the other hand, if there were gruesome aliens who perceived the world in terms of the *schwarmer-than* relation, then those aliens may well find *Gruesome Thermometer* more analog than *Mercury Thermometer*.

8.2 | The Small Systems Objection

Consider the following system:

Arrows—**Vehicles**: Arrows placed on trail signs. Arrows come in three sizes: Long, Medium, and Short. **Contents**: Distances. **Interpretation**: A long arrow represents "more than 1 km," a medium arrow "200-1000 m," and a short arrow "less than 200 m."

Arrows uses arrows placed on wooden signs along a walking path to indicate how far the walker is from the next landmark. Longer arrows represent greater distances. The rulebound structure theory gives *Arrows* maximum mirroring and purity (but low mass). However, a perfect score on two out of three measures may strike the reader as overly generous for a system that does not seem very analog at all. More generally, the small systems objection says that our theory incorrectly counts small systems as highly analog.

First, we think it is somewhat misleading to say that *Arrows* is highly analog. Though *Arrows* scores highly on two out of three measures, the fact that it scores extremely low on the third measure is significant for assessing intuitions about analogicity. In fact, we think that ordinary ascriptions of analogicity typically require relatively high scores on all three measures. Since *Arrows* has low mass, we would expect most people to judge *Arrows* to be not particularly analog.

Second, there is an alternative way of precisifying how the *Arrows* system works, and this alternative precisification may confound intuitions about analogicity. On the *stipulative* precisification, *Arrows*' interpretation function assigns the contents "more than 1 km," "200-1000 m," and "less than 200 m" to the long, medium, and short arrows directly. On the *structural* precisification, *Arrows*' interpretation function relies on a *longer-than* relation between the arrow symbols. Only on the second precisification does *Arrows* count as having maximum mirroring and purity.

In fact, if one were told only the inputs and outputs of *Arrows*, the stipulative interpretation would be the most natural way to understand how *Arrows* works. If tasked with assigning the three arrows to their contents, most would take the easy route and simply stipulate which arrow has which content. The alternative method of inventing rules that relate length relations between arrows to distance relations between paths is more convoluted for systems as simple as *Arrows*. Imagine going on a hike and encountering the trail signs of the *Arrows* system. Typically, such trail signs include a key that captures the system's interpretation function. The key is overwhelmingly likely to simply list each vehicle next to its content. This is evidence that small systems like *Arrows* will, in general, tend to be symbolic, and not use interpretative rules that rely on structure.

8.3 | The Verbal Dispute Objection

What exactly turns on which theory of analog representation is true? To some readers, it may seem that there are many ways of precisifying the notion of analog representation and that nothing significant turns on which notion we adopt. Perhaps the rulebound structure theory targets one representational kind, magnitude mirroring theories target another, density accounts yet another, and so on.

While we agree that some disputes about analog representation may be partly verbal, we also think that theorists of analog representation have a common target. After all, no philosophers of analog representation take their claims to be stipulatively true. If developing a theory of analog representation were a matter of stipulation, then it would make little sense to appeal to intuitions about cases. Instead, philosophers develop theories that purport to identify what they take to be an interesting representational kind and to explain the nature of that kind. It is easy to see that not all representational kinds are equally interesting: the category "representations that are accurate on Tuesdays" is far less interesting than the category "representations that involve magnitude mirroring," which (in our view) is less interesting still than the category "representations whose interpretation functions map rulebound structure to rulebound structure." In appealing to intuitions about cases, we have sought to evoke intuitions about joints between representational kinds rather than intuitions about the extension of the English word 'analog'. Our theorizing is thus better thought of as conceptually engineering the notion of the analog, rather than as trying to capture its ordinary usage.

Moreover, we think the rulebound structure theory does a better job of capturing the target representational kind than other theories. In some cases, existing theories of analog representation classify as symbolic systems that seem fundamentally similar to the systems they classify as analog. For example, density theories classify dense mercury thermometers as analog while classifying *Mercury Thermometer*, which has only 100 contents, as symbolic. In other cases, existing theories leave open how to apply their account to challenge cases. For example, it is unobvious how to make magnitude theories work for dimensions that lack ordinal structure (such as hue). Moreover, as far as we know, no prior discussion of analog representation explains how to deal with the lucky scrambled systems and hybrid systems discussed earlier. By contrast, the rulebound structure theory provides intuitive and satisfying verdicts on all of these issues. For these reasons, we believe the rulebound structure theory to be a genuine improvement over extant theories.

Further evidence for the substantivity of these disputes comes from the observation that which of these theories is correct will have downstream implications for formal semantics. There is a budding field concerned with the project of developing formal semantics for non-symbolic, non-language-like representational systems.[31] Providing a formal semantics requires providing an interpretation function. The rulebound structure theory identifies what form a formal semantics for analog systems should take. In particular, the formal semantic apparatus should consist of a set of rules that map syntactic relations between vehicles to semantic relations between contents.

Finally, a virtue of the rulebound structure theory is that it explains some of the costs and benefits of using analog versus symbolic systems. We turn to this issue next.

9 | FECUNDITY AND FLEXIBILITY

The rulebound structure theory explains the fact that analog systems are more fecund but less flexible than their symbolic counterparts. A representational system is *semantically fecund* to the extent that it has high expressive power relative to the complexity of its interpretation function. In other words, fecund systems generate greater expressive bang for their interpretive buck. To illustrate, contrast *Structural Population Map*, which uses 100 vehicles (shades of red) to represent 100 contents (population densities) with *Inkblots*:

Inkblots—**Vehicles**: 100 random inkblots. **Contents**: 100 types of fruit. **Interpretation**: 100 rules, each of which maps a particular inkblot to a particular fruit.

Both systems map the same number of vehicles to the same number of contents. However, while *Structural Population Map* achieves this mapping largely through a single simple rule (darker shades represent denser populations), *Inkblots* has as many rules as it has vehicles, mapping each vehicle to its content one by one. This gives *Structural Population Map* greater fecundity than *Inkblots*. Although both systems map the same number of vehicles to the same number of contents, *Structural Population Map* does so more efficiently.

Let us say that two systems are *counterparts* just in case exactly the same contents are expressible in each. Analog systems are more fecund than their symbolic counterparts.[32] This is evident when we observe that analog representations exhibit two of the main symptoms of fecundity: they are easy to learn and easy to create. To learn *Inkblots*, you must memorize the content of each inkblot individually. Knowing the content of one inkblot gives you no purchase on the content of any other. By contrast, learning *Structural Population Map* is comparatively easy. Once you learn the contents of a few colors and that darker shades of greater represent greater population densities, you have basically learned the system. Furthermore, creating an analog system such as *Structural Population* is easier than creating a symbolic system such as *Inkblots*.[33] If you were tasked with assigning 100 inkblots to 100 fruits while your philosophical nemesis was tasked with assigning 100 shades of red to 100 population densities, your nemesis would finish first every time.

---

[31] See Casati and Varzi [1999], Cohn [2014], Cumming, Greenberg, and Kelly [2017], and Greenberg [2018] for some examples.

[32] We suspect that the restriction to counterpart systems is unnecessary for systems without complex representations involving multiple parts.

[33] In small systems with very few vehicles and contents, it may sometimes be easier to stipulate contents one-by-one rather than lay down rules capable of handling hundreds of contents. *Arrows* from §8.2 illustrates this situation.

The rulebound structure theory explains why analog systems are fecund: namely, because their interpretive rules take advantage of the structure of their vehicles and contents. Analog systems use structure as a resource to assign many contents in one fell swoop. A single simple rule (e.g. "darker shades represent greater densities") can do a large amount of interpretative work, generating many, perhaps infinitely many, contents. By contrast, the interpretation functions of symbolic systems become increasingly complex as more vehicles and contents are added. The core reason for the ease of learning and ease of creation of analog systems is fundamentally the same: shades of red and populations have analogous structures that are easy to exploit for representational purposes. The learners and creators of an analog system can take advantage of this pre-existing structure; the metasemantic scaffolding is already in place.

If analog systems have all of these benefits, why use symbolic systems at all? In spite of their virtues, analog systems have a significant flaw: they are less *semantically flexible* than symbolic systems. This lack of flexibility manifests in two ways. First, analog systems are less *modular* than symbolic systems. Symbolic systems tend to use separate rules to map each vehicle to its content, whereas the rules of analog systems map many vehicles to contents all at once. This makes the alteration and addition of content assignments in analog systems more difficult than in symbolic systems. In *Inkblots*, any particular content assignment can be modified without changing the rest of the system, and it is simple to add a new inkblot to represent kumquat. By contrast, in *Structural Population Map*, it is difficult to change a particular assignment of color to population density without massively altering many other assignments, and the most straightforward way of extending the system would be to simply tack an extra symbolic rule onto the analog system. Second, analog systems are less *general* than symbolic systems, in that analog systems are limited in what we can use them to represent. For analog representation to be possible, the representational target must have some exploitable structure in the first place, and some vehicles with analogous structure must be available to do the representing. Symbolic representation imposes no such requirements.[34]

Fecundity and flexibility provide insight into where we should or should not expect to encounter analog representation. Even when appropriately structured subject matter and vehicles are available, there may still be advantages to constructing a symbolic system. If one wants a flexible system with modular rules, a symbolic system is preferable. On the other hand, if one wants a system that is easy to create and learn, then analog representation will likely work best. These considerations apply not only to whole systems, but also to particular features of systems. Imagine a scenario in which representation along one parameter (such as spatial location) has exploitable structure, and where the designers of the system do not expect their representational needs to change. That parameter would be best represented analogically. But suppose there is also a second parameter (such as type of building) that does not have easily exploitable structure, or for which the designers expect there to be future modifications of the system (such as new types of buildings). That second parameter is best represented symbolically.

10 | CONCLUSION

We began this paper with the question, "What is the mark of the analog?" We explained where to look for that mark—in interpretation functions—as well as what to look for—rules that map structure to structure.

[34] See Camp [2007] for a related discussion of the trade-offs between different types of representational systems.

On our view, interpretation functions should be understood as sets of rules that map vehicles to contents, and not simply as vehicle-content pairs. A system's structure is provided by the relations codified in the rules of its interpretation function. Analog systems are those marked by the use of interpretive rules that map syntactic structural features onto semantic structural features. Since systems can do better or worse in mapping structure onto structure, analogicity comes in degrees. Furthermore, systems can do better or worse in at least three different ways, corresponding to the measures we have called mirroring, purity, and mass. Mirroring measures the extent to which the system's structure is preserved by its interpretation function; purity measures the proportion of representational distinctions that are due to structure; and mass measures the amount of representational distinctions due to structure.

We have illustrated how the rulebound structure theory applies in a range of cases and compares to its main rivals. We think our theory does better on a variety of counts, including giving satisfying verdicts on systems that alternative theories misclassify. The rulebound structure theory's three degree-theoretic measures allow considerable flexibility in accounting for which systems are (or are not) analog and to what extent. To our knowledge, these fine-grained comparisons between representational systems, including hybrid systems that are partially analog and partially symbolic, are not captured by any extant theory of analog representation. We have also responded to what we take to be the main objections to the rulebound structure theory and used the theory to explain why analog systems are fecund but inflexible.

Analog systems are those that take advantage of a very basic representational resource: namely, the system's structure. That structure is provided by the rules of its interpretation function. We have explained how these simple ideas yield a systematic theory with greater predictive accuracy and explanatory scope than its rivals. As a result, we believe the rulebound structure theory provides a better—and not merely different—theory of analog representation.[35]

ORCID

Andrew Y. Lee https://orcid.org/0000-0002-1899-0511
Joshua Myers https://orcid.org/0000-0002-2803-1939
Gabriel Oak Rabin https://orcid.org/0000-0003-0239-5328

REFERENCES

Barwise, J., & Hammer, E. (1996). Diagrams and the concept of logical system. In G. Allwein & J. Barwise (Eds.), *Logical reasoning with diagrams*. Oxford University Press.

Beck, J. (2015). Analogue magnitude representations: A philosophical introduction. *British Journal for the Philosophy of Science*, *66* (4), 829–855.

Beck, J. (2019). Perception is analog: The argument from weber's law. *Journal of Philosophy*, *116* (6), 319–349.

Blachowicz, J. (1997). Analog representation beyond mental imagery. *Journal of Philosophy*, *94* (2), 55–84.

Block, N. (forthcoming). *The border between seeing and thinking*. MIT Press.

Burge, T. (2018). Iconic representation: Maps, pictures, and perception. In W. Shyam & F. A. Dorio (Eds.), *The map and the territory: Exploring the foundations of science, thought and reality*. pp. 79–100. Springer.

Camp, E. (2007). Thinking with maps. *Philosophical Perspectives*, *21* (1), 145–182.

Casati, R., & Varzi, A. (1999). *Parts and places: The structure of spatial representation*. MIT Press.

Cohn, N. (2014). *The visual language of comics*. Bloomsbury.

[35] ACKNOWLEDGMENTS: We are grateful for helpful comments from Nikola Andonovski, Gabriel Greenberg, John Kulvicki, Tamar Weber, Assaf Weksler, two anonymous referees for *Noûs*, and audiences at Dartmouth College, Tel-Hai College, UAE University Al Ain, the University of Grenoble Centre for Philosophy of Memory, and the 2019 meeting of the Anglo-German Picture Theory Group.

Couch, L. (2013). *Digital and analog communication systems* (8th ed.). Pearson.
Cumming, S., Greenberg, G., & Kelly, R. (2017). Conventions of viewpoint coherence in film. *Philosopher's Imprint*, *17*.
Dorr, C., & Hawthorne, J. (2013). Naturalness. In K. Bennett & D. Zimmerman (Eds.), *Oxford studies in metaphysics: Volume 8* (pp. 1). Oxford University Press.
Dretske, F. I. (1981). *Knowledge and the flow of information*. MIT Press.
Euler, L. (1748). *Introductio in analysin infinitorum*.
Feigenson, L., Carey, S., & Hauser, M. D. (2002). The representations underlying infants' choice of more: object files vs. analog magnitudes. *Psychological Science*, *13*, 150–156.
Frigerio, A., Giordani, A., & Mari, L. (2013). On Representing information: A characterization of the analog/digital distinction. *Dialectica*, *67* (4), 455–483.
Giardino and Greenberg (2015). Introduction: Varieties of iconicity. *Review of Philosophy and Psychology*, *6* (1), 1–25.
Goodman, N. (1968). *Languages of art*. Bobbs-Merrill.
Greenberg, G. (2018). Content and target in pictorial representation. *Ergo*, *5*.
Haugeland, J. (1981). Analog and analog. *Philosophical Topics*, *12*, 213–225.
Haugeland, J. (1991). "Representational Genera" from *Having Thought: Essays in the Metaphysics of Mind*, ed. Lawrence Erlbaum.
Helmholtz, H. (1887). *Counting and measuring*, C.L. Bryan (trans.), New Jersey: D. Van Nostrand, 1930.
Kleiner, I. (2009). Evolution of the function concept: A brief survey. In M. Anderson, V. Katz, R. Wilson (Eds.), *Who gave you the epsilon?: And other tales of mathematical history* (pp. 14–26). MAA.
Kulvicki, John (2015). Analog Representation and the Parts Principle. *Review of Philosophy and Psychology*, *64* (1), 165–180.
Kyburg, H. (1997). Quantities, magnitudes, and numbers. *Philosophy of Science*, *64* (3), 377–410.
Lewis, D. (1971). Analog and digital. *Nous*, *5*, 321–327.
Lewis, D. (1983). New work for a theory of universals. *Australasian Journal of Philosophy*, *61* (4), 343–377.
Lewis, D. (1984). Putnam's paradox. *Australasian Journal of Philosophy*, *62* (3), 221–236.
Lewis, D. (2009). Ramseyan humility. In D. Braddon-Mitchell & R. Nola (Eds.), *Conceptual analysis and philosophical naturalism* (pp. 203–222). Cambridge, MA: MIT Press.
Lopes, D. M. (2016). In the eye of the beholder. In J. Dodd (Ed.), *Art, mind, and narrative: Themes from the work of peter goldie* (pp. 223–340). Oxford.
Maley, C. J. (2011). Analog and digital, continuous and discrete. *Philosophical Studies*, *155* (1), 117–131.
Papayannopoulos, P. (2020). Computing and modelling: Analog vs. Analogue. *Studies in History and Philosophy of Science Part A*, *83*, 103–120.
Peacocke, C. (1986). The inaugural address: Analogue content. *Aristotelian Society Supplementary Volume*, *60* (1), 1–17.
Peacocke, C. (2019). *The primacy of metaphysics*. Oxford University Press.
Quilty-Dunn, J. (2016). Iconicity and the format of perception. *Journal of Consciousness Studies*, *23* (3-4), 255–263.
Quilty-Dunn, J. (2017). Syntax and Semantics of Perceptual Representation. Ph.D. dissertation, The Graduate Center, City University of New York.
Quilty-Dunn, J. (2019). Perceptual pluralism. *Noûs*.
Schonbein, W. (2014). Varieties of analog and digital representation. *Minds and Machines*, *24*(4), 415–438.
Schönfinkel, M. (1924). *On the building blocks of mathematical logic*. Reprinted in van Heijenoort, J. [1967]. (1976). *From Frege to Godel: A Source Book in Mathematical Logic, 1879–1931* (3rd ed.). Harvard University Press.
Shagrir, O. (2010). Brains as analog-model computers. *Studies in History and Philosophy of Science Part A*, *41* (3), 271–279.
Shagrir, O. (2012). Structural representations and the brain. *British Journal for the Philosophy of Science*, *63* (3), 519–545.
Shea, N. (2014). Exploitable isomorphism and structural representation. *Proceedings of the Aristotelian Society*, *114* (2), 123–144.
Shepard, R. N. (1978). The mental image. *American Psychologist*, 33, 125–137.
Siegelmann, H. T., & Fishman, S. (1998). Analog computation with dynamical systems. *Physica D: Non- linear Phenomena*, *120*(1–2), 214–235.

Sterrett, S. G. (2017). Experimentation on analogue models. In L. Magnani & T. Bertolotti (Eds.), *Springer handbook of model-based science* (pp. 857–878). Springer International Publishing.
Swoyer, C. (1991). Structural representation and surrogative reasoning. *Synthese*, *87* (3), 449–508.
Tao, T. (2011). *An introduction to measure theory Vol. 126*. Amer Mathematical Society.
Thomson-Jones, K. (2015). Philosophy of digital art. *Stanford Encyclopedia of Philosophy*.
Trenholme, R. (1994). Analog simulation. *Philosophy of Science*, *61* (1), 115–131.
Vergis, A., Steiglitz, K., & Dickinson, B. (1986). The complexity of analog computation, Mathematics and Computers in Simulation, *28*(2), 91–113.
Williams, J. R. G. (2015). Lewis on reference and eligibility. In B. Loewer & J. Schaffer (Eds.), *A companion to david lewis* (pp. 367–382). Wiley-Blackwell.
Wittgenstein, L. (2009). *Philosophical investigations (PI)* (4th ed.). P.M.S. Hacker and J. Schulte (eds. and trans.), Oxford: Wiley-Blackwell.
Xue, Y. (2016). Recent development in analog computation: A brief overview. *Analog Integrated Circuits and Signal Processing*, *86*(2), 181–187.
Youschkevitch, A. P. (1976). The concept of function up to the middle of the 19th century. *Archive for History of Exact Sciences*, *16*(1), 37–85.

How to cite this article: Lee, A. Y., Myers, J., & Rabin, G. The Structure of Analog Representation. *Noûs*. 2022;1–28. https://doi.org/10.1111/nous.12404